Costly Information Processing and the No-Trade Theorem

Sebastian Monroy Taborda[a]

[a] *Universidad ORT Uruguay, Montevideo, Uruguay*

Abstract

This paper revisits the no-trade theorem by explicitly modeling a margin that is usually kept implicit: agents must be able to process the informational content of what they observe without cost. The paper formalizes the argument in a minimal bilateral risky-asset environment with a two-stage structure. The buyer's problem reduces to a one-dimensional maximization of an indirect certainty-equivalent objective; the interior optimum equates the marginal processing cost with the marginal value of precision. The key mechanism is the role of equilibrium objects (i.e., prices) in aligning beliefs under rational expectations. We show that swamping's cognitive prerequisite — costless conditioning on the public object — fails whenever the processing intensity is finite, and that this failure is endogenous: agents optimally choose the degree of failure by equating the marginal value and marginal cost of precision; this failure has equilibrium consequences.

Keywords: Rational inattention, no-trade theorem, speculative trade, information processing, belief disagreement, sufficient statistics, processing costs
JEL: D82, D83, G14

---

1. Introduction

A central benchmark in information economics is the no-trade theorem of Milgrom and Stokey (1982). If agents share common or concordant priors, start from an ex-ante Pareto-efficient allocation, and the feasibility and mutual desirability of trade are common knowledge, then speculative trade cannot be strictly mutually beneficial. The theorem is powerful because it converts a simple epistemic argument into a sharp market prediction: once information is properly incorporated into beliefs, disagreement-based trade disappears.

This paper identifies a cognitive assumption that the no-trade logic takes for granted and that has received little formal attention: agents must be able to process the informational content of what they observe without cost. Prices, disclosures, and counterpart communications may be informative in principle, but their informational content is not automatically internalized. Interpreting such objects may require attention, computation, or analytical effort. When processing is costly, rational agents choose finite precision. The key implication is that posterior beliefs need not fully align even when agents share common priors and update Bayesianly on what they actually process.

The mechanism unfolds in four stages. First, an observable object (i.e., a price, a message, a disclosure) embeds information that is relevant for payoffs. Second, an agent selects the precision with which to process this object, balancing the certainty-equivalent benefit of higher precision against a convex cost of information processing. Third, because the optimal level of processing is finite, the resulting effective signal that enters posterior updating remains noisy. Fourth, since agents condition on different effective

---

Email address: `monroy@ort.edu.uy` (Sebastian Monroy Taborda)

statistics, their posteriors do not fully align, and trade occurs with positive probability. The paper does not refute the no-trade theorem; instead, it delineates a condition on the processing margin under which the posterior-alignment step underlying the no-trade result may fail to hold.

We formalize the argument in a minimal bilateral risky-asset environment with a two-stage structure. In Stage 1, the buyer commits to a processing intensity $\kappa \geq 0$, anticipating its effect on the Stage 2 clearing price. In Stage 2, markets clear bilaterally, given the agents' information sets. The seller processes her information frictionlessly; the buyer conditions on $\mathcal{I}_B(\kappa) = \sigma(s_B, \tilde{m}_B)$, where $\tilde{m}_B = m_{S\to B} + \xi_B$ and $\xi_B \sim \mathcal{N}(0, \kappa^{-1})$ are endogenous processing noise. This reduced-form representation implements the linear-quadratic rational-inattention framework of Sims (2003): in Gaussian environments with convex information costs, optimal behavior is observationally equivalent to conditioning on the true signal plus independent Gaussian noise, whose variance is pinned down by the cost structure.[1]

Under CARA-normality, bilateral clearing yields closed-form expressions for the transaction price and quantity. The traded quantity equals the posterior-mean gap divided by aggregate risk-bearing capacity. The buyer's Stage 1 problem reduces to a one-dimensional maximization of an indirect certainty-equivalent objective; the interior optimum equates the marginal processing cost with the marginal value of precision.

The paper's analytical contribution is organized around a three-step contrast. In the *frictionless-processing benchmark* ($\kappa, \tau \to \infty$), agents extract the full informational content of equilibrium objects, posteriors align almost surely, and speculative trade is eliminated—recovering the Milgrom and Stokey prediction. The *costly-processing economy* breaks this alignment: when processing is expensive, the buyer optimally adopts finite precision $\kappa^* < \infty$, conditions on a noisy processed statistic, and forms a posterior that generically differs from the seller's. Trade occurs with positive probability—not because priors disagree, but because posteriors do. The *comparative static* connects costs to volume: a higher processing cost induces lower equilibrium precision, widens the posterior-mean gap, and, under the stated risk-capacity condition, raises expected trade volume. This last step gives the mechanism empirical content—it predicts that reducing the interpretive complexity of public disclosures should compress trading volume, even without changing the underlying information.

Three results anchor the paper. First, Lemma 2 shows that under finite processing precision and non-degenerate information sets, the posterior-mean gap $\mu_B(\kappa) - \mu_S$ is a non-degenerate Gaussian random variable, so a nonzero desired trade quantity arises almost surely. Second, Lemma 3 establishes that the frictionless limit $\kappa \to \infty$ does not in general recover no-trade: processing noise and message noise are distinct frictions, and eliminating the former leaves the latter intact. No-trade requires the joint limit $(\kappa, \tau) \to (\infty, \infty)$. Third, Proposition 1 characterizes the two-stage equilibrium and shows that the processing cost $C(\kappa)$ governs the equilibrium magnitude of disagreement via κ^*, while Proposition 2 makes the comparative static precise: higher processing costs shift the equilibrium toward more disagreement and, when the denominator effect is dominated, greater expected trade volume.

Table 1 summarizes the contribution relative to three benchmarks. Relative to Milgrom and Stokey (1982), the paper keeps common priors and Bayesian updating but relaxes frictionless processing: agents condition on processed statistics that are strictly less informative than the raw public object, so Theorem 3 of MS loses its cognitive foundation. Relative to Grossman and Stiglitz (1980), the friction is interpretation, not acquisition: agents may observe an informative object but process it only imperfectly. Relative to Kyle (1985), trade does not require strategic order submission or noise trading. The primitive friction is the cost of internalizing information already in the environment.

---

[1] See Sims (2003), Section 3. The richer mutual-information formulation of Matějka (2016) — with general priors and discrete optimal actions — shares the same economic force but delivers a qualitatively different equilibrium object (a discrete price distribution). The Gaussian implementation is adopted here for analytical tractability.

Table 1: Positioning of the mechanism

Framework	Main friction	Source of trade or no-trade	Relation to this paper
Milgrom-Stokey	Common priors; frictionless conditioning	Posterior alignment eliminates speculative trade	Processing costs prevent alignment
Grossman-Stiglitz	Costly information acquisition	Full revelation removes acquisition incentives	Friction is interpretation, not acquisition
Kyle	Strategic trading and noise order flow	Partial revelation through market microstructure	No strategic order-flow channel
This paper	Costly information processing (Sims 2003)	Finite precision sustains posterior disagreement	Processing is the primitive margin

Notes: The table is intended to clarify the mechanism, not to classify the full literatures associated with each benchmark.

The counterpart message in the model is a reduced-form stand-in for any observable market object (i.e., price change, quote, disclosure, and counterparty communication) that contains payoff-relevant information but must be interpreted. The modeling choice is deliberately parsimonious: it isolates the processing channel from a full rational-expectations price-formation mechanism. The relevant question for no-trade logic is whether agents can costlessly condition on the informational content of observed objects. With processing costs, they cannot; they condition on processed statistics, and those statistics are generically insufficient.

The processing cost channel generates a prediction that differentiates it from heterogeneous-priors and noise-trading alternatives: simplifying the format or presentation of a public disclosure—without changing its informational content—should reduce trading volume by lowering effective processing costs and narrowing posterior disagreement. This comparative static does not follow from models where trade depends on the distribution of information rather than its interpretive complexity.

Section 2 reviews the Aumann (1976) epistemic foundations. Section 3 discusses the Milgrom–Stokey theorem, the swamping result, and the implicit processing assumption. Section 4 introduces the two-stage trading environment and derives the main equilibrium results. Section 5 connects the mechanism back to no-trade and processing-stage sufficiency failure. Section 6 discusses scope, limitations, and empirical implications. Section 7 concludes.

2. Disagreement, Common Knowledge and “Agreeing to Disagree”

This section builds on the theory underlying Aumann (1976)’s "agreeing to disagree" result. Consider a finite state space Ω with a common prior p, and agents who observe private information. For each agent i, let $\mathcal{I}_i(\omega)$ denote the information set in state ω (i.e., the information partition). Then, an event E is common knowledge in state ω if it holds throughout the set of all states consistent with everyone’s information and everyone’s knowledge of others’ information (Aumann, 1976; Samuelson, 2004).

With all this, one can present the Aumann (1976) theorem of impossibility to agree to disagree:

Theorem 1 (Theorem A in Aumann (1976)). *If two people have the same priors, and their posteriors for an event E are common knowledge, then these posteriors are equal.*

The economic interpretation is that disagreement cannot survive when beliefs become effectively shared. This benchmark is central to the no-trade reasoning: in rational-expectations environments, equilibrium objects (i.e., prices) serve as public statistics that encode others’ information, thus promoting posterior alignment.

3. Milgrom—Stokey and the Informational Role of Prices

Milgrom and Stokey (1982) provide a benchmark no-trade result: starting from an ex-ante Pareto-optimal allocation under common (or concordant) priors, speculative trade cannot be strictly mutually beneficial once feasibility and mutual acceptability are common knowledge.[2]

For our purposes, the key mechanism is the role of equilibrium objects (i.e., prices) in aligning beliefs under rational expectations. Under Milgrom and Stokey (1982), reopening markets after private information arrives can yield fully revealing equilibrium prices. That is, the price vector can act as a sufficient statistic for the payoff-relevant uncertainty: conditioning on the equilibrium price vector can reproduce the same posterior that would result from conditioning on the pooled information of all agents. Once this alignment of posteriors is achieved, speculative trading motives disappear.

A related implication is a swamping property: when prices are fully revealing, additional private signals become redundant conditional on prices. This is precisely the step that motivates our wedge. The logic in Milgrom and Stokey (1982) implicitly assumes that agents can extract and condition without cost on the full informational content of the observed equilibrium object. If processing is costly and optimally limited, conditioning on a processed statistic need not replicate the frictionless posterior, so sufficiency, swamping, and posterior alignment can fail.

We next introduce costly information processing (i.e., motivated by rational inattention) as a wedge that prevents observed objects from functioning as sufficient statistics in Sections 4.2 and 4.3. Section 4 integrates it into a minimally risky-asset trading environment.

4. An Illustrative Trading Environment with Costly Information Processing

This section centers the discussion on a canonical risky-asset trading problem that: (i) makes trade the object of interest, (ii) preserves the informational logic behind the no-trade arguments, and (iii) introduces friction at the processing stage rather than at the acquisition stage.

The model is deliberately minimal. There is one risky asset, two CARA agents, and Gaussian signals. The seller processes information frictionlessly, while the buyer chooses how precisely to process a noisy message about the seller's signal. The section proceeds in four steps: first, it describes the environment and information structure; second, it introduces the endogenous processing technology; third, it derives demands and market clearing conditional on processing; and fourth, it characterizes the buyer's optimal processing intensity.

Throughout the model, κ denotes the buyer's processing precision, with higher values corresponding to a finer interpretation of the message. We use x_i^* for individual optimal demands and x^* for the market-clearing traded quantity. The notation CE denotes the buyer's certainty-equivalent trading gain.

4.1. Environment

There is a single risky asset with a terminal payoff

$$v \sim \mathcal{N}(\mu_0, \sigma_v^2) \tag{1}$$

and a risk-free numeraire.

Two agents trade bilaterally: a seller S initially holds $q_0 > 0$ units of the risky asset, while the buyer B initially holds zero units. The agents' preferences over terminal wealth are represented by a CARA function given by:

$$U_i(W) = -\exp(-\gamma_i W) \quad \forall i \in \{S, B\} \tag{2}$$

with constant absolute risk aversion given by $\gamma_i > 0$.

---

[2] We refer readers to Milgrom and Stokey (1982) for formal statements and proofs. Our purpose here is to isolate the mechanism that links no-trade to posterior alignment.

Each agent i observes private signals about v in the form of

$$s_i = v + \varepsilon_i \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_i^2) \tag{3}$$

where $\varepsilon_S, \varepsilon_B$ are independent of each other and of v. In addition, each side receives a noisy counterpart message, specifically:

$$m_{i \to j} = s_i + \eta_i \quad \eta_i \sim \mathcal{N}(0, \tau_i^{-1}), i \neq j \tag{4}$$

where η_S, η_B are independent of all primitives. This message is interpreted broadly: it can represent cheap talk contaminated by noise, imperfect disclosure, or any frictional communication about counterpart information. The trade occurs at a price p in the numeraire. Let x denote the net demand of the buyer for the risky asset, such that the seller's net position changes by $-x$. Then, the feasibility constraint is $x \in [-q_0, q_0]$.

The timing is as follows. Nature draws the asset payoff. Each agent observes a private signal. Agents exchange noisy messages about their signals. The buyer chooses a processing intensity for the seller's message. Given beliefs and processing, agents trade at a clearing price. Finally, the asset payoff is realized.

4.2. Costly processing as endogenous precision

The key departure from the frictionless benchmark is that the buyer does not fully process the seller's message. The model assigns an explicit processing decision only to the buyer, with the seller assumed to process her information frictionlessly ($\kappa_S = \infty$). This asymmetry is a deliberate proof-of-concept choice rather than a claim about how processing costs are distributed across market participants. It isolates the processing channel cleanly: by making one agent's effective information set a function of κ while holding the other's fixed, the model admits a one-dimensional Stage-1 optimization problem that yields the sharp comparative statics in Proposition 2 without additional notation.

The asymmetry does not drive the main qualitative result. Posterior disagreement arises because the two agents condition on different effective statistics at the time of trading—not because one processes more precisely than the other. In a two-sided extension with precisions (κ_B, κ_S), both agents would condition on noisy processed signals $\tilde{m}_B = m_{S \to B} + \xi_B$ and $\tilde{m}_S = m_{B \to S} + \xi_S$, and the posterior-mean gap $\mu_B(\kappa_B) - \mu_S(\kappa_S)$ would remain non-degenerate for any finite pair (κ_B, κ_S): both processing-noise terms contribute independently to disagreement, and the no-trade benchmark still requires the joint limit $(\kappa_B, \kappa_S, \tau) \to (\infty, \infty, \infty)$. The one-sided model is the limiting case $\kappa_S \to \infty$ and serves as the simplest environment in which the mechanism operates. Section 6.4 discusses the two-sided extension further.

The buyer chooses an attention or processing level $\kappa \geq 0$ that determines the precision with which $m_{S \to B}$ is transformed into an effective processed signal. Higher κ means that the buyer devotes more attention to interpreting the message, thereby reducing processing noise. We model this as adding endogenous Gaussian noise:

$$\tilde{m}_B = m_{S \to B} + \xi_B \quad \xi_B \sim \mathcal{N}(0, \kappa^{-1}) \tag{5}$$

independent of all primitives. Processing is expensive.

Assumption 1 (Processing costs). *The processing cost function* $C : \mathbb{R}_+ \to \mathbb{R}_+$ *is continuous, strictly increasing, and convex, with* $C(0) = 0$ *and* $\lim_{\kappa \to \infty} C(\kappa) = \infty$.

Let $\mathcal{I}_B(\kappa) = \sigma(s_B, \tilde{m}_B)$ denote the buyer's information set after processing. By normality, the buyer's posterior is Gaussian in the form of

$$v \mid \mathcal{I}_B(\kappa) \sim \mathcal{N}(\mu_B(\kappa), \sigma_B^2(\kappa)) \tag{6}$$

where $\sigma_B^2(\kappa)$ is the conditional variance, which is strictly decreasing in κ whenever $\tau_B > 0$.

4.3. Rational inattention foundations: Sims (2003)

The endogenous-precision technology introduced above has a direct antecedent in the rational-inattention literature. The key reference is Sims (2003).

Sims (2003) proposes that agents process information through a communication channel with finite Shannon capacity and shows that the optimal strategy can be characterized in terms of the mutual information between the state and the action. In the linear-quadratic case with Gaussian priors, this constraint reduces to a particularly tractable form: Sims (2003) demonstrates that the optimal behavior of the action Y given state X is observationally equivalent to a signal-extraction problem in which X is observed with i.i.d. Gaussian noise whose variance is pinned down by the capacity constraint.[3] Our reduced-form representation is a direct implementation of this result: the processed message $\tilde{m}_B = m_{S\to B} + \xi_B$ with $\xi_B \sim \mathcal{N}(0, \kappa^{-1})$ and a convex processing cost $C(\kappa)$ captures exactly this structure, where κ indexes the precision of the effective Gaussian channel chosen by the buyer.

The broader rational-inattention literature extends this framework to general priors and discrete action spaces. Most notably, Matějka (2016) shows that a rationally inattentive seller subject to a Shannon mutual-information constraint will optimally set prices from a finite discrete support—a result driven by the interaction of the mutual-information constraint with non-Gaussian prior distributions and qualitatively distinct from what a Gaussian-channel implementation delivers. The economic force is the same: costly processing leads agents to condition on coarser representations of the environment, so that observed objects need not be treated as sufficient statistics. The present paper, however, follows Sims (2003) exclusively; the Gaussian endogenous-precision representation is the linear-quadratic special case that delivers closed-form posterior moments and keeps the no-trade mechanism transparent. The relation to Matějka (2016) and the broader rational-inattention literature is taken up in Section 6.

4.4. *Demands and Bilateral Clearing*

Given a price p and a processing choice κ, the buyer selects x to maximize expected utility from the trading position, taking the processing cost as fixed at the chosen κ:

$$\max_{x\in[-q_0,\,q_0]} \mathbb{E}\Big[-\exp\big(-\gamma_B\, x\,(v-p)\big)\Big|\,\mathcal{I}_B(\kappa)\Big] - C(\kappa) \tag{7}$$

Under CARA-normality, the buyer's interior optimum takes the mean-variance form:

$$x_B^*(p,\kappa) = \frac{\mu_B(\kappa) - p}{\gamma_B \sigma_B^2(\kappa)} \tag{8}$$

whenever this lies in $(-q_0, q_0)$, otherwise the solution is truncated at the relevant feasibility bound. A positive x_B^* means the buyer acquires units of the risky asset; a negative x_B^* means the buyer sells units short (up to the endowment constraint). In what follows, we focus on interior solutions; corner solutions arise only when the posterior gap is large relative to aggregate risk-bearing capacity.

Similarly, let $\mathcal{I}_S = \sigma(s_S, m_{B\to S})$ denote the seller's information set. [4] Then

$$v \mid \mathcal{I}_S \sim \mathcal{N}(\mu_S, \sigma_S^2) \tag{9}$$

---

[3] See Sims (2003), Section 3. When X is Gaussian and the objective is quadratic, the optimal conditional distribution $q(Y \mid X)$ is Gaussian, so the action behaves as if X were observed through a Gaussian channel with variance determined by the information constraint.

[4] The seller's frictionless processing is an identifying assumption, not an innocuous simplification. Its purpose is to isolate the buyer's processing margin as the sole source of posterior heterogeneity. If the seller also chose a processing intensity $\kappa_S \geq 0$, the model would have two endogenous precisions, and the Stage 1 choice would become a simultaneous-move game. In that extended game, both posterior means would be affected by processing noise, but the direction of the result is preserved: as long as at least one agent processes at finite precision and the information sets remain non-identical (Assumption 2), posterior means differ almost surely, and trade occurs with positive probability. The bilateral-asymmetric specification is therefore the minimal one that delivers the mechanism cleanly.

and the seller's interior optimal supply is

$$x_S^*(p) = -\frac{\mu_S - p}{\gamma_S \sigma_S^2} \tag{10}$$

truncated to respect the endowment constraint.

Definition 1 (Two-stage equilibrium with endogenous processing). *The trading game proceeds in two stages.*

Stage 1 (processing choice). *Before observing signals, the buyer commits to a processing intensity $\kappa \geq 0$. Given κ, the buyer's effective information set is $\mathcal{I}_B(\kappa) = \sigma(s_B, \tilde{m}_B)$, where $\tilde{m}_B = m_{S\to B} + \xi_B$ and $\xi_B \sim \mathcal{N}(0, \kappa^{-1})$. The buyer takes as given the mapping $\kappa \mapsto p^*(\kappa)$ induced by Stage 2 clearing and chooses κ^* to maximize the net expected certainty-equivalent objective $\Phi(\kappa) = \mathbb{E}[CE(p^*(\kappa), \kappa)] - C(\kappa)$.*

Stage 2 (bilateral clearing given κ). *Fix a processing intensity $\kappa \geq 0$. A bilateral clearing equilibrium given κ is a pair of measurable functions $(p^*(\kappa), x^*(\kappa))$ such that:*

(i) ***Optimality.*** *Given price $p^*(\kappa)$ and information sets $\mathcal{I}_B(\kappa)$ and $\mathcal{I}_S$, the buyer and seller choose CARA—normal optimal demands as in* (8) *and* (10) *(interior whenever feasible);*

(ii) ***Market clearing.*** $x_B(p^*(\kappa), \kappa) + x_S(p^*(\kappa)) = 0$*;*

(iii) ***Feasibility.*** *The traded quantity satisfies $|x^*(\kappa)| \leq q_0$.*

An equilibrium of the two-stage game is a processing choice κ^* solving Stage 1 and a clearing pair $(p^*(\kappa^*), x^*(\kappa^*))$ solving Stage 2 given κ^*.

Conditional on a processing choice κ, trade arises if and only if posterior means differ. The next two results show, first, how disagreement maps into the clearing quantity and, second, why such disagreement is generic under finite processing.

Lemma 1 (Closed-form price and trade given κ). *Under Definition 1, the unique clearing price and quantity are*

$$p^*(\kappa) = \frac{\gamma_S \sigma_S^2 \, \mu_B(\kappa) + \gamma_B \sigma_B^2(\kappa) \mu_S}{\gamma_B \sigma_B^2(\kappa) + \gamma_S \sigma_S^2}, \tag{11}$$

$$x^*(\kappa) = \frac{\mu_B(\kappa) - \mu_S}{\gamma_B \sigma_B^2(\kappa) + \gamma_S \sigma_S^2}. \tag{12}$$

Hence $x^*(\kappa) \neq 0$ if and only if $\mu_B(\kappa) \neq \mu_S$.

Assumption 2 (Nondegenerate information and processing). *The following conditions rule out degenerate information and processing cases:*

i Assume that the counterpart message is informative. That is, $\tau_B \in (0, \infty)$,

ii Assume that the buyer cannot perfectly recover the seller's sufficient statistic from the message: the counterpart message and/or processing is noisy (e.g., $\tau_B < \infty$ and/or the optimal choice satisfies $\kappa^* < \infty$), and

iii Assume that the two agents' information sets are not identical in the Gaussian sense, i.e. $\mathcal{I}_B(\kappa) \neq \mathcal{I}_S$. More generally, the linear projections defining $\mu_B(\kappa)$ and μ_S use different sufficient statistics.

The following lemma shows that posterior disagreement is generic in the Gaussian environment.

Lemma 2 (Posterior disagreement is generic). *Under the Assumption 2 and the Gaussian primitives described in equations* (1), (3) *and* (4)*, the posterior mean gap $\mu_B(\kappa) - \mu_S$ is a normal non-degenerate random variable for*

$\kappa \in [0, \infty)$. *In particular,*

$$\mathbb{P}(\mu_B(\kappa) \neq \mu_S) = 1, \tag{13}$$

$$\mathbb{P}(\mu_B(\kappa) > \mu_S) \in (0, 1). \tag{14}$$

The proof of this lemma is found in Appendix Appendix C. An eagle-eye view is presented here. Under the Gaussian primitives described in equations (1), (3) and (4), the payoff v and all observed signals/messages are jointly Gaussian. In Gaussian environments, conditional expectations are affine functions of the conditioning variables, so both posterior means can be written as

$$\mu_S = \mathbb{E}[v \mid \mathcal{I}_S] = a_S + b_S^\top Y_S, \qquad \mu_B(\kappa) = \mathbb{E}[v \mid \mathcal{I}_B(\kappa)] = a_B(\kappa) + b_B(\kappa)^\top Y_B(\kappa),$$

for appropriate Gaussian vectors Y_S and $Y_B(\kappa)$. Hence the posterior-mean gap $\mu_B(\kappa) - \mu_S$ is itself an affine transformation of a jointly Gaussian vector and is therefore normally distributed. Assumption 2 rules out perfect spanning (the buyer does not perfectly recover the seller's sufficient statistic because the counterpart message and/or processing is noisy, and the information sets are not identical), which implies $\mathrm{Var}(\mu_B(\kappa) - \mu_S) > 0$. The gap is thus a non-degenerate normal random variable, so $\mathbb{P}(\mu_B(\kappa) \neq \mu_S) = 1$ and $\mathbb{P}(\mu_B(\kappa) > \mu_S) \in (0, 1)$.

The following lemma separates the two sources of posterior disagreement by comparing three limiting configurations and shows that no-trade requires both processing noise and message noise to vanish.

Lemma 3 (Three limits and the approach to no-trade). *Maintain the Gaussian primitives of equations* (1)—(4) *and the symmetry conditions:*

(a) ***Symmetric message precision:*** $\tau_B = \tau_S \equiv \tau \in (0, \infty)$;

(b) ***Symmetric private signal precision:*** $\sigma_S^2 = \sigma_B^2 \equiv \sigma^2$;

(c) ***Symmetric risk aversion:*** $\gamma_S = \gamma_B \equiv \gamma$.

The model admits three distinct limiting configurations that separate the two sources of posterior disagreement:

(i) ***Finite processing, noisy message*** ($\kappa < \infty$, $\tau < \infty$)*: The buyer's effective signal* $\tilde{m}_B = m_{S \to B} + \xi_B$ *retains processing noise* $\xi_B \sim \mathcal{N}(0, \kappa^{-1})$ *in addition to the message noise already embedded in* $m_{S \to B}$. *The posterior-mean gap variance satisfies* $V(\kappa) > V(\infty) > 0$, *and posterior disagreement is generic:*

$$\mathbb{P}(\mu_B(\kappa) \neq \mu_S) = 1.$$

This is the costly-processing economy studied throughout the paper. Trade occurs with positive probability.

(ii) ***Perfect processing, noisy message*** ($\kappa \to \infty$, $\tau < \infty$)*: Processing noise vanishes* ($\xi_B \to 0$ *a.s.*)*, so the buyer conditions on* $m_{S \to B} = s_S + \eta_S$ *directly. However, message noise* (η_S, η_B) *remains. The posterior means converge to*

$$\mu_B(\infty) = \alpha\, s_B + \beta\,(s_S + \eta_S), \tag{15}$$

$$\mu_S = \alpha\, s_S + \beta\,(s_B + \eta_B), \tag{16}$$

and the residual gap variance satisfies

$$V(\infty) \;\equiv\; \mathrm{Var}(\mu_B(\infty) - \mu_S) = 2(\alpha - \beta)^2 \sigma^2 + 2\beta^2 \tau^{-1} > 0$$

for all $\tau < \infty$. *Disagreement and trade persist even under costless processing. The processing margin has been eliminated; the remaining source of disagreement is the primitive message noise, which is a feature of the information structure and independent of* $C(\kappa)$.

(iii) ***Perfect processing, perfect message*** *($\kappa \to \infty, \tau \to \infty$): Both noise sources vanish. Noiseless messages convey each agent's private signal exactly ($m_{S \to B} \to s_S$, $m_{B \to S} \to s_B$), so both agents' information sets converge to $\mathcal{I}_B(\infty) = \mathcal{I}_S = \sigma(s_S, s_B)$. Posterior means coincide almost surely:*

$$\mathbb{P}(\mu_B(\infty) = \mu_S) = 1, \qquad \mathbb{P}(x^*(\infty) = 0) = 1.$$

This joint limit recovers the no-trade benchmark of Milgrom and Stokey: with common priors and perfectly shared information, speculative trade is eliminated.

Proof 1. *Case (i). Under $\kappa < \infty$ and $\tau < \infty$, the buyer's effective conditioning set $\mathcal{I}_B(\kappa) = \sigma(s_B, \tilde{m}_B)$ differs from the seller's $\mathcal{I}_S = \sigma(s_S, m_{B \to S})$ due to the presence of independent processing noise ξ_B. Lemma 2 establishes that $\mathbb{P}(\mu_B(\kappa) \neq \mu_S) = 1$ under Assumption 2, and $V(\kappa) > V(\infty)$ follows from the decomposition in Proposition 2.*

Case (ii). As $\kappa \to \infty$, $\xi_B \sim \mathcal{N}(0, \kappa^{-1}) \to 0$ almost surely, so $\tilde{m}_B \to m_{S \to B} = s_S + \eta_S$ almost surely and the buyer's information set converges to $\mathcal{I}_B(\infty) = \sigma(s_B, m_{S \to B})$.

Under conditions (a)–(c), the posterior means take the form in (15)–(16), *where the coefficients from the standard Gaussian projection are:*

$$\omega(\tau) = (\sigma^2 + \tau^{-1})^{-1}, \qquad \alpha = \frac{\sigma^{-2}}{\sigma_v^{-2} + \sigma^{-2} + \omega(\tau)}, \qquad \beta = \frac{\omega(\tau)}{\sigma_v^{-2} + \sigma^{-2} + \omega(\tau)}. \tag{17}$$

Taking the difference:

$$\mu_B(\infty) - \mu_S = (\alpha - \beta)(s_B - s_S) + \beta(\eta_S - \eta_B), \tag{18}$$

which has variance $V(\infty) = 2(\alpha - \beta)^2\sigma^2 + 2\beta^2\tau^{-1}$. Both terms are nonnegative, and their sum is strictly positive for any $\sigma^2 < \infty$ and $\tau \in (0, \infty)$: the first term can vanish only in the noiseless-message limit $\omega(\tau) = \sigma^{-2}$, while the second term vanishes only when message noise disappears. Hence $V(\infty) > 0$ for all finite τ, and $\mathbb{P}(\mu_B(\infty) \neq \mu_S) = 1$.

Case (iii). As additionally $\tau \to \infty$, the message noise $\eta_S, \eta_B \to 0$ almost surely. Noiseless messages yield $m_{S \to B} \to s_S$ and $m_{B \to S} \to s_B$, so $\mathcal{I}_B(\infty) = \mathcal{I}_S = \sigma(s_S, s_B)$. Both agents condition on the same sufficient statistic, and by standard Bayesian updating, $\mu_B(\infty) = \mu_S$ almost surely. Equivalently, $V(\infty) \to 2(\alpha - \beta)^2\sigma^2 + 2\beta^2\tau^{-1} \to 0$ as $\tau \to \infty$, since $\omega(\tau) \to \sigma^{-2}$, hence $\alpha - \beta \to 0$, while β remains bounded and $\beta^2\tau^{-1} \to 0$. Hence $\mathbb{P}(x^*(\infty) = 0) = 1$.

Summary. *Two distinct frictions govern posterior disagreement: processing noise (endogenous, eliminated as $\kappa \to \infty$) and message noise (primitive, eliminated only as $\tau \to \infty$). The no-trade benchmark requires the joint limit $(\kappa, \tau) \to (\infty, \infty)$.*

Remark 1 (Endogenous versus primitive disagreement). *Cases (i)–(iii) of Lemma 3 identify two conceptually distinct sources of posterior disagreement. Message noise ($\tau < \infty$) is a primitive of the information structure: it generates the irreducible floor $V(\infty) > 0$ and is independent of processing costs. Processing noise ($\kappa^{-1} > 0$) is endogenous: it is disciplined by $C(\kappa)$ through the first-order condition $C'(\kappa^*) = \Delta(\kappa^*)$. The paper's mechanism operates exclusively through the endogenous component. The processing cost $C(\kappa)$ determines how far above $V(\infty)$ the equilibrium gap variance $V(\kappa^*)$ lies, and hence the magnitude—but not the binary existence—of trade.*

The following proposition presents the equilibrium with endogenous information processing and the resulting trade.

Proposition 1 (Equilibrium with endogenous processing and trade). *Maintain Assumption 1 (convex processing costs) and Assumption 2 (Nondegenerate information processing)*

Define the buyer's objective

$$\Phi(\kappa) \equiv \mathbb{E}[CE(p^*(\kappa), \kappa)] - C(\kappa), \tag{19}$$

where $CE(\cdot)$ is the certainty-equivalent gain from trading at the clearing price.[5]

---

[5] An explicit derivation of CE and the processing FOC is given in Appendix Appendix B.

Then:

1. ***Existence of an optimal processing choice.*** *There exists $\kappa^* \in [0, \infty)$ that maximizes $\Phi(\kappa)$.*
2. ***Interior characterization.*** *If $\kappa^* > 0$ is interior, it satisfies the first-order condition*

$$C'(\kappa^*) = \Delta(\kappa^*) \tag{20}$$

where $\Delta(\kappa) \equiv \frac{d}{d\kappa}\mathbb{E}[CE(p^*(\kappa), \kappa)]$.[6]. In addition, if Φ is concave in a neighborhood of κ^*, the condition is sufficient, and κ^* is locally unique.

3. ***Sources of trade and the role of processing costs.*** *The clearing quantity at κ^* is given by* (12). *Posterior disagreement, and hence trade, arises from two distinct sources that Lemma 3 disentangles:*
    - *(i)* Message noise *($\tau_B < \infty$): even with perfect processing ($\kappa = \infty$), the buyer cannot recover the seller's private signal s_S exactly because the message $m_{S\to B} = s_S + \eta_S$ contains irreducible noise. This source of disagreement is a primitive of the information structure and does not depend on processing costs.*
    - *(ii)* Processing noise *($\kappa^* < \infty$): the buyer's optimal choice of finite precision adds an additional layer of noise $\xi_B \sim \mathcal{N}(0, (\kappa^*)^{-1})$ to the message. This is the endogenous margin governed by $C(\kappa)$.*

The processing cost $C(\kappa)$ determines κ^* and thereby controls the magnitude of source (ii). Under Assumption 2 and Lemma 2, the combined effect of both sources yields:

$$\mathbb{P}(x^*(\kappa^*) \neq 0) = 1. \tag{21}$$

Moreover, for any $q_0 > 0$,

$$\mathbb{P}(|x^*(\kappa^*)| < q_0) > 0, \tag{22}$$

so interior feasible trade occurs with positive probability.

The processing cost pins down equilibrium trade volume. *Higher $C'(\cdot)$ implies lower κ^*, more processing noise, and a wider posterior gap on average. Lower $C'(\cdot)$ implies higher κ^*, less processing noise, and a narrower gap. The processing cost, therefore, determines the equilibrium level of disagreement attributable to source (ii), while source (i) sets a lower bound on disagreement that persists even in the frictionless-processing limit $\kappa \to \infty$ (see Lemma 3).*

Proposition 1 establishes two distinct facts that should not be conflated.

First, posterior means differ almost surely (Lemma 2), so the unconstrained clearing quantity is nonzero almost surely. This follows from Assumption 2 and holds for any $\kappa \in [0, \infty)$, including the frictionless limit $\kappa \to \infty$, so long as message noise $\tau_B < \infty$ (see Lemma 3). In other words, the existence of trade is a consequence of the primitive information structure, not of the processing cost per se.

Second, the processing cost $C(\kappa)$ determines the equilibrium magnitude of disagreement by pinning down κ^*. Higher $C'(\cdot)$ lowers κ^*, increases the processing-noise contribution to the posterior gap, and raises expected trade volume when the numerator effect dominates the induced movement in aggregate risk-bearing capacity. The comparative statics of trade volume with respect to processing costs are therefore the empirically relevant prediction of the model, not the binary existence of trade.

Whether realized trade is interior, that is $0 < |x^*(\kappa^*)| < q_0$, additionally depends on the feasibility constraint. Since $\mu_B(\kappa^*) - \mu_S$ is normally distributed with mean zero and variance $V^* > 0$, the event

$$\left\{ 0 < \frac{|\mu_B(\kappa^*) - \mu_S|}{\gamma_B \sigma_B^2(\kappa^*) + \gamma_S \sigma_S^2} < q_0 \right\} \tag{23}$$

---

[6] See Appendix B

has a positive probability for any $q_0 > 0$. When q_0 is large enough that the endowment constraint never binds almost surely, interior trade has a probability approaching one.

Remark 2 (From posterior disagreement to mutually beneficial trade). *The results operate at four distinct levels that carry different epistemic weight, and it is useful to keep them separate. Posterior disagreement: Lemma 2 establishes* $\mathbb{P}(\mu_B(\kappa^*) \neq \mu_S) = 1$*—the strongest result, holding in every state of the world under Assumption 2. Desired trade: because* $x^*(\kappa) = (\mu_B(\kappa) - \mu_S)/D(\kappa)$*, any nonzero gap maps one-for-one into a nonzero desired trade quantity, so* $\mathbb{P}(x^*(\kappa^*) \neq 0) = 1$ *as well; this includes realizations where the desired quantity would exceed the endowment bound. Interior trade: the event* $\{0 < |x^*(\kappa^*)| < q_0\}$ *has strictly positive probability for any* $q_0 > 0$*, as* (23) *establishes—the model is somewhat weaker here, showing that interior trade is possible rather than guaranteed. Mutually voluntary exchange: at the clearing price* $p^*(\kappa^*)$*, which by* (11) *lies strictly between* $\mu_B(\kappa^*)$ *and* μ_S *whenever they differ, each agent trades in the direction that increases conditional expected utility, so the exchange is mutually voluntary. The model does not, however, establish a full Pareto welfare comparison relative to an ex-ante optimal allocation—that would require a richer specification of initial endowments and outside options beyond the scope of the present analysis.*

We define the marginal value of precision as the exact derivative of the buyer's expected certainty-equivalent trading gains, evaluated at the clearing price.

$$\Delta(\kappa) \equiv \frac{d}{d\kappa}\mathbb{E}[\, CE(p^*(\kappa), \kappa)] , \tag{24}$$

where $CE(\cdot)$ is the maximized certainty-equivalent under CARA-normality. Appendix Appendix B derives $\Delta(\kappa)$ and the processing FOC formally. For intuition, Appendix Appendix B also reports a sufficient statistic expression that highlights the risk-reduction channel under additional restrictions.

Lemma 1 shows that the equilibrium trade is driven by posterior disagreement, as in (12). Thus, conditional on κ, the model suggests that speculative positions scale with the difference in the posterior means and are dampened by the aggregate risk-bearing capacity. The clearing price in (11) is a risk-adjusted average of the posterior means, so the prices partially reflect both agents' valuations, with weights increasing in the effective risk tolerance of the other side.

Information processing affects trade through the buyer's posterior moments, $\mu_B(\kappa)$ and $\sigma_B^2(\kappa)$. The higher processing intensity increases the precision with which the buyer interprets the counterpart message, shifting $\mu_B(\kappa)$ towards the seller's valuation and lowering $\sigma_B^2(\kappa)$. The endogenous choice of κ^* equates marginal processing costs and the marginal value of precision in (20), thereby pinning down the extent to which disagreement and residual risk persist in equilibrium.

Whenever processing information is effectively frictionless, the buyer chooses high precision, thus shrinking posterior disagreement and dampening trade. On the other hand, whenever processing is prohibitively expensive, the buyer ignores the counterpart message. In this case, the informational gains from interpreting counterpart information vanish, and trade is again reduced. These limiting cases connect the illustrative environment back to the posterior logic of alignment underlying the no-trade arguments: when the mapping from observables to posteriors becomes effectively frictionless, speculative motives collapse.

If $\kappa^* > 0$ is an interior solution, a necessary condition is $C'(\kappa^*) = \Delta(\kappa^*)$. Under local concavity of the indirect objective, this condition is also sufficient (See Appendix Appendix B for the derivation and sufficient conditions).

4.5. Messages as stand-ins for informational prices

In MS, posterior alignment is driven by a public object—the equilibrium price vector—that can become a sufficient statistic for payoff-relevant uncertainty under rational expectations. Our illustrative environment is deliberately more parsimonious and does not model a full rational-expectations price-formation mechanism. Instead, counterpart messages serve as reduced-form proxies for observable market objects—prices, quotes, disclosures, or communications—that aggregate information but must themselves be interpreted.

The results should therefore be read as statements about the informational content of observed objects when agents face costly processing, rather than as a full price-formation model. The crucial issue for no-trade is whether agents can costlessly extract and condition on the informational content of what they observe. When processing is costly, a processed statistic (a message here, prices in a fuller rational-expectations equilibrium) need not be sufficient. As a result, posterior alignment can fail even under common priors, and trade can arise in equilibrium.

The following lemma states the same point using a public statistic.

Lemma 4 (Processing-stage disagreement with a public statistic). *Let Z be a public statistic observed by both agents, and suppose that v and Z are jointly Gaussian with $\mathrm{Var}(Z) \equiv \sigma_Z^2 > 0$ and $\mathrm{Cov}(v, Z) \equiv \rho \neq 0$. The seller conditions on Z frictionlessly, while the buyer conditions on a processed version $\tilde{Z} = Z + \xi(\kappa)$, where $\xi(\kappa) \sim \mathcal{N}(0, \kappa^{-1})$ is independent of (v, Z).*

Then, for any finite $\kappa < \infty$, the posterior means of the two agents differ:

$$\mathbb{E}[v \mid Z] \neq \mathbb{E}[v \mid \tilde{Z}] \quad \textit{with probability one.} \tag{25}$$

In particular, posterior alignment fails even though the underlying informational object Z is public.

Proof 2. *Under joint Gaussianity, posterior means are linear projections. The seller's posterior mean is*

$$\mathbb{E}[v \mid Z] = \mu_0 + \frac{\rho}{\sigma_Z^2}(Z - \mu_Z), \tag{26}$$

where $\mu_Z = \mathbb{E}[Z]$.

For the buyer, the processed signal $\tilde{Z} = Z + \xi(\kappa)$ has variance $\mathrm{Var}(\tilde{Z}) = \sigma_Z^2 + \kappa^{-1}$, and $\mathrm{Cov}(v, \tilde{Z}) = \mathrm{Cov}(v, Z) = \rho$ (since $\xi(\kappa)$ is independent of v). The buyer's posterior mean is therefore

$$\mathbb{E}[v \mid \tilde{Z}] = \mu_0 + \frac{\rho}{\sigma_Z^2 + \kappa^{-1}}(\tilde{Z} - \mu_Z). \tag{27}$$

Taking the difference:

$$\mathbb{E}[v \mid Z] - \mathbb{E}[v \mid \tilde{Z}] = \rho\left(\frac{1}{\sigma_Z^2} - \frac{1}{\sigma_Z^2 + \kappa^{-1}}\right)(Z - \mu_Z) - \frac{\rho}{\sigma_Z^2 + \kappa^{-1}}\,\xi(\kappa). \tag{28}$$

The attenuation factor satisfies

$$\frac{1}{\sigma_Z^2} - \frac{1}{\sigma_Z^2 + \kappa^{-1}} = \frac{\kappa^{-1}}{\sigma_Z^2(\sigma_Z^2 + \kappa^{-1})} > 0 \qquad \textit{for all } \kappa < \infty, \tag{29}$$

and the second term $\frac{\rho}{\sigma_Z^2+\kappa^{-1}}\,\xi(\kappa)$ is a nondegenerate Gaussian random variable (since $\rho \neq 0$ and $\kappa < \infty$). Together, the right-hand side of (28) is a Gaussian random variable with strictly positive variance, so the two posterior means differ with probability one. The degenerate cases, whenever $\kappa = \infty$ (no processing noise) or $\sigma_Z^2 = 0$ (uninformative statistic) or $\rho = 0$ (irrelevant statistic), are excluded by assumption.

4.6. How This Differs from Kyle (1985) and Grossman and Stiglitz (1980)

Our mechanism is closest in spirit to the literature on information frictions, but it operates at a different margin from the two canonical frameworks.

In Grossman and Stiglitz (1980), the key wedge is information acquisition: agents decide whether to become informed by paying a cost, and equilibrium prices cannot be fully revealed because full disclosure would eliminate incentives to acquire information. By contrast, our wedge is information processing: agents may have access to an informative statistic (a counterpart message here and prices in a fuller rational-expectations environment), but they optimally choose limited precision in interpreting it. The Lemma 4 illustrates the distinction sharply: posterior disagreement can persist even when the underlying informational object is public, simply because agents condition on processed versions of that object. In this sense, our mechanism targets information overload and interpretive complexity, rather than incentives to acquire information.

Kyle (1985) generates trade and partial revelation through the submission of strategic orders by an insider, a market maker who sets prices, and noise trading that prevents full inference. Our environment abstracts from these institutional features. Trade does not rely on strategic behavior or market making; instead, it emerges because limited processing prevents observed statistics from becoming sufficient, sustaining endogenous disagreement even in a stripped-down bilateral clearing setup.

The contribution is therefore orthogonal to both classic mechanisms: we provide a tractable bridge from no-trade logic to observed trading by relaxing a knife-edge cognitive benchmark (i.e., frictionless processing) without abandoning common priors or introducing strategic microstructure.

5. Implications for the No-Trade Theorem

This section connects the trading environment back to the logic in Sections 2 and 3. The analytical structure reveals something sharp: the force of no-trade arguments depends critically on an implicit cognitive requirement—that agents can extract and condition on the informational content of observed objects without cost. Once processing is costly and optimally limited, this requirement fails, and posterior alignment is no longer guaranteed. The argument proceeds in three steps: the frictionless benchmark in Section 5.1, the costly-processing economy in Section 5.2, and the comparative statics connecting processing costs to trade volume in Section 5.5.

5.1. *The frictionless-processing benchmark*

To isolate the processing margin, we compare the costly-processing economy against two progressively more frictionless benchmarks. Lemma 3 separates the three cases.

Case (i): Finite processing, noisy message ($\kappa < \infty$, $\tau < \infty$). This is the economy studied throughout the paper. The buyer's effective signal $\tilde{m}_B = m_{S\to B} + \xi_B$ carries both processing noise $\xi_B \sim \mathcal{N}(0, \kappa^{-1})$ and the message noise already embedded in $m_{S\to B}$. The two agents' information sets are strictly non-identical, posterior disagreement is generic, and trade occurs with positive probability. The processing cost $C(\kappa)$ disciplines the endogenous component of disagreement by pinning down κ^*.

Case (ii): Perfect processing, noisy message ($\kappa \to \infty$, $\tau < \infty$). Suppose processing costs vanish and the buyer conditions on the message $m_{S\to B}$ directly. Processing noise is gone, yet posterior means still differ almost surely. Lemma 3(ii) shows that the residual gap variance is

$$V(\infty) = 2(\alpha - \beta)^2\sigma^2 + 2\beta^2\tau^{-1} > 0, \tag{30}$$

which is strictly positive for any finite message precision τ. The source of disagreement here is the primitive message noise (η_S, η_B), which is a feature of the information structure and orthogonal to the processing margin. Trade would persist even if processing were free. This case makes precise why the processing cost $C(\kappa)$ governs only part of the posterior-mean gap: the component $V(\infty)$ is irreducible from the perspective of the processing margin.

Case (iii): Perfect processing, perfect message ($\kappa \to \infty$, $\tau \to \infty$). When both sources of noise vanish simultaneously, each agent learns the other's private signal exactly, and the two information sets converge to the same object: $\mathcal{I}_B(\infty) = \mathcal{I}_S = \sigma(s_S, s_B)$ almost surely. Posterior means then coincide:

$$\mathbb{P}(\mu_B(\infty) = \mu_S) = 1 \qquad \text{and} \qquad \mathbb{P}(x^*(\infty) = 0) = 1. \tag{31}$$

This joint limit recovers the no-trade benchmark of Milgrom and Stokey. When agents can extract the full informational content of what they receive without effort, and when what they receive is itself noiseless, common priors drive out speculative trade.

The comparison across the three cases clarifies what the no-trade logic requires and where the paper's mechanism operates. The processing cost $C(\kappa)$ is an endogenous friction that shifts the economy from case (iii) toward case (i)—it determines how far the equilibrium sits from the no-trade benchmark. Message noise τ is a primitive friction that sets the floor $V(\infty)$: even if processing were free, some disagreement would remain. The paper's comparative static is entirely about the endogenous component: higher costs reduce κ^*, raise $\mathrm{Var}(\delta(\kappa^*))$, and push $V(\kappa^*)$ further above its irreducible floor.

5.2. *The costly-processing economy: finite precision and posterior disagreement*

We now depart from the frictionless limit and ask: what happens when processing is expensive and the buyer optimally chooses finite precision?

The answer is direct. Under finite $\kappa^* \in (0, \infty)$, the buyer's effective signal $\tilde{m}_B = m_{S\to B} + \xi_B$ with $\xi_B \sim \mathcal{N}(0, (\kappa^*)^{-1})$ remains noisy. Combined with the seller's information set $\mathcal{I}_S = \sigma(s_S, m_{B\to S})$, the two agents' information sets are strictly non-identical. Lemma 2 establishes that under Assumption 2, this non-identity leads to:

$$\mathbb{P}(\mu_B(\kappa^*) \neq \mu_S) = 1. \tag{32}$$

That is, posterior disagreement is *generic*. Moreover, by Lemma 1, this disagreement maps directly into trade:

$$x^*(\kappa^*) = \frac{\mu_B(\kappa^*) - \mu_S}{\gamma_B \sigma_B^2(\kappa^*) + \gamma_S \sigma_S^2}. \tag{33}$$

Because the posterior-mean gap is a non-degenerate Gaussian random variable, interior trade (i.e., $0 < |x^*(\kappa^*)| < q_0$) occurs with positive probability. The mechanism is straightforward: finite processing precision prevents the buyer from fully extracting the informational content of the seller's message, leaving a residual disagreement that sustains trade.

5.3. *Isolating the processing wedge: the role of sufficiency*

The contrast between the frictionless limit and the costly-processing case pinpoints where the no-trade logic breaks down. The key is the notion of a *sufficient statistic*.

In the benchmark setting (Milgrom-Stokey), the price change (or equivalently, in our bilateral setting, the public message) serves as a sufficient statistic: conditioning on it reproduces the posterior one would form from conditioning on all available information in the market. This sufficiency is what makes the price a powerful aggregator: no matter what private signal an agent held, the price contains all relevant information, making private information redundant for belief updating. This is the swamping property mentioned in Section 3.

Processing frictions break this sufficiency. To see this precisely, consider the following thought experiment. Suppose there is a public object Z (a price, a message, a disclosure) that would be sufficient—in the information-theoretic sense—if agents could condition on it perfectly. However, suppose agents only observe a processed version $\tilde{Z} = Z + \xi(\kappa)$ where the noise $\xi(\kappa) \sim \mathcal{N}(0, \kappa^{-1})$ is independent of the payoff-relevant state. Lemma 4 shows that for any $\kappa < \infty$,

$$\mathbb{E}[v \mid Z] \neq \mathbb{E}[v \mid \tilde{Z}] \quad \text{with probability one.} \tag{34}$$

The processed signal $\tilde{Z}$ is strictly less informative than the raw object Z. More importantly, conditioning on $\tilde{Z}$ does not subsume the agent's private signal:

$$\pi(\cdot \mid \tilde{Z}, s) \neq \pi(\cdot \mid \tilde{Z}) \quad \text{generically.} \tag{35}$$

This is what we call *processing-stage sufficiency failure*. It is distinct from the Milgrom-Stokey swamping result, which is an equilibrium statement about competitive markets with many traders and fully revealing prices. Rather, processing-stage sufficiency failure captures the same underlying economic logic at the cognitive level: when agents cannot costlessly extract the informational content of a public object, that object ceases to align beliefs.

The point is not that common priors or Bayesian rationality breaks down. Agents in our model continue to update Bayesianly on what they observe. Rather, what they observe—the effective information on which they condition—is coarser than the raw public object. Because their effective information sets differ, their posteriors differ. And because posteriors differ, trade occurs.

5.4. Which assumption of Milgrom-Stokey is relaxed?

With the above analysis in hand, we can state precisely how this paper relates to the Milgrom-Stokey framework. The Milgrom-Stokey theorem, as discussed in Section 3, establishes that speculative trade cannot be strictly mutually beneficial once feasibility and mutual acceptability are common knowledge, provided agents share common or concordant priors.

The theorem's power derives from a two-step argument: (i) agents observe public information (prices, announcements), and (ii) they then condition on this public information, updating beliefs in a way that aligns posterior valuations. The second step is typically treated as automatic—an agent observes a price and instantly uses it to infer what others must know. This paper makes the second step explicit: conditioning on public information is not automatic; it requires cognitive effort, and when that effort is costly, agents optimally limit their precision.

The implicit assumption we relax is therefore best stated as follows: *agents can costlessly extract and condition on the full informational content of observed equilibrium objects (prices, messages, disclosures).* This is the "frictionless conditioning" assumption.

Under costly processing, this assumption fails. Agents observe the public object but process it imperfectly, effectively observing only a noisy version. As a result, the public object does not achieve sufficiency, posterior means do not align, and trade can occur even under common priors.

To be clear, we do not claim that Milgrom-Stokey is wrong or that it has been overthrown by adding processing costs. Rather, we identify a condition—costless cognitive processing—that the theorem implicitly relies on. When this condition fails endogenously (because agents choose finite precision to balance informational benefits against costs), the posterior-alignment mechanism breaks down, and the no-trade conclusion need not hold.

5.5. The role of processing costs in determining trade volume

While the above results establish the existence of trade when processing is costly, they do not yet address the quantitative relationship between processing costs and trading volume. This relationship is central to understanding the paper's empirical implications.

Proposition 1 characterizes the two-stage equilibrium and highlights two distinct facts that should not be conflated. First, posterior disagreement exists almost surely: the posterior-mean gap $\mu_B(\kappa^*) - \mu_S$ is a non-degenerate Gaussian random variable. This binary fact—that trade *can* occur—follows from the information structure (specifically, from Assumption 2) and obtains for any $\kappa \in [0, \infty)$, including the frictionless limit $\kappa \to \infty$, provided message noise $\tau < \infty$.

Second, and more importantly, the processing cost $C(\kappa)$ determines the equilibrium magnitude of disagreement by pinning down κ^*. The first-order condition $C'(\kappa^*) = \Delta(\kappa^*)$ equates the marginal cost of processing precision with the marginal value of that precision. Higher marginal processing costs (larger $C'(\kappa^*)$) lead to a lower optimal precision κ^*. Lower precision means more processing noise κ^{-1}, which widens the posterior gap. Since trade volume is proportional to the posterior-mean gap divided by the aggregate risk-bearing term, this raises expected trade volume whenever the numerator effect dominates the induced movement in $D(\kappa)$.

Conversely, if processing becomes cheaper (lower $C'(\kappa^*)$), the buyer optimally chooses higher κ^*, reducing processing noise and narrowing the posterior-mean gap. In the limit, as $C(\kappa) \to 0$ (costless processing), the buyer chooses $\kappa^* \to \infty$, the endogenous component of the posterior gap shrinks, and trade volumes decline toward the no-trade benchmark under the same risk-capacity condition.

This comparative static—that processing costs determine equilibrium posterior disagreement and, under the stated condition, equilibrium trading volume—is the model's central empirical prediction. It differentiates the mechanism from alternative explanations based on heterogeneous priors or noise trading. In those alternative theories, trading volume is governed by how information is distributed or how much noise is present. In the present framework, the key determinant is whether agents can internalize the information they receive. This leads to a distinctive comparative static: **simplifying the format or presentation of public disclosure, while leaving its informational content unchanged, should reduce trading volume by lowering effective processing costs and thus diminishing posterior disagreement**.

This prediction does not follow from heterogeneous-priors or differential-information models. It provides an empirically testable way to distinguish the processing-cost channel from its alternatives.

The following proposition formalizes the comparative statics of the processing margin and establishes the chain from cost level to posterior disagreement and, under a transparent sufficient condition, equilibrium trade volume.

Proposition 2 (Processing costs and equilibrium trade: comparative statics). *Maintain Assumption 1 and Assumption 2. Consider a one-parameter family of cost functions* $C_\lambda(\kappa) = \lambda \cdot c(\kappa)$*, where* $\lambda > 0$ *scales the level of processing costs and* $c(\cdot)$ *satisfies Assumption 1. Let* $\kappa^*(\lambda)$ *be the interior solution to* $\lambda c'(\kappa^*) = \Delta(\kappa^*)$*, with the second-order condition* $\Delta'(\kappa^*) - \lambda c''(\kappa^*) < 0$ *maintained throughout. Define* $V(\kappa) \equiv \mathrm{Var}(\mu_B(\kappa) - \mu_S)$ *as the variance of the posterior mean gap and* $D(\kappa) \equiv \gamma_B \sigma_B^2(\kappa) + \gamma_S \sigma_S^2$ *as the aggregate risk-bearing capacity in* (12).

Then:

1. ***Higher processing costs reduce optimal precision.*** $\kappa^*(\lambda)$ *is strictly decreasing in* λ*. By the implicit function theorem applied to the first-order condition* $F(\kappa^*, \lambda) \equiv \Delta(\kappa^*) - \lambda c'(\kappa^*) = 0$*:*

$$\frac{d\kappa^*(\lambda)}{d\lambda} = \frac{c'(\kappa^*)}{\Delta'(\kappa^*) - \lambda c''(\kappa^*)} < 0, \tag{36}$$

   where the sign follows from $c'(\kappa^*) > 0$ *and the second-order condition. When processing is more costly, the buyer optimally invests less effort in interpreting the counterpart message.*

2. ***Lower precision widens posterior disagreement.*** *Let* $\mu_B(\infty) \equiv \mathbb{E}[v \mid s_B, m_{S\to B}]$ *be the frictionless-processing posterior mean and* $\delta(\kappa) \equiv \mu_B(\kappa) - \mu_B(\infty)$ *the processing error attributable to finite* κ*. Because* ξ_B *is mean-zero and independent of all primitives, the tower property gives* $\mathbb{E}[\delta(\kappa) \mid s_B, m_{S\to B}] = 0$*, so the gap variance decomposes as*

$$V(\kappa) = \underbrace{V(\infty)}_{\text{message noise}} + \underbrace{\mathrm{Var}(\delta(\kappa))}_{\text{processing noise}}, \tag{37}$$

   where $V(\infty) \equiv \mathrm{Var}(\mu_B(\infty) - \mu_S) > 0$ *is the irreducible disagreement that survives even under frictionless processing (Lemma 3), and* $\mathrm{Var}(\delta(\kappa))$ *is strictly positive for any finite* κ *and strictly decreasing in* κ*, vanishing as* $\kappa \to \infty$*. Thus* $V(\kappa)$ *is strictly decreasing in* κ*: higher precision brings the buyer's effective belief closer to the frictionless benchmark, narrowing the gap with the seller.*

3. ***Processing costs determine disagreement and conditionally determine trade volume.*** *Parts (i) and (ii) together yield the chain*

$$\lambda \uparrow \implies \kappa^*(\lambda) \downarrow \implies V(\kappa^*(\lambda)) \uparrow \implies \mathbb{E}[|\mu_B(\kappa^*(\lambda)) - \mu_S|] \uparrow, \tag{38}$$

where the last step uses $\mathbb{E}[|\mu_B - \mu_S|] = \sqrt{2V/\pi}$*. The expected absolute trade volume satisfies*

$$\mathbb{E}[|x^*(\kappa^*(\lambda))|] = \frac{\sqrt{2V(\kappa^*(\lambda))/\pi}}{D(\kappa^*(\lambda))}, \tag{39}$$

and inherits the same direction whenever the seller's risk-capacity term dominates the denominator—that is, when $\gamma_S\sigma_S^2 \geq \gamma_B\sigma_B^2(\kappa^*)$*—so that* $D(\kappa^*)$ *is approximately constant and the numerator effect in* (39) *determines the sign.*[7]

Proposition 2 establishes the comparative statics that Proposition 1 left implicit. The decomposition (37) is the key step: it isolates the processing-noise component of disagreement—$\mathrm{Var}(\delta(\kappa))$— from the irreducible message-noise component $V(\infty)$. The former is the margin the buyer controls; the latter is not. Higher processing costs compress the buyer's choice of κ^*, expanding $\mathrm{Var}(\delta(\kappa^*))$ and hence widening the posterior gap. The chain (38) makes precise the sense in which the cost function $C(\cdot)$—not the information structure—governs equilibrium posterior disagreement and, under the condition in (39), equilibrium trade volume. This is the empirically relevant prediction of the model: two markets with identical information structures but different interpretive complexity should exhibit different trading activity, with the more complex market sustaining wider disagreement and larger positions.

5.6. Formal characterization: processing-stage sufficiency failure and the MS framework

The above intuitions can be formalized more precisely. This subsection develops a formal proposition that specifies exactly which assumption underlying the Milgrom-Stokey theorem is relaxed when processing is costly.

Recall that in Milgrom-Stokey, the key mechanism is that equilibrium prices (or more generally, equilibrium objects) serve as sufficient statistics for payoff-relevant uncertainty. When an agent observes the equilibrium object and conditions on it, she effectively observes the entire information structure that any agent in the economy has access to. Private signals become redundant conditional on prices. This is swamping.

Our framework shows that this sufficiency property depends fundamentally on agents being able to extract the informational content of the observed object without cognitive friction. When such friction is present, the observed object—even if it is theoretically fully informative—no longer serves as a sufficient statistic in practice because agents condition on noisy processed versions of it.

Proposition 3 (Sufficiency failure under costly processing). *Consider a public object Z (e.g., a price, a message, a disclosure) that aggregates payoff-relevant information. Suppose that under frictionless conditioning, Z is a sufficient statistic for the payoff v in the following sense:*

$$\mathbb{E}[v \mid Z, s_i] = \mathbb{E}[v \mid Z] \quad \textit{for all agents } i, \tag{40}$$

where s_i is any private signal. This is the sufficiency property that underlies the Milgrom-Stokey swamping result.

Now introduce costly processing. Suppose agents observe a noisy processed version of Z given by

$$\tilde{Z} = Z + \xi(\kappa), \tag{41}$$

where $\xi(\kappa) \sim \mathcal{N}(0, \kappa^{-1})$ *is independent of all payoff-relevant states, and κ represents processing precision.*

Then, for any finite $\kappa < \infty$*, the processed object* $\tilde{Z}$ *is* not *a sufficient statistic:*

$$\mathbb{E}[v \mid \tilde{Z}, s_i] \neq \mathbb{E}[v \mid \tilde{Z}] \quad \textit{with positive probability.} \tag{42}$$

---

[7] The denominator $D(\kappa)$ is itself decreasing in κ (since $\sigma_B^2(\kappa)$ falls as processing precision rises), which would, in isolation, amplify trade. The stated condition ensures the numerator channel dominates.

Consequently, private signals s_i do not become redundant. Instead, agents condition on a coarser partition than the full information set, and posterior disagreement can persist even though the raw object Z would have unified beliefs under frictionless conditioning.

Proof 3. *This follows directly from Lemma 4. The processed signal $\tilde{Z}$ contains less information than Z because it is contaminated by independent noise $\xi(\kappa)$. Therefore, the posterior mean conditional on $\tilde{Z}$ differs from that conditional on Z:*

$$\mathbb{E}[v \mid Z] \neq \mathbb{E}[v \mid \tilde{Z}] \quad \textit{almost surely for } \kappa < \infty. \tag{43}$$

This immediately implies that $\tilde{Z}$ cannot be sufficient, since a sufficient statistic must contain all relevant information. Moreover, private signals retain information value after conditioning on $\tilde{Z}$, so $\mathbb{E}[v \mid \tilde{Z}, s_i] \neq \mathbb{E}[v \mid \tilde{Z}]$ generically.

The significance of Proposition 3 is that it isolates the exact point at which the no-trade logic breaks: *not* the existence of private information, *not* the logical structure of Bayesian updating, but rather the ability to extract the full informational content of observed objects. Once this ability is limited by processing costs, the public object loses its power to align beliefs.

5.7. Swamping and the cognitive prerequisite

Theorem 3 of Milgrom and Stokey (1983)—the swamping result—is perhaps the most elegant statement in the no-trade literature. It shows that in a competitive equilibrium, the change in prices becomes a sufficient statistic that makes private information redundant. The theorem is powerful precisely because it says: no matter what you privately learned, the price change contains all that matters for your valuation.

Our framework clarifies what this result *requires*: that agents can costlessly extract and condition on the full informational content embedded in the price change. If Z denotes the price change and agents observe it perfectly (frictionlessly), then swamping follows. But when agents observe a processed version $\tilde{Z} = Z + \xi(\kappa)$ with finite κ, the cognitive prerequisite of swamping fails.

The distinction is subtle but important. The claim is not that common priors are violated, that Bayesian updating fails, or that the relevant information is unavailable—in principle, the price change is there for agents to observe. The claim is narrower: the informational content of the price change is not automatically internalized. Agents must expend cognitive effort to interpret it, and when that effort is costly, they optimally choose limited precision. At that limited precision, the price change no longer functions as a sufficient statistic, and the posterior-alignment mechanism of the theorem fails.

Processing frictions therefore do not contradict the Milgrom-Stokey theorem; they *condition* its applicability on an implicit assumption—frictionless information processing—that may fail in markets where information is plentiful but cognitive capacity is scarce.

Definition 2 (The frictionless-conditioning assumption). *The Milgrom-Stokey result implicitly assumes that agents can costlessly extract and condition on the full informational content of observed equilibrium objects. We call this the* frictionless-conditioning assumption. *Formally, it asserts that there is no gap between what an object Z theoretically contains and what an agent actually extracts from observing Z. The agent's effective information after observing Z is the same as the full information content of Z.*

When processing is costly and agents optimally choose finite precision, the frictionless-conditioning assumption fails endogenously. *The agent's effective information becomes a coarsened version of the raw object, and the posterior-alignment machinery of Milgrom-Stokey breaks down.*

To formalize the comparison with the MS framework, consider the following setup. Suppose agents share common priors and receive private signals about a payoff-relevant state. They trade in an equilibrium where prices form. The Milgrom-Stokey conclusion is that if prices are fully revealing (which they can be under certain equilibrium conditions), and if agents can condition on prices costlessly, then posteriors align and trade ceases.

Our result shows what happens when the second condition fails. Let p^* denote the equilibrium price (or price change). Under frictionless conditioning, agents observe p^* exactly and extract all its informational content:

$$\mathcal{I}_i \text{ (frictionless)} = \sigma(s_i, p^*). \tag{44}$$

Under costly processing, agents observe a noisy version $\tilde{p}^* = p^* + \xi(\kappa)$:

$$\mathcal{I}_i \text{ (with processing costs)} = \sigma(s_i, \tilde{p}^*). \tag{45}$$

The information sets differ: $\sigma(s_i, p^*) \neq \sigma(s_i, \tilde{p}^*)$ for finite κ. This difference, though seemingly small, is sufficient to break posterior alignment and allow trade. The magnitude of the effect is determined by the processing cost function $C(\kappa)$: higher costs lead to lower κ^*, wider posterior gaps, and more trade.

Proposition 4 (Processing costs and the Milgrom-Stokey no-trade condition). *Let Z be a payoff-relevant equilibrium object (a price, message, or public disclosure) with $\mathrm{Cov}(v, Z) \neq 0$, and suppose agents share common priors, start from an ex-ante Pareto-efficient allocation, and update Bayesianly.*

(i) ***Frictionless conditioning.*** *Suppose Z is a sufficient statistic for the payoff-relevant state in the sense that $\mathbb{E}[v \mid Z, s_i] = \mathbb{E}[v \mid Z]$ for all agents i, and agents observe Z without processing cost (Definition 2 holds). Then posterior means coincide, and speculative trade cannot be strictly mutually beneficial.*

(ii) ***Costly processing.*** *Suppose agents observe a processed version $\tilde{Z} = Z + \xi(\kappa)$ with $\xi(\kappa) \sim \mathcal{N}(0, \kappa^{-1})$ independent of (v, Z), and $\kappa < \infty$. Then $\tilde{Z}$ is not a sufficient statistic: $\mathbb{E}[v \mid \tilde{Z}, s_i] \neq \mathbb{E}[v \mid \tilde{Z}]$ with positive probability, posterior means generically differ, and speculative trade occurs with positive probability.*

The sole operative difference between the two cases is the processing precision κ. All other conditions—common priors, ex-ante Pareto optimality, and Bayesian updating—remain intact.

Proof 4. *Part (i). Under frictionless conditioning, Z is sufficient by hypothesis: $\mathbb{E}[v \mid Z, s_i] = \mathbb{E}[v \mid Z]$ for all i. Both agents then form the same posterior mean $\mathbb{E}[v \mid Z]$, so $\mu_B - \mu_S = 0$ almost surely. By the Aumann-type argument in Section 2, common-knowledge coincident posteriors under common priors leave no speculative trading motive, and the clearing quantity $x^* = (\mu_B - \mu_S)/D = 0$ almost surely.*

Part (ii). Under $\kappa < \infty$, Lemma 4 establishes that $\mathbb{E}[v \mid Z] \neq \mathbb{E}[v \mid \tilde{Z}]$ almost surely. Because $\xi(\kappa)$ is independent of (v, Z, s_i), conditioning on $\tilde{Z}$ does not recover the information in Z, so private signals s_i retain information value after conditioning on $\tilde{Z}$: $\mathbb{E}[v \mid \tilde{Z}, s_i] \neq \mathbb{E}[v \mid \tilde{Z}]$ with positive probability. Sufficiency fails. Proposition 1 then maps this sufficiency failure into positive-probability trade in the bilateral equilibrium under common priors.

The Milgrom-Stokey no-trade theorem is therefore a conditional statement: *if* agents share common priors *and* can extract equilibrium information costlessly, *then* speculative trade cannot be strictly mutually beneficial. By relaxing the second condition—not the first—trade can persist under common priors. This is a within-model question about information processing, not a rejection of common-prior logic.

6. Discussion, Scope and Relation to the Literature

This section presents the reach of our mechanism, places it in relation to the literature, and highlights the empirical and theoretical implications that follow from interpreting the no-trade result as driven by frictionless processing.

6.1. Reach of our mechanism

Our contribution is both conceptual and mechanism-focused. We re-examine the strength of the no-trade theorem by explicitly modeling a margin that is usually kept implicit: agents' limited ability to process the information contained in prices and in their counterparties' communications. The trading setting in Section 4 demonstrates that when information processing is costly and therefore optimally restricted, posterior beliefs need not almost surely converge, and trade can occur with positive probability even when priors are common.

At the same time, our analysis is deliberately streamlined. We do not construct a full rational-expectations equilibrium with endogenously informative prices and a large number of traders, nor do we incorporate strategic order placement or explicit market-making rules. Our argument differs in scope: even in a bare-bones bilateral trading setting, limited processing can prevent informational objects from serving as sufficient statistics, thereby undermining the posterior-alignment logic that drives no-trade theorems.

6.2. Place in the literature

Our approach is most closely aligned with the interpretation of no-trade, in which common priors and common knowledge underpin agreement. We stress that in market environments,the primary channel through which beliefs may be aligned is frequently the equilibrium price, which, assuming rational expectations, can in principle be fully informative. Processing frictions disrupt the mapping from observables to posteriors: the meet/join structure has less bite when agents' effective information sets are coarser than the partitions generated by the raw observables.

Grossman–Stiglitz-type frameworks yield only partial revelation because full informational efficiency would remove incentives to acquire information. Our focus is different: agents may have access to informative signals but incur costs in interpreting them. This "processing" margin is particularly natural in environments where information is plentiful yet intricate, and the critical friction concerns not the existence of information but whether it is actually internalized.

Kyle-type models produce trade and partial revelation via strategic order choice, market-making, and noise trading. We set aside these institutional details to isolate a distinct mechanism: trade may occur because processing constraints keep prices or other communications from serving as sufficient statistics, even when no strategic trading protocol is present.

Our reduced-form setup with endogenous precision follows the linear-quadratic rational inattention framework of Sims (2003), as outlined in Section 4.3. In that framework, optimal behavior under a Shannon capacity constraint is observationally equivalent to conditioning on the state through a Gaussian channel: the effective information set is a coarsened version of the underlying observables, and the degree of coarsening is determined by the cost structure. This is precisely the representation $\tilde{m}_B = m_{S\to B} + \xi_B$ adopted in the model, where κ^{-1} plays the role of channel noise.

A closely related paper is Matějka (2016), who studies a rationally inattentive seller choosing what information to acquire subject to a full mutual-information constraint. The central finding—that the optimal price distribution has finite discrete support even when the underlying demand state is continuous—reflects the same economic force: when processing is costly, agents optimally condition on coarser representations of the environment. The formal implementations differ. Matějka (2016) works with general priors and an explicit mutual-information constraint, which delivers the discreteness result but at the expense of the closed-form posterior moments that make the no-trade mechanism transparent here. The present paper uses the Gaussian special case of Sims (2003); Matějka (2016) is better understood as establishing the same economic logic in a richer setting, rather than as a direct foundation for the Gaussian representation. The broader rational inattention literature (Sims, 2006) likewise demonstrates that optimal attention allocation produces effective information sets that are strictly coarser than the underlying observables, which is the key property the present model relies on.

The processing-cost channel connects naturally to the literature on rational inattention in financial markets. Peng and Xiong (2006) show that when investors have limited attention and categorize information by asset class rather than processing it stock by stock, comovement and inattention-driven price patterns emerge. Van Nieuwerburgh and Veldkamp (2009) study portfolio underdiversification under information-processing constraints, showing that investors optimally choose to learn about a subset of assets, consistent with the idea that processing costs determine the effective information set rather than the set of available signals. Van Nieuwerburgh and Veldkamp (2010) extend this to an equilibrium setting where information choice and portfolio choice interact, producing home bias as a rational outcome of processing constraints.

These papers share with the present model the idea that the relevant friction is the internalization of available information rather than its existence. The present paper differs in focus: rather than characterizing portfolio or pricing distortions, it identifies the minimal condition under which the no-trade result breaks down.

6.3. *Qualitative implications and empirical relevance*

Although our contribution is primarily theoretical, processing-based interpretation yields qualitative predictions that can be subjected to empirical scrutiny.

For fixed fundamentals, trading volume should be greater when payoff-relevant information is more complex or more difficult to interpret (i.e., entails higher effective processing costs), because posterior disagreements will be less fully arbitraged away.

Markets in which agents differ more in their cognitive or technological processing capabilities should feature more persistent differences in beliefs and, as a result, higher trading activity.

If the key bottleneck is processing (rather than acquisition) of information, then disclosure rules or information design that lower interpretation costs can reduce disagreement and trading volume, even when the quantity of raw information disclosed remains constant.

These qualitative predictions are not exclusive to the processing-cost explanation; models with heterogeneous priors and models with noise traders likewise imply positive trading volume. What is more distinctive to the processing channel is its prediction about how trading volume reacts to changes in disclosure design, holding the amount of disclosed information constant. In standard heterogeneous-priors or differential-information frameworks, trading volume is governed by how information is distributed across agents. In contrast, in the current model, the key determinant is whether agents are able to internalize the information they receive. This leads to a comparative static that the alternative theories do not generate: simplifying the format or presentation of a public disclosure—while leaving its informational content unchanged—should lower trading volume by reducing effective processing costs and thus diminishing posterior disagreement.

Evidence consistent with this prediction appears in the literature on plain-language financial disclosure and on the trading response to earnings announcements of varying complexity. While a structural test of this channel would require a fuller model, the prediction provides a differentiating empirical implication that could, in principle, separate the processing-cost story from its alternatives.

6.4. *Limitations and natural extensions*

Two limitations merit emphasis. First, the illustrative trading environment assumes that transaction prices clear a bilateral market conditional on agents' beliefs; it does not endogenize the informational role of prices through a complete equilibrium of rational-expectations with a large population of traders. Second, we capture processing frictions using an endogenous-precision technology; more elaborate rational-inattention formulations (for instance, explicit mutual-information constraints) could be incorporated, but only at the expense of heavier notation and reduced clarity.

These limitations naturally suggest several extensions. One is to embed processing costs in an equilibrium setting with endogenously informative prices (e.g., with a market maker or in an auction framework), allowing for a sharper comparison with the MS "full revelation" benchmark. Another is to analyze heterogeneous or state-contingent processing costs, which can produce endogenous time variation in belief dispersion and trading volume.

A third extension—and the one most directly relevant to Comment 6 of the referee—is to relax the one-sided processing assumption and allow the seller also to choose a processing precision κ_S, so that both agents simultaneously optimize over (κ_B, κ_S). In this two-sided environment, both agents condition on processed signals $\tilde{m}_B = m_{S\to B} + \xi_B$ and $\tilde{m}_S = m_{B\to S} + \xi_S$ with $\xi_i \sim \mathcal{N}(0, \kappa_i^{-1})$. The posterior-mean gap becomes

$$\mu_B(\kappa_B) - \mu_S(\kappa_S) = [\mathbb{E}[v \mid s_B, \tilde{m}_B]] - [\mathbb{E}[v \mid s_S, \tilde{m}_S]], \tag{46}$$

which is a non-degenerate Gaussian random variable for any finite pair (κ_B, κ_S) under Assumption 2, since the two processing noise terms ξ_B and ξ_S contribute to disagreement independently of one another. The qualitative result—generic posterior disagreement and positive-probability trade under common priors—therefore survives the two-sided generalization. What changes is the comparative static: equilibrium trade volume is now a function of both cost structures $C_B(\kappa_B)$ and $C_S(\kappa_S)$, and the relevant policy experiment is a joint reduction in processing costs that narrows disagreement from both sides. The one-sided model studied in the text is the limiting case $\kappa_S \to \infty$ (or equivalently $C_S \equiv 0$), which is most natural when one counterpart has substantially lower processing costs than the other—as, for instance, when a dealer or institutional investor faces a retail or non-specialist counterpart in an over-the-counter market. We leave the full two-sided equilibrium analysis to future work.

6.5. *Summary and limitations*

The no-trade theorem continues to serve as a central reference point. We contend that its empirical relevance hinges on an implicit cognitive assumption: agents must be able to costlessly infer and condition on the informational content of market signals. Once this information-processing margin is made explicit and modeled endogenously, posterior beliefs need not coincide, and trade can arise in a stripped-down setting without relaxing the common-prior assumption. This offers a straightforward way to connect no-trade reasoning with the trading activity actually observed in markets where information is abundant but attention and interpretive capacity are limited.

7. Conclusion

This paper re-examines the no-trade theorem by foregrounding an assumption that is typically left implicit: agents are assumed to process the informational content of observed market signals at zero cost. The power of the theorem rests on two steps: i) agents observe a public signal (such as a price movement or a message), and ii) they then costlessly update their beliefs based on it, thereby bringing their posteriors into alignment. The paper demonstrates that this second step is itself costly, and that these processing costs have implications for equilibrium outcomes.

The mechanism is developed in a minimal two-stage bilateral environment. In Stage 1, the buyer commits to a processing precision κ^* by equating marginal processing cost with marginal certainty-equivalent value. In Stage 2, markets clear bilaterally given the resulting information sets. Three results follow. Lemma 2 shows that under finite κ and non-degenerate information, posterior means differ almost surely. Lemma 3 establishes that the no-trade benchmark requires the joint limit $(\kappa, \tau) \to (\infty, \infty)$: eliminating processing noise alone does not restore posterior coincidence when messages are noisy. Proposition 1 characterizes the equilibrium and identifies the processing cost as the determinant of disagreement magnitude rather than trade existence.

Two conceptual clarifications emerge from this analysis. First, the mere presence or absence of trade is driven by the primitive information structure—distinct information sets generically lead to distinct posterior beliefs—whereas the cost of processing information determines the equilibrium dispersion of posterior beliefs and, under transparent risk-capacity conditions, the equilibrium volume of trade. Second, the paper's connection to the Milgrom–Stokey swamping result (Theorem 3 of MS) is complementary rather than contradictory: MS shows that swamping occurs when agents can costlessly condition on price movements in a competitive equilibrium. This paper identifies the cognitive requirement underlying that result—frictionless information processing—and demonstrates that when this requirement fails endogenously, the public signal no longer serves as a sufficient statistic.

The channel is modeled using the linear–quadratic rational inattention framework of Sims (2003), which yields closed-form expressions for posterior moments in the Gaussian setting. More general specifications, such as those in Matějka (2016) and Sims (2006), could be adopted, but only at the expense of analytical tractability. Natural directions for future work include embedding information-processing costs into a full rational-expectations equilibrium with endogenously informative prices, studying heterogeneity in processing capacity across agents, and confronting the model's distinctive empirical implication—that, holding content fixed, the format of disclosure should influence trading volume—with data on plain-language regulation and trading behavior around earnings announcements.

The broader takeaway is that no-trade outcomes depend not only on how information is structured, but also on the cognitive standard that determines how that information is absorbed. Assuming frictionless information processing is demanding in markets where signals are plentiful but attention and interpretive capacity are limited. Making this margin explicit and treating it as an endogenous decision offers a tractable link between the no-trade benchmark and the actual trading activity seen in today's information-dense markets.

Data availability statement

No datasets were generated or analyzed during the current study.

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies

During the preparation of this work, the author used AI-assisted tools for language editing and formatting suggestions. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Matějka, F. (2016). Rationally inattentive seller: Sales and discrete pricing. *Review of Economic Studies*, 83(3):1125–1155.

Milgrom, P. and Stokey, N. (1982). Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26:17–27.

Peng, L. and Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80(3):563–602.

Samuelson, L. (2004). Modeling knowledge in economic analysis. *Journal of Economic Literature*, 42(2):367–403.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Sims, C. A. (2006). Rational Inattention: Beyond the Linear-Quadratic Case. *American Economic Review*, 96(2):158–163.

Van Nieuwerburgh, S. and Veldkamp, L. (2009). Information immobility and the home bias puzzle. *Journal of Finance*, 64(3):1187–1215.

Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information acquisition and under-diversification. *Review of Economic Studies*, 77(2):779–805.

Appendix A. Posterior Coincidence and the Collapse to No-Trade Logic

This appendix formalizes two complementary results about posterior coincidence. Section Appendix A.2 shows that posterior coincidence implies no speculative trade, establishing the reduced-form no-trade benchmark. Section Appendix A.5 characterizes the conditions under which the frictionless-processing limit $\kappa \to \infty$ achieves posterior coincidence and shows that, in general, both processing noise and message noise must vanish. Together, these results anchor the claim in Lemma 3 and Remark 1 that the joint limit $(\kappa, \tau) \to (\infty, \infty)$ is required for the model to recover the no-trade benchmark.

Appendix A.1. Setup

Consider the bilateral risky-asset trade environment in Section 4. Under CARA preferences and conditional normality, each agent $i \in \{B, S\}$ has (interior) optimal demand

$$x_i^*(p) \;=\; \frac{\mu_i - p}{\gamma_i \, \sigma_i^2}, \tag{A.1}$$

where (μ_i, σ_i^2) are the agent's posterior mean and variance of v given the respective information set.

Appendix A.2. Posterior coincidence implies no speculative trade

Define posterior coincidence as

$$\mathbb{P}\Big(\mu_B = \mu_S \text{ and } \sigma_B^2 = \sigma_S^2\Big) = 1. \tag{A.2}$$

The condition (A.2) is the formal way to state that both sides evaluate the asset identically in all states of the world (the same posterior distribution over v).

Under (A.2), let μ and σ^2 denote the common posterior mean and variance (random variables measurable with respect to the common posterior). Consider the price $p = \mu$. Substituting into (A.1) yields

$$x_B^*(\mu) = \frac{\mu - \mu}{\gamma_B \, \sigma^2} = 0 \qquad \text{and} \qquad x_S^*(\mu) = \frac{\mu - \mu}{\gamma_S \, \sigma^2} = 0. \tag{A.3}$$

Hence, at the price equal to the common conditional expectation of the payoff, neither side wishes to take a speculative position. More generally, when posteriors coincide, both agents share the same valuation and disagree neither about the expected payoff nor about residual risk; therefore, any non-zero position would have to be justified purely by exogenous hedging needs or constraints unrelated to information. In the absence of such additional motives, the informational channel for trade shuts down.

Appendix A.3. Connection to no-trade logic

The classic no-trade logic hinges on the observation that, under rational expectations and common knowledge of rationality, equilibrium prices can become (or approximate) sufficient statistics for the payoff. When that happens, conditioning on the price leads agents' posteriors to coincide, and the scope for mutually beneficial speculative trade vanishes. Condition (A.2) captures precisely this degeneracy in reduced form: once posterior coincidence holds almost surely, the model collapses back to the no-trade benchmark along the informational margin.

Appendix A.4. Why the paper imposes posterior non-coincidence

Imposing $\mathbb{P}(\mu_B \neq \mu_S) > 0$ (or, more strongly, $\mathbb{P}((\mu_B, \sigma_B^2) \neq (\mu_S, \sigma_S^2)) > 0$) is therefore a minimal non-degeneracy condition that ensures there exists a set of states with genuine posterior disagreement. This is what allows the clearing price to balance non-zero demand and supply, generating trade with positive probability in the illustrative environment.

Appendix A.5. Frictionless limit and the conditions for posterior coincidence

Lemma 3 in the main text provides the full formal statement and proof. This subsection records the key sufficient condition for posterior coincidence for reference.

Remark 3 (Sufficient condition for no-trade in the limit). *Under the Gaussian primitives of equations* (1)—(4)*, posterior coincidence* $\mathbb{P}(\mu_B = \mu_S) = 1$ *holds if and only if* $\mathcal{I}_B = \mathcal{I}_S$ *almost surely. That is, both agents condition on the same sufficient statistic for* v*. In the bilateral model, this requires:*

(i) $\kappa \to \infty$ *(no processing noise), so that* $\tilde{m}_B \to m_{S\to B}$ *almost surely; and*

(ii) $\tau \to \infty$ *(noiseless messages), so that* $m_{S\to B} \to s_S$ *and* $m_{B\to S} \to s_B$ *almost surely.*

Under conditions (i) and (ii) together, $\mathcal{I}_B(\infty) = \mathcal{I}_S = \sigma(s_S, s_B)$*, and the two agents share the same information set, implying* $\mu_B = \mu_S$ *almost surely and hence* $x^* = 0$ *almost surely. Neither condition alone is sufficient:* $\kappa \to \infty$ *with* $\tau < \infty$ *leaves residual message noise, while* $\tau \to \infty$ *with* $\kappa < \infty$ *leaves residual processing noise.*

Appendix B. Derivation of the Interior Processing Condition

This appendix derives the first-order condition characterizing an interior Stage 1 processing choice $\kappa^* > 0$, as stated in Proposition 1. The buyer's Stage 1 problem takes the equilibrium mapping $\kappa \mapsto p^*(\kappa)$ — induced by Stage 2 clearing — as given, and maximizes a one-dimensional indirect certainty-equivalent objective. The derivation proceeds in three steps: (i) reduce the buyer's problem to a scalar choice via the certainty-equivalent representation; (ii) differentiate through the equilibrium price mapping using the chain rule; and (iii) characterize the interior optimum via the condition $C'(\kappa^*) = \Delta(\kappa^*)$.

Appendix B.1. From the trading problem to an indirect objective in κ

Fix any processing level $\kappa \geq 0$. Given a transaction price p and an information set $\mathcal{I}_B(\kappa)$, the buyer chooses a position $x \in [0, q_0]$ to maximize the net expected utility of the processing cost:

$$\max_{x\in[-q_0,q_0]} \mathbb{E}[-\exp(-\gamma_B x(v-p)) \,|\, \mathcal{I}_B(\kappa)] \; - \; C(\kappa). \tag{B.1}$$

Let $\mathcal{V}(p, \kappa)$ denote the maximized value (utility) before subtracting $C(\kappa)$:

$$\mathcal{V}(p,\kappa) \equiv \max_{x\in[-q_0,q_0]} \mathbb{E}\left[-\exp(-\gamma_B x(v-p)) \,|\, \mathcal{I}_B(\kappa)\right] \tag{B.2}$$

In equilibrium, the transaction price clears the bilateral market, resulting in a (measurable) mapping $\kappa \mapsto p^*(\kappa)$ via (11). The buyer's processing choice can therefore be written as a problem with one dimension

$$\max_{\kappa\geq 0} \Phi(\kappa) \; \equiv \; \mathbb{E}\Big[\mathcal{V}(p^*(\kappa),\kappa)\Big] \; - \; C(\kappa), \tag{B.3}$$

where the outer expectation is taken over the equilibrium-relevant random variables.

Appendix B.2. Certainty-equivalent representation under CARA-normality

Under the conditional normality of v given $\mathcal{I}_B(\kappa)$, the buyer's posterior is

$$v \mid \mathcal{I}_B(\kappa) \sim \mathcal{N}(\mu_B(\kappa), \sigma_B^2(\kappa)).$$

For any fixed (p, κ), the buyer's (unconstrained) optimal position is the standard CARA-normal demand:

$$x_B^*(p,\kappa) \ = \ \frac{\mu_B(\kappa) - p}{\gamma_B \, \sigma_B^2(\kappa)}, \tag{B.4}$$

which is truncated if it falls outside $[0, q_0]$. When the optimum is interior, substituting $x_B^*(p,\kappa)$ into the objective yields a certainty-equivalent gain from trade (up to an additive constant independent of κ):

$$CE(p,\kappa) \ = \ \frac{(\mu_B(\kappa) - p)^2}{2\,\gamma_B \, \sigma_B^2(\kappa)}. \tag{B.5}$$

Thus, in the interior region, maximizing the expected utility is equivalent to maximizing the expected certainty equivalent.

Appendix B.3. Exact marginal value of precision

Define the marginal value of precision as

$$\Delta(\kappa) \ \equiv \ \frac{d}{d\kappa}\,\mathbb{E}[\,CE(p^*(\kappa),\kappa)]\,, \tag{B.6}$$

where $p^*(\kappa)$ is the clearing price induced by κ and $CE(p,\kappa)$ is the maximized certainty equivalent in (B.5). Differentiating $CE(p^*(\kappa),\kappa)$ yields

$$\Delta(\kappa) = \mathbb{E}\left[\frac{\mu_B(\kappa) - p^*(\kappa)}{\gamma_B\,\sigma_B^2(\kappa)}(\mu_B'(\kappa) - p^{*\prime}(\kappa)) - \frac{(\mu_B(\kappa) - p^*(\kappa))^2}{2\,\gamma_B\,(\sigma_B^2(\kappa))^2}\,(\sigma_B^2)'(\kappa)\right], \tag{B.7}$$

which captures both the direct effect of processing on posterior moments and the equilibrium-mediated effect through the price response $p^{*\prime}(\kappa)$.

Remark 4 (Risk-reduction sufficient statistic). *If the equilibrium price response term is negligible (e.g., under a partial-equilibrium view in which $p^*(\kappa)$ is treated as locally fixed, or under conditions implying $\mu_B'(\kappa) \approx p^{*\prime}(\kappa)$), then* (B.7) *is well-approximated by the risk-reduction channel:*

$$\Delta(\kappa) \ \approx \ -\frac{1}{2}\,\frac{\partial}{\partial\kappa}\Big(\gamma_B\,\sigma_B^2(\kappa)\Big)^{-1} \cdot \mathbb{E}\Big[(\mu_B(\kappa) - p^*(\kappa))^2\Big]. \tag{B.8}$$

We use (B.8) *only as an interpretive device; the definition of* $\Delta(\kappa)$ *through* (B.6) *is exact.*

Appendix B.4. Envelope argument and the marginal value of precision

Consider the composition $CE(p^*(\kappa),\kappa)$. Differentiating with respect to κ and using (B.5) gives

$$\frac{d}{d\kappa}CE(p^*(\kappa),\kappa) = \frac{\partial CE}{\partial \mu_B}\,\mu_B'(\kappa) + \frac{\partial CE}{\partial \sigma_B^2}\,(\sigma_B^2)'(\kappa) + \frac{\partial CE}{\partial p}\,p^{*\prime}(\kappa), \tag{B.9}$$

$$\frac{\partial CE}{\partial \mu_B} = \frac{\mu_B(\kappa) - p^*(\kappa)}{\gamma_B\,\sigma_B^2(\kappa)}, \qquad \frac{\partial CE}{\partial \sigma_B^2} = -\frac{(\mu_B(\kappa) - p^*(\kappa))^2}{2\,\gamma_B\,(\sigma_B^2(\kappa))^2}, \qquad \frac{\partial CE}{\partial p} = -\frac{\mu_B(\kappa) - p^*(\kappa)}{\gamma_B\,\sigma_B^2(\kappa)}. \tag{B.10}$$

Combining these expressions yields

$$\frac{d}{d\kappa}CE(p^*(\kappa),\kappa) = \frac{\mu_B(\kappa) - p^*(\kappa)}{\gamma_B\,\sigma_B^2(\kappa)}\Big(\mu_B'(\kappa) - p^{*\prime}(\kappa)\Big) \;-\; \frac{(\mu_B(\kappa) - p^*(\kappa))^2}{2\,\gamma_B\,(\sigma_B^2(\kappa))^2}\,(\sigma_B^2)'(\kappa). \tag{B.11}$$

Define the marginal value of precision as the derivative of the expected certainty-equivalent gain with respect to κ:

$$\Delta(\kappa) \;\equiv\; \frac{d}{d\kappa}\,\mathbb{E}\Big[\,CE(p^*(\kappa),\kappa)\Big]. \tag{B.12}$$

The expression (B.11) makes clear that $\Delta(\kappa)$ captures both (i) the direct effect of processing on posterior moments (μ_B, σ_B^2) and (ii) the equilibrium-mediated effect through the price response $p^{*\prime}(\kappa)$.

For exposition purposes, it is often useful to highlight the risk-reduction channel implied by the Gaussian precision choice. When the effect of κ operates primarily through $\sigma_B^2(\kappa)$ (or when $\mu_B'(\kappa) - p^{*\prime}(\kappa)$ is second-order), a convenient sufficient-statistical approximation to (B.12) is provided.

$$\Delta(\kappa) \approx -\,\frac{1}{2}\,\frac{\partial}{\partial\kappa}\Big(\gamma_B\,\sigma_B^2(\kappa)\Big)^{-1}\cdot\mathbb{E}\Big[\big(\mu_B(\kappa) - p^*(\kappa)\big)^2\Big], \tag{B.13}$$

using $(\sigma_B^2)'(\kappa) < 0$ whenever the processed message contains information.

Appendix B.5. Interior FOC: necessity and sufficiency

The assumptions on the cost of processing the information imply that $C(\kappa)$ is differentiable for $\kappa > 0$ and convex. Consider the indirect objective (B.3). If the optimal processing choice is interior at $\kappa^* > 0$, then the first-order condition is

$$\Phi'(\kappa^*) = 0 \quad\Longleftrightarrow\quad \frac{d}{d\kappa}\mathbb{E}\Big[CE(p^*(\kappa),\kappa)\Big]\Big|_{\kappa=\kappa^*} = C'(\kappa^*). \tag{B.14}$$

By the definition of $\Delta(\kappa)$ in (B.12), (B.14) is equivalent to

$$C'(\kappa^*) \;=\; \Delta(\kappa^*), \tag{B.15}$$

This is the interior processing condition reported in the main text.

The condition is also sufficient for optimality whenever $\Phi(\kappa)$ is concave in a neighborhood of κ^*. A convenient sufficient condition is

$$\Phi''(\kappa^*) = \Delta'(\kappa^*) - C''(\kappa^*) < 0, \quad \text{i.e.} \quad C''(\kappa^*) > \Delta'(\kappa^*). \tag{B.16}$$

Under strict concavity, the interior solution is unique.

Appendix C. Proof of Lemma 2 (Posterior Disagreement is Generic)

The following proof establishes that, under finite processing precision and non-degenerate information, the posterior-mean gap is a non-degenerate Gaussian. Combined with Lemma 3, which characterizes the limiting case, this shows that disagreement is a generic feature of the model for all $(\kappa, \tau) \in (0, \infty)^2$.

Proof 5. *Under the Gaussian primitives described in equations* (1), (3) *and* (4), *the vector of observables entering each agent's information set is jointly Gaussian with* v. *In Gaussian environments, conditional expectations are affine in the conditioning variables. Hence there exist constants* $a_S \in \mathbb{R}$, *vectors* $b_S \in \mathbb{R}^{k_S}$, $a_B(\kappa) \in \mathbb{R}$, *and* $b_B(\kappa) \in \mathbb{R}^{k_B}$ *such that*

$$\mu_S \equiv \mathbb{E}[v \mid \mathcal{I}_S] = a_S + b_S^\top Y_S, \qquad \mu_B(\kappa) \equiv \mathbb{E}[v \mid \mathcal{I}_B(\kappa)] = a_B(\kappa) + b_B(\kappa)^\top Y_B(\kappa), \tag{C.1}$$

where Y_S *collects the seller's signals/messages, and* $Y_B(\kappa)$ *collects the buyer's signals/messages after processing. In particular,* Y_S *and* $Y_B(\kappa)$ *are jointly Gaussian.*

Therefore, the posterior mean gap is itself affine in a jointly Gaussian vector:

$$\mu_B(\kappa) - \mu_S = (a_B(\kappa) - a_S) + b_B(\kappa)^\top Y_B(\kappa) - b_S^\top Y_S.$$

It follows that $\mu_B(\kappa) - \mu_S$ *is normally distributed.*

It remains to show non-degeneracy. Under Assumption 2, the buyer does not perfectly recover the seller's sufficient statistic: either the message noise is nonzero ($\tau_B < \infty$) *and/or processing noise is nonzero* ($\kappa < \infty$)*. Combined with the condition that the two information sets are not identical in the Gaussian sense (e.g.,* $\sigma_S^2 \neq \sigma_B^2$ *or, more generally, distinct sufficient statistics enter* Y_S *and* $Y_B(\kappa)$*), there exists at least one Gaussian component that enters the affine representation of* μ_S *but is not perfectly spanned by the components entering* $\mu_B(\kappa)$ *(or vice versa). Consequently,*

$$\mathrm{Var}(\mu_B(\kappa) - \mu_S) > 0,$$

so the distribution is nondegenerate. This implies $\mathbb{P}(\mu_B(\kappa) \neq \mu_S) = 1$ *and, by symmetry and continuity of the normal distribution,* $\mathbb{P}(\mu_B(\kappa) > \mu_S) \in (0, 1)$.