# Polynomial Welfare Decompositions Fail for Single-Moving-Entry Belief Mis-Specifications

[Author]*

April 23, 2026

**Abstract**

In linear-Gaussian Kalman-filter problems with a single-moving-entry equilibrium-level (EL) belief mis-specification – where a scalar bias parameter enters the perceived signal covariance through exactly one entry, rationally with a linear denominator – the true-measure welfare function is a rational function of the bias with a squared linear denominator, and generically not a polynomial. This structural rational-versus-polynomial result (Theorem 1) follows from Cramer's rule applied to a rank-one perturbation. Any corresponding *decision-level* (DL) specification – one in which the agent's action is a convex combination of correct $\mathbb{P}$-projections – admits a clean polynomial welfare decomposition by the tower property. The paper works out three applications: cursed managerial disclosure in an Eyster-Rabin-Vayanos-consistent noisy-REE (Application 1, derived in full); partial-equilibrium thinking in the sense of Bastianello and Fontanier (2025) (Application 2, conjecture-level derivation); and managerial overconfidence about own-signal precision (Application 3, reduced-form primitive). The result is a structural template: modelers seeking polynomial welfare decompositions under single-moving-entry EL primitives cannot have them; the polynomial form is the tower-property counterpart available only at the decision stage.

**Keywords:** Kalman filter, belief mis-specification, welfare decomposition, cursedness, partial-equilibrium thinking, overconfidence.
**JEL:** D82, D83, G14, G18.

*[Affiliations and acknowledgments]

# 1 Introduction

Behavioral finance and information economics routinely embed a scalar bias parameter into a linear-Gaussian Kalman-filtering problem: a cursed trader attenuates the informational content of the price (Eyster et al., 2019), a partial-equilibrium-thinking agent partially ignores the aggregate mapping (Bastianello and Fontanier, 2025), an overconfident manager inflates the precision of her own signal. In each case, the bias operates at the belief-formation stage: the agent forms a single Kalman posterior under a mis-specified perceived covariance and acts on that posterior. A natural tractability question is whether the resulting welfare function admits a clean polynomial decomposition in the bias parameter, an attenuated price-channel identity of the form $W = W_{\text{prior}} - f(\chi)W_{\text{diff}} + \text{cost}$ for some polynomial $f$. The question matters because the answer determines whether a modeler can work in closed form or must carry rational functions through disclosure-design, calibration, and policy exercises.

This paper answers the question in the negative for a broad class of belief-stage misspecifications. The main result (Theorem 1) is a structural rational-versus-polynomial characterization: when a scalar bias $\chi$ enters the perceived signal covariance $\Sigma^{\text{perc}}_{\text{sig}}$ through a *single moving entry*, rationally in a transformed variable $u = g(\chi)$ with a linear denominator, and the perceived cross-covariance $\Sigma^{\text{perc}}_{\theta,\text{sig}}$ is $\chi$-free, the perceived Kalman weights are rational in $u$ with a common linear denominator $D(u)$, and the true-measure mean-squared error is rational in $u$ with denominator $D(u)^2$. Generically, $D(u)^2$ does not divide the numerator, so the welfare function is not a polynomial in $\chi$ on a nonempty open set of primitive configurations.

The mechanism is transparent. Cramer's rule applied to a 3x3 (or higher) matrix in which one entry is linear in $u$ and the rest are constant produces cofactors that are at most linear in $u$ and a determinant that is linear in $u$. The resulting Kalman weights are therefore linear-over-linear in $u$, and squaring them to form the true-measure MSE produces the $D(u)^2$ denominator. Whether the numerator is divisible by $D(u)^2$ is a single polynomial condition on primitives, which cuts out a closed set of measure zero; on the open complement, the welfare function is strictly rational, not polynomial.

The obstruction is non-obvious despite the short proof. Nobody in the behavioral-finance literature has predicted that EL-style welfare functions would be polynomial, but nobody has ruled it out either; the literature typically works at $\chi = 0$ or computes specific equilibrium objects without asking the decomposition question. A modeler confronting a new EL primitive would plausibly conjecture, by analogy with the attenuated price-channel form available at $\chi = 0$, that a polynomial-in-$\chi$ decomposition extends to the interior. Theorem 1 shows that conjecture fails structurally for any single-moving-entry EL primitive, regardless of the underlying behavioral story.

As a foil, the paper records that the tower property delivers a clean polynomial decomposition for any *decision-level* (DL) specification (one in which the agent's action is a convex-combination-in-$\chi$ of two or more correct $\mathbb{P}$-projections of $\theta$). DL is not a behavioral model; it is the polynomial-achieving algebraic counterpart that any single-moving-entry EL primitive would fail to match. The DL form is useful for two reasons: (i) it locates the precise algebraic object that the EL primitive cannot reproduce, and (ii) it gives a tractable approximation whose bounded deviation from EL is informative about how much tractability costs.

The paper works out three applications. Application 1 (cursed managerial disclosure). In the Goldstein and Yang (2019) disclosure environment with a manager cursed in the sense of Eyster et al. (2019), the cursedness parameter $\kappa_M$ enters the perceived signal covariance through the single $(3,3)$ entry $\sigma_{33} = V_0 + 1/[(1-\kappa_M)^2\tau_z]$. This satisfies the single-moving-entry template with $g(\kappa_M) = (1-\kappa_M)^2$; Theorem 1 applies. I derive the cursed-disclosure consequence in full: the rational structure of the Kalman weights, the open-set non-divisibility at the canonical symmetric scenario, and the numerical magnitudes of the DL-EL gap across calibrated scenarios. Applications 2–3 (PET; overconfidence). I describe how partial-equilibrium thinking and managerial overconfidence about own-signal precision plausibly fit the template, but flag these applications as conjecture-level: I defer rigorous verification that each primitive reduces to a single moving entry of $\Sigma^{\text{perc}}_{\text{sig}}$ in a fully-worked equilibrium setup.

The paper positions against two anchors. Eyster et al. (2019) fix the belief-stage convention for cursedness that my Application 1 specializes. Goldstein and Yang (2019) give the $\kappa_M = 0$ rational Bayesian benchmark whose polynomial-in-$\kappa_M$ extension Theorem 1 rules out. Bastianello and Fontanier (2025) give an EL primitive for traders reading prices that Application 2 conjectures fits the template. de Clippel and Zhang (2022) analyze non-Bayesian receivers in an abstract persuasion setting whose primitive lies outside the Kalman-filter class. Banerjee et al. (2025, 2026); Ostrizek and Sartori (2024) study adjacent settings without raising the polynomial-decomposition question.

The contribution is a two-page Cramer's-rule-plus-tower-property observation about rank-one perturbations of a Kalman inverse covariance. The observation is short, but the paper's structure adds value beyond a single application: by abstracting to the single-moving-entry class, the result converts what would otherwise be a series of one-off negative checks into a portable template with a clear sufficient condition. Modelers working in the linear-Gaussian behavioral-Kalman class can read off from the primitive's algebraic structure whether a polynomial welfare decomposition is available, without deriving the welfare function.

Section 2 sets up the linear-Gaussian Kalman framework and the disclosure environment used in Application 1. Section 3 states and proves Theorem 1 (the general single-moving-entry obstruction) and records the DL polynomial counterpart. Section 4 works out Application 1 (cursed managerial disclosure) in full, including numerics, and describes Applications 2–3 at the conjecture level. Section 5 gives a closed-form welfare-loss identity at the DL counterpart, conditional on Assumption 1, with a brief diagnostic on disclosure-rank reversal. Section 6 discusses limitations. Section 7 concludes. Appendices collect the REE existence proof (A), auxiliary lemmas (B), and supporting numerical output (C).

# 2 Setup

This section fixes two layers. Section 2.1 gives the general linear-Gaussian Kalman-filter layer at which the main theorem operates. Section 2.2 fixes the cursed-managerial-disclosure environment used in Application 1; Section 2.3 defines the equilibrium-level and decision-level specifications in that environment. Readers interested only in the general result can stop after Section 2.1.

## 2.1 The Kalman-filter layer: single-moving-entry EL primitives

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and consider an agent who observes an $n$-vector of signals $s \in \mathbb{R}^n$ and computes a linear estimate of a latent $\theta \in \mathbb{R}^d$. Under the true joint distribution, $(\theta, s)$ is jointly Gaussian with covariance blocks $\Sigma_{\text{sig}}$ and $\Sigma_{\theta,\text{sig}}$; under the agent's (mis-specified) perceived distribution, the analogous blocks are $\Sigma^{\text{perc}}_{\text{sig}}(\chi)$ and $\Sigma^{\text{perc}}_{\theta,\text{sig}}(\chi)$, parameterized by a scalar bias $\chi \in [0, 1)$. The perceived Kalman weight vector is

$$b(\chi)^\top := \Sigma^{\text{perc}}_{\theta,\text{sig}}(\chi)\, \Sigma^{\text{perc}}_{\text{sig}}(\chi)^{-1}. \tag{1}$$

The agent's action is $b(\chi)^\top s$. Welfare is the true-measure mean-squared error $W(\chi) := \mathbb{E}_{\mathbb{P}}[(b(\chi)^\top s - L\theta)^2]$ for a fixed linear functional $L : \mathbb{R}^d \to \mathbb{R}$, plus any $\chi$-independent cost term.

**Definition 1** (Single-moving-entry EL primitive)**.** The EL primitive $(\Sigma^{\text{perc}}_{\text{sig}}, \Sigma^{\text{perc}}_{\theta,\text{sig}})$ is *single-moving-entry* if there exist indices $(i^*, j^*) \in \{1, \ldots, n\}^2$ and a polynomial change of variable $u = g(\chi)$, with $g$ a polynomial in $\chi$, such that:

(a) $(\Sigma^{\text{perc}}_{\text{sig}}(\chi))_{i^*j^*}$ and $(\Sigma^{\text{perc}}_{\text{sig}}(\chi))_{j^*i^*}$ are rational in $u$ with denominator linear in $u$, while every other entry of $\Sigma^{\text{perc}}_{\text{sig}}(\chi)$ is independent of $\chi$;

(b) $\Sigma^{\text{perc}}_{\theta,\text{sig}}(\chi)$ is independent of $\chi$.

Condition (a) is the algebraic footprint of a belief-stage mis-specification that attenuates the precision (or inflates the variance) of one signal in isolation. Condition (b) says that the agent's mental model of how $\theta$ correlates with the signal vector is itself unbiased; only the agent's model of how the signals correlate with each other is wrong, and wrong in exactly one place.

The decision-level counterpart of an EL primitive is defined algebraically, not behaviorally:

**Definition 2** (Decision-level counterpart)**.** A *decision-level (DL) counterpart* to an EL primitive is any action rule of the form $k_{\text{DL}}(\chi) := \sum_\ell \lambda_\ell(\chi)\, m_\ell$, where each $m_\ell := \mathbb{E}_{\mathbb{P}}[L\theta \mid \mathcal{G}_\ell]$ is a correct $\mathbb{P}$-projection of $L\theta$ onto a sub-sigma-field $\mathcal{G}_\ell \subseteq \sigma(s)$, and the mixing weights $\lambda_\ell(\chi)$ are polynomials in $\chi$ summing to one at every $\chi$ and matching the EL action at the endpoints $\chi \in \{0, 1\}$.

DL is not a behavioral model of the agent; it is the polynomial-in-$\chi$ object that a single-moving-entry EL primitive would fail to match. When I speak of "the DL counterpart" in the rest of the paper, I mean the unique two-posterior mixture that matches the EL endpoints in the cursed-disclosure application of Section 2.2; abstractly, any polynomial mixture of correct $\mathbb{P}$-projections inherits the tower-property polynomial decomposition.

## 2.2 Application-1 environment: cursed managerial disclosure

The cursed-disclosure application uses a single-period noisy rational-expectations market with public disclosure, informed trading, and managerial feedback to real investment, in the Goldstein and Yang (2019) tradition. Readers interested in the general theorem (Section 3) can skip this subsection; it is the concrete setup in which Application 1 lives.

**Fundamentals and disclosure.** The asset payoff is a sum of two independent components,

$$\theta = \theta_A + \theta_B, \qquad \theta_A \sim \mathcal{N}(0, 1/\tau_{\theta A}), \qquad \theta_B \sim \mathcal{N}(0, 1/\tau_{\theta B}), \tag{2}$$

with $\tau_{\theta A}, \tau_{\theta B} > 0$. The planner chooses disclosure precisions $(\tau_A, \tau_B) \in \mathbb{R}^2_+$ and releases

$$y_X = \theta_X + \varepsilon_X, \qquad \varepsilon_X \sim \mathcal{N}(0, 1/\tau_X), \qquad X \in \{A, B\}, \tag{3}$$

with $(\varepsilon_A, \varepsilon_B)$ independent of each other and of all other primitives. The two-component structure is minimal but load-bearing for the paper's welfare trade-off (Lemma 2).

**Traders and noisy-REE.** A continuum $i \in [0,1]$ of CARA informed traders with risk aversion $a > 0$ observes a private signal

$$s_i = \theta + \eta_i, \qquad \eta_i \sim \mathcal{N}(0, 1/\alpha),$$

i.i.d. across traders and independent of the public signals and of noise trade $z \sim \mathcal{N}(0, 1/\tau_z)$. Trader cursedness $\kappa_T \in [0,1]$ enters through the Eyster et al. (2019) attenuation of the price-channel loading: a trader treats the price as if its loading vector $\rho \in \mathbb{R}^2_+$ were $(1-\kappa_T)\rho$. The cursed-trader information matrix is

$$\Lambda(\kappa_T) = \operatorname{diag}(\tau_{\theta A} + \tau_A, \tau_{\theta B} + \tau_B) + \alpha\, \mathbf{1}\mathbf{1}^\top + (1-\kappa_T)^2 \tau_z \rho\rho^\top. \tag{4}$$

Market clearing produces a linear noisy-REE whose informational reduced form is the scalar price signal

$$\tilde{p} = \rho_A^* \theta_A + \rho_B^* \theta_B + \tilde{z}, \qquad \tilde{z} \sim \mathcal{N}(0, 1/\tau_z). \tag{5}$$

**Lemma 1** (Existence of a linear noisy-REE)**.** *For every admissible primitive vector* $(\tau_A, \tau_B, \tau_{\theta A}, \tau_{\theta B}, \alpha, \tau_z, a)$ *and each* $\kappa_T \in [0,1]$*, there exists a linear noisy-REE characterized by a fixed point* $\rho = T(\rho; \tau, \kappa_T)$ *of a continuous self-map* $T$ *on a compact invariant subset of* $\mathbb{R}^2_+$*. On the canonical parameter grid, the fixed point* $\rho^*(\tau; \kappa_T)$ *is unique and* $C^1$ *in all arguments.*

Appendix A uses Brouwer's theorem for existence and the implicit function theorem for local uniqueness and regularity. I verify uniqueness and $C^1$ regularity numerically on the canonical grid; a global analytical argument is not available.

## 2.3 Managerial action: EL and DL in the disclosure environment

The manager observes $(y_A, y_B, \tilde{p})$ and takes an action $k \in \mathbb{R}$ to minimize the true expected squared loss $\mathbb{E}[(k-\theta)^2]$. The two specifications are:

**Definition 3** (Equilibrium-level cursedness (EL))**.** The EL $\kappa_M$-cursed manager computes her posterior using a mis-specified perceived noise precision $(1-\kappa_M)^2\tau_z$ in place of $\tau_z$ for the price signal, and sets her action equal to that cursed posterior mean:

$$k_{\mathrm{EL}}(\kappa_M) := \tilde{m}_B(\kappa_M), \tag{EL}$$

where $\tilde{m}_B(\kappa_M)$ is the linear Kalman estimate of $\theta$ given $(y_A, y_B, \tilde{p})$ under the perceived signal covariance with $\operatorname{Var}(\tilde{p})$ replaced by $V_0 + 1/[(1-\kappa_M)^2\tau_z]$ and $V_0 := (\rho_A^*)^2/\tau_{\theta A} + (\rho_B^*)^2/\tau_{\theta B}$.

The EL specification is a direct transcription of the Eyster et al. (2019) belief-stage primitive to the manager's side: the agent forms a single Kalman posterior under a perceived joint Gaussian in which the price-signal noise precision is attenuated.

**Definition 4** (Decision-level counterpart (DL))**.** Let

$$m_B := \mathbb{E}[\theta \mid y_A, y_B, \tilde{p}], \qquad m_{NP} := \mathbb{E}[\theta \mid y_A, y_B],$$

be correct $\mathbb{P}$-projections. The DL counterpart with mixing weight $\kappa_M$ is

$$k_{\mathrm{DL}}(\kappa_M) := (1 - \kappa_M)\, m_B + \kappa_M\, m_{NP}. \tag{DL}$$

The two specifications agree at the endpoints $\kappa_M \in \{0, 1\}$ and differ in the interior. DL is the tower-property polynomial counterpart of the EL primitive in the sense of Definition 2: both projections $m_B$ and $m_{NP}$ are correct under $\mathbb{P}$, so the tower property applies and the DL action induces a polynomial welfare expression. I use DL not as a competing behavioral model (it has no micro-foundation in the ERV tradition; no published paper computes welfare in a cursed-disclosure setting by mixing two correct Bayesian posteriors with ERV weight) but as the algebraic object that a modeler hoping for a polynomial decomposition would have to reach, and that Theorem 1 shows the EL primitive cannot match. Numerically, DL approximates EL with a parameter-dependent error (Section 4); a modeler willing to accept the approximation gains closed-form tractability at the cost of the documented deviation.

## 2.4 Planner

The planner minimizes the expected squared action error plus a quadratic disclosure cost,

$$W(\tau_A, \tau_B; \kappa_M, \kappa_T) = \mathbb{E}[(k - \theta)^2] + \frac{c_0}{2}(\tau_A^2 + \tau_B^2), \qquad c_0 > 0, \tag{W}$$

where $k$ is produced by (EL) or (DL) and the expectation is under $\mathbb{P}$. The parameter $c_0 > 0$ is a mathematical regularization that guarantees interior optima; the qualitative claims of the paper do not depend on it, and magnitudes are reported as sensitivity bands. I write $W_{\mathrm{EL}}$ and $W_{\mathrm{DL}}$ for welfare under the two specifications.

# 3 Main Theorem: Structural Rational-versus-Polynomial Characterization

This section states and proves the paper's main result (Theorem 1): under any single-moving-entry EL primitive in the sense of Definition 1, the true-measure welfare function is rational, not polynomial, in the bias parameter on a nonempty open set of primitive configurations. Theorem 2 records the polynomial-in-$\chi$ decomposition available at the DL counterpart via the tower property.

## 3.1 Statement

**Theorem 1** (Structural rational-versus-polynomial characterization)**.** *Consider any single-moving-entry EL primitive in the sense of Definition 1, embedded in a linear-Gaussian welfare problem of the form $W(\chi) = \mathbb{E}_{\mathbb{P}}[(b(\chi)^\top s - L\theta)^2]$ for a fixed linear functional $L$, with expectation taken under the true joint distribution and with the perceived primitives evaluated at some equilibrium fixed point $\rho^*$ that is $C^1$ in a surrounding primitive vector. Then:*

(i) *The perceived Kalman weights $b(\chi)$ are rational functions of $u = g(\chi)$ with common linear denominator $D(u)$, where $D$ is a polynomial of degree 1 in $u$ whose coefficients are $C^1$ in the primitive vector.*

(ii) *The true-measure welfare $W(\chi)$ is a rational function of $u$ with denominator $D(u)^2$ and numerator $N(u)$ of degree at most 2.*

(iii) *There exists a nonempty open set $\mathcal{V}$ of primitive configurations on which $D(u)^2$ does not divide $N(u)$. On $\mathcal{V}$, $W(\chi)$ is a strictly rational, non-polynomial function of $u$, and hence not a polynomial in $\chi$.*

Theorem 1 is the paper's central mathematical claim. Its content is a structural feature of Cramer's rule applied to a rank-one perturbation of a signal covariance: a single moving entry propagates through the cofactor expansion to give linear-over-linear weights, which in turn give a rational-with-$D(u)^2$ welfare. The non-divisibility condition of (iii) is a single polynomial equation in primitives, which generically does not hold, so the welfare function is generically rational-not-polynomial.

**Theorem 2** (DL polynomial counterpart)**.** *Any DL counterpart in the sense of Definition 2 yields a polynomial-in-$\chi$ welfare expression. Specifically, if $k_{\mathrm{DL}}(\chi) = \sum_\ell \lambda_\ell(\chi)\, m_\ell$ with each $\lambda_\ell$ a polynomial of degree $\leq d$ in $\chi$ and each $m_\ell$ a correct $\mathbb{P}$-projection of $L\theta$, then*

$\mathbb{E}_{\mathbb{P}}[(k_{\mathrm{DL}}(\chi) - L\theta)^2]$ *is a polynomial in* $\chi$ *of degree at most* $2d$*, and the coefficients are determined by the cross-moments* $\mathbb{E}_{\mathbb{P}}[(m_\ell - L\theta)(m_{\ell'} - L\theta)]$*, which by the tower property collapse to differences of the projection variances.*

Theorem 2 is the algebraic baseline. The contrast (i)-(ii)-(iii) of Theorem 1 versus Theorem 2 says: the polynomial decomposition that a DL counterpart admits automatically is precisely the object that no single-moving-entry EL primitive can reproduce. This is the paper's structural contribution.

## 3.2 Proof of Theorem 1

**Step 1: Cramer's-rule form of the Kalman weights.** Under Definition 1, only the $(i^*, j^*)$ entry of $\Sigma^{\mathrm{perc}}_{\mathrm{sig}}(\chi)$ depends on $\chi$, and rationally in $u = g(\chi)$ with linear denominator. Write that entry as $\sigma_{i^*j^*}(u) = (a_1 u + a_0)/(d_1 u + d_0)$ for constants $a_0, a_1, d_0, d_1$. Each cofactor of $\Sigma^{\mathrm{perc}}_{\mathrm{sig}}$ either omits row $i^*$ and column $j^*$ (and is then constant in $u$) or includes that entry exactly once by multilinearity of the determinant (and is then linear in $u$ after clearing the denominator $d_1 u + d_0$). The determinant $\det \Sigma^{\mathrm{perc}}_{\mathrm{sig}}(u)$ is linear in $\sigma_{i^*j^*}$ by cofactor expansion along row $i^*$, hence after clearing denominators is a linear polynomial in $u$ divided by $d_1 u + d_0$. Combining, the Kalman weights $b(\chi)^\top = \Sigma^{\mathrm{perc}}_{\theta,\mathrm{sig}} \operatorname{adj}(\Sigma^{\mathrm{perc}}_{\mathrm{sig}}) / \det \Sigma^{\mathrm{perc}}_{\mathrm{sig}}$ are each of the form

$$b_j(u) = \frac{\tilde{P}_j\, u + \tilde{Q}_j}{D(u)}, \qquad D(u) := (\text{linear in } u), \tag{6}$$

for constants $\tilde{P}_j, \tilde{Q}_j$ depending smoothly on primitives and on the equilibrium fixed point $\rho^*$ (which inherits $C^1$ regularity from the hypothesis). This proves (i).

**Step 2: MSE has denominator $D(u)^2$ with numerator of degree $\leq 2$.** Under $\mathbb{P}$, the action error is a linear combination of the Gaussian primitives with coefficients that are either of the form $b_j(u)$ or $b_j(u) - (\text{constant})$, each linear-over-$D(u)$. Taking the expectation of the square, each of the finitely many variance terms contributes a ratio $(\text{linear in } u)^2 / D(u)^2$. Summing gives $W(\chi) = N(u)/D(u)^2$ with $N$ a polynomial in $u$ of degree at most 2. This proves (ii).

**Step 3: open-set non-divisibility.** $D(u)^2$ divides $N(u)$ if and only if $D$ has a double root that is also a double root of $N$; since $\deg D = 1$, the only possibility is that $D$ has some root $u_*$ at which $N(u_*) = 0$ and that root is of multiplicity 2 in $N$ (requiring $N = cD^2$ for some $c$, i.e., three polynomial equations in primitives). Generically, $N(u_*) \neq 0$, and non-divisibility is a single strict polynomial inequality $N(u_*) \neq 0$. The coefficient map from primitives to

$(N, D)$ is $C^1$ by the $C^1$ regularity of $\rho^*$ and the smoothness of the Kalman algebra. Hence the set $\mathcal{V} := \{N(u_*) \neq 0\}$ is open. It is nonempty as long as the single-moving-entry structure is nondegenerate: at any primitive configuration at which the moving entry $\sigma_{i^*j^*}$ is not constant in $u$ (i.e., $a_1 d_0 - a_0 d_1 \neq 0$), a generic choice of remaining primitives gives $N(u_*) \neq 0$ because a nondegenerate rank-one perturbation generically does not produce a degenerate numerator. Application 1 in Section 4 verifies non-divisibility explicitly at a canonical configuration, demonstrating non-emptiness. This proves (iii). □

**Proof of Theorem 2.** Each $m_\ell - L\theta$ has $\mathbb{E}_{\mathbb{P}}[m_\ell - L\theta] = 0$, and by the tower property $\mathbb{E}_{\mathbb{P}}[(m_\ell - L\theta)(m_{\ell'} - L\theta)] = \mathbb{E}_{\mathbb{P}}[(m_{\ell_{\min}} - L\theta)^2]$, where $\mathcal{G}_{\ell_{\min}}$ is the coarser of the two sub-sigma-fields (correct projections onto nested sigma-fields have the nested tower property). The welfare is $\sum_{\ell,\ell'} \lambda_\ell(\chi)\lambda_{\ell'}(\chi)\, \mathbb{E}_{\mathbb{P}}[(m_\ell - L\theta)(m_{\ell'} - L\theta)]$, which is polynomial in $\chi$ of degree at most $2d$ because each $\lambda_\ell$ is polynomial of degree $\leq d$. □

## 3.3 Economic reading

The DL counterpart computes correct $\mathbb{P}$-projections and mixes them at the action stage. Because each projection is the orthogonal projection of $L\theta$ onto its information set, the tower-cross identity holds and the bias multiplies cleanly into a polynomial. The EL primitive, by contrast, miscomputes a single projection with a wrong perceived covariance. The action is no longer an orthogonal projection of $L\theta$ onto any information set, so the tower property fails. The mis-specified Kalman weights are rational in $\chi$ through the single moving entry of the perceived signal covariance, and the true-measure MSE inherits that rational structure. The structural rational-versus-polynomial dichotomy is a signature of the *level* at which the bias enters the projection machinery: belief-stage (EL) gives rational; action-stage (DL) gives polynomial. The characterization is tight: Theorem 1 uses only that exactly one entry of the perceived covariance moves; Theorem 2 uses only that each projection in the mixture is correct under $\mathbb{P}$.

**What falls outside the template.** Two kinds of EL primitives lie outside Definition 1. First, primitives that move the cross-covariance $\Sigma^{\text{perc}}_{\theta,\text{sig}}$ break condition (b); the Cramer's-rule argument still gives rational weights, but the numerator-denominator structure changes and welfare can have more complex rational dependence on $\chi$ (including polynomial as a special algebraic coincidence). Second, primitives that move two or more entries of $\Sigma^{\text{perc}}_{\text{sig}}$ with independent parameters break (a); the denominator becomes multivariate rational and the conclusions of Theorem 1 do not follow without further argument. The single-moving-entry condition is tight: it is the precise structure that forces the linear-denominator, $D(u)^2$-squared-denominator, rational-not-polynomial pattern.

**Depth and value of the observation.** The proof of Theorem 1 is a two-page Cramer's-rule-plus-tower-property observation about rank-one perturbations of a Kalman inverse covariance. It is short. Two features make it non-obvious and useful. First, the behavioral-finance literature routinely works with EL-style belief-stage mis-specifications but does not typically ask the welfare-decomposition question at interior $\chi$; a modeler confronting a new EL primitive would plausibly conjecture, by analogy with the $\chi = 0$ benchmark, that a polynomial-in-$\chi$ decomposition extends. Theorem 1 rules out that conjecture structurally, without requiring the modeler to derive the welfare function of the specific primitive. Second, the class of single-moving-entry EL primitives is broad enough to cover several distinct behavioral stories (cursedness, PET, overconfidence, base-rate neglect on a single prior precision), so the template converts what would otherwise be a series of case-by-case checks into a single portable sufficient condition. The paper's value beyond any one application is exactly this portability.

# 4 Applications

Three behavioral primitives in the linear-Gaussian-Kalman class fit Definition 1. Application 1 (cursed managerial disclosure, Section 4.1) is derived in full, including the explicit construction of $N(u)$ and $D(u)$, the non-divisibility check at the canonical scenario, and the numerical magnitude of the DL-EL gap. Applications 2 and 3 (PET; overconfidence, Section 4.2) are described at the conjecture level; I defer rigorous embedding in fully-worked equilibrium settings.

## 4.1 Application 1: Cursed Managerial Disclosure

In the disclosure environment of Section 2.2, the EL manager's perceived $3 \times 3$ signal covariance is

$$\Sigma_{\text{sig}}^{\text{perc}}(\kappa_M) = \begin{pmatrix} 1/\tau_{\theta A} + 1/\tau_A & 0 & \rho_A^*/\tau_{\theta A} \\ 0 & 1/\tau_{\theta B} + 1/\tau_B & \rho_B^*/\tau_{\theta B} \\ \rho_A^*/\tau_{\theta A} & \rho_B^*/\tau_{\theta B} & V_0 + \dfrac{1}{(1-\kappa_M)^2 \tau_z} \end{pmatrix}, \tag{7}$$

with $V_0 = (\rho_A^*)^2/\tau_{\theta A} + (\rho_B^*)^2/\tau_{\theta B}$. The perceived cross-covariance $\Sigma_{\theta,\text{sig}}^{\text{perc}} = (1/\tau_{\theta A},\, 1/\tau_{\theta B},\, \rho_A^*/\tau_{\theta A} + \rho_B^*/\tau_{\theta B})$ is $\kappa_M$-free. Only the $(3,3)$ entry of $\Sigma_{\text{sig}}^{\text{perc}}$ depends on $\kappa_M$, with $g(\kappa_M) = (1-\kappa_M)^2$ and moving indices $(i^*, j^*) = (3,3)$. Definition 1 is satisfied, and Theorem 1 applies.

**Corollary 1** (Cursed-disclosure obstruction)**.** *In the noisy-REE of Section 2.2, for every fixed* $(\tau, \kappa_T)$ *in a nonempty open set* $\mathcal{U}$ *of primitive parameter configurations containing the canonical scenario* $S_1$ $(\tau_{\theta A} = \tau_{\theta B} = \alpha = \tau_z = a = 1,\ \tau_A = \tau_B = 0.5,\ \kappa_T = 0)$*, the function* $\kappa_M \mapsto W_{\mathrm{EL}}(\tau; \kappa_M, \kappa_T) - (c_0/2)\|\tau\|^2$ *is a strictly rational, non-polynomial function of* $\kappa_M \in [0, 1)$ *with denominator* $D(u)^2$*,* $u = (1 - \kappa_M)^2$*. Consequently, on* $\mathcal{U}$ *there is no polynomial* $f$ *in* $\kappa_M$ *for which* $W_{\mathrm{EL}}$ *admits the representation* $W_{\mathrm{prior}}(\tau) - f(\kappa_M)W_{\mathrm{diff}}(\tau; \kappa_T) + (c_0/2)\|\tau\|^2$*. In contrast, the DL counterpart* (DL) *admits the polynomial decomposition*

$$W_{\mathrm{DL}}(\tau; \kappa_M, \kappa_T) = W_{\mathrm{prior}}(\tau) - (1 - \kappa_M^2)\, W_{\mathrm{diff}}(\tau; \kappa_T) + \frac{c_0}{2}(\tau_A^2 + \tau_B^2), \tag{8}$$

*with* $W_{\mathrm{prior}}(\tau) := 1/(\tau_{\theta A} + \tau_A) + 1/(\tau_{\theta B} + \tau_B)$ *and* $W_{\mathrm{diff}}(\tau; \kappa_T) := \mathbb{E}[(m_B^{\kappa_T} - m_{NP})^2]$*.*

Corollary 1 follows from Theorem 1 once the single-moving-entry structure of (7) is verified, and from Theorem 2 applied to the two-posterior DL mixture.

**Explicit non-divisibility check at $S_1$.** I verify the nonemptiness of $\mathcal{U}$ by exhibiting $S_1 \in \mathcal{U}$ explicitly. In the variable $u = (1 - \kappa_M)^2$, the $(3, 3)$ entry is $\sigma_{33}(u) = V_0 + 1/(u\tau_z) = (V_0 u\tau_z + 1)/(u\tau_z)$. The determinant of $\Sigma_{\mathrm{sig}}^{\mathrm{perc}}(u)$ is $R\,\sigma_{33}(u) + S = [(RV_0 + S)u\tau_z + R]/(u\tau_z)$, with $R = (1/\tau_{\theta A} + 1/\tau_A)(1/\tau_{\theta B} + 1/\tau_B) > 0$ and $S$ determined by primitives and $\rho^*$. Write $D(u) := (RV_0 + S)u\tau_z + R$. Cofactor expansion gives the Kalman weights $b_j(u) = (\tilde{P}_j u + \tilde{Q}_j)/D(u)$ as in (6). The action error

$$k_{\mathrm{EL}}(u) - \theta = (b_A + b_p\rho_A^* - 1)\theta_A + (b_B + b_p\rho_B^* - 1)\theta_B + b_A\varepsilon_A + b_B\varepsilon_B + b_p\tilde{z}$$

has independent terms under $\mathbb{P}$; taking the expectation of the square gives

$$\mathrm{MSE}_{\mathrm{EL}}(u) = \frac{(b_A + b_p\rho_A^* - 1)^2}{\tau_{\theta A}} + \frac{(b_B + b_p\rho_B^* - 1)^2}{\tau_{\theta B}} + \frac{b_A^2}{\tau_A} + \frac{b_B^2}{\tau_B} + \frac{b_p^2}{\tau_z}, \tag{9}$$

a rational function of $u$ with denominator $D(u)^2$ and numerator $N(u)$ of degree at most 2. At $S_1$ the committed script `prop2_2component_verify.py` evaluates $N$ at the root $u_* = -R/[(RV_0 + S)\tau_z] < 0$ of $D$ and returns $N(u_*) \neq 0$: non-divisibility holds at $S_1$. Continuity of the coefficient map from primitives and the $C^1$ fixed point $\rho^*$ (Lemma 1) into $(N, D)$ extends the nonzero margin to an open neighborhood $\mathcal{U}$ of $S_1$ in the primitive vector, and to an open extension in $\kappa_T$ via smoothness of $\rho^*$ in $\kappa_T$. The committed script verifies non-divisibility at two additional scenarios $S_2, S_3$, each inside $\mathcal{U}$.

A subtlety in the logical bridge between "rational in $u$" and "not polynomial in $\kappa_M$" deserves attention. If $W$ were polynomial in $\kappa_M$ of degree $d$, it would be polynomial in $u = (1 - \kappa_M)^2$ only if every odd-degree term in $\kappa_M$ vanished; under the EL specification,

Table 1: Level-specificity gap in the two-component REE. Source: `prop2_2component_verify.txt`. $S_1$: $\tau_{\theta A} = \tau_{\theta B} = \alpha = \tau_z = a = 1$, $\tau_A = \tau_B = 0.5$, $\kappa_T = 0$. $S_3$: $\tau_{\theta A} = \tau_{\theta B} = 1, \alpha = 0.3, \tau_z = 2, a = 0.5$, $\tau_A = \tau_B = 0.5$, $\kappa_T = 0$. DL prediction is $W_{\text{prior}} - (1 - \kappa_M^2) W_{\text{diff}}$. Gap is $\text{MSE}_{\text{EL}}$ − DL prediction.

| Scenario | $\kappa_M$ | $\text{MSE}_{\text{EL}}$ | DL pred | Gap | Gap % | $f_{\text{EL}}$ | $1 - \kappa_M^2$ |
|---|---|---|---|---|---|---|---|
| $S_1$ | 0.00 | 0.9200 | 0.9200 | 0.0000 | 0.00% | 1.000 | 1.000 |
| $S_1$ | 0.25 | 0.9510 | 0.9458 | 0.0052 | 0.55% | 0.925 | 0.938 |
| $S_1$ | 0.50 | 0.9998 | 1.0233 | −0.0236 | −2.30% | 0.807 | 0.750 |
| $S_1$ | 0.75 | 1.1721 | 1.1525 | 0.0197 | 1.71% | 0.390 | 0.438 |
| $S_1$ | 1.00 | 1.3333 | 1.3333 | 0.0000 | 0.00% | 0.000 | 0.000 |
| $S_3$ | 0.00 | 0.6803 | 0.6803 | 0.0000 | 0.00% | 1.000 | 1.000 |
| $S_3$ | 0.25 | 0.7330 | 0.7211 | 0.0119 | 1.65% | 0.919 | 0.938 |
| $S_3$ | 0.50 | 0.9192 | 0.8435 | 0.0756 | 8.97% | 0.634 | 0.750 |
| $S_3$ | 0.75 | 1.1911 | 1.0476 | 0.1435 | 13.70% | 0.218 | 0.438 |
| $S_3$ | 1.00 | 1.3333 | 1.3333 | 0.0000 | 0.00% | 0.000 | 0.000 |

$W$ enters $\kappa_M$ only through $u$ (since $\sigma_{33}$ is a function of $(1 - \kappa_M)^2$), so polynomial-in-$\kappa_M$ is equivalent to polynomial-in-$u$. The rational-not-polynomial-in-$u$ conclusion therefore implies rational-not-polynomial-in-$\kappa_M$.

**Endpoint consistency.** At $\kappa_M = 0$ ($u = 1$), $\sigma_{33}(1) = V_0 + 1/\tau_z$ equals the true variance of $\tilde{p}$; cursed-Kalman weights coincide with true Bayesian weights; $k_{\text{EL}} = m_B$. At $\kappa_M = 1$ ($u = 0$), $\sigma_{33} \to \infty$; $b_p \to 0$ and $(b_A, b_B)$ converge to Bayesian weights on $(y_A, y_B)$; $k_{\text{EL}} = m_{NP}$. Both endpoints match the DL action, consistent with agreement at $\kappa_M \in \{0, 1\}$ in (8). Rational-versus-polynomial is compatible with agreement at isolated points.

**Proof of** (8) **(DL side).** By the tower property, $m_{NP} = \mathbb{E}[m_B^{\kappa_T} \mid \sigma(y_A, y_B)]$ under $\mathbb{P}$, so $\mathbb{E}[(m_B^{\kappa_T} - m_{NP}) m_{NP}] = 0$ and $\mathbb{E}[(m_B^{\kappa_T} - m_{NP})(m_{NP} - \theta)] = -W_{\text{diff}}$. Squaring $k_{\text{DL}} - \theta = (1 - \kappa_M)(m_B^{\kappa_T} - \theta) + \kappa_M (m_{NP} - \theta)$, using these identities, the coefficient on $W_{\text{prior}}$ collapses to 1 and the coefficient on $W_{\text{diff}}$ to $-(1 - \kappa_M^2)$. Adding the cost term gives (8). The identity is indifferent to $\kappa_T$ because tower uses only the true measure, not the trader-side specification. □

**Numerical magnitudes: the DL-EL gap.** Table 1 reports, for each $\kappa_M \in \{0, 0.25, 0.5, 0.75, 1.0\}$, the EL manager's actual MSE, the DL baseline $W_{\text{prior}} - (1 - \kappa_M^2) W_{\text{diff}}$, and the implied weight $f_{\text{EL}} := (W_{\text{prior}} - \text{MSE}_{\text{EL}})/W_{\text{diff}}$, at scenarios $S_1$ (symmetric) and $S_3$ (low $\alpha$, high $\tau_z$).

In $S_1$ the interior gap is small (±2%) and oscillates in sign; the implied $f_{\text{EL}}$ crosses $1 - \kappa_M^2$ at two interior points. In $S_3$ the gap is uniformly positive and monotone in magnitude, reaching $+13.7\%$ at $\kappa_M = 0.75$. The qualitative contrast (oscillation versus monotonicity) is driven by the ratio $V_0\tau_z$ in (7): a large ratio makes the $1/[(1-\kappa_M)^2\tau_z]$ term dominate at interior $\kappa_M$ and produces a one-signed deviation; a moderate ratio gives oscillation. The DL counterpart is therefore a bounded-error approximation to EL, with error magnitude ranging from $\pm 2\%$ to $+13.7\%$ across the tested scenarios. A modeler willing to accept this error gains the closed-form decomposition (8).

**Pedagogical one-component illustration.** A one-component reduction ($\tau_\theta = \tau_y = \tau_z = \rho = 1$) gives the EL implied weight $f_{\mathrm{EL}}(0.5) = 23/27 \approx 0.852$, while the DL scalar is $1 - 0.25 = 0.75$. The closed-form gap 0.102 mirrors the rational-versus-polynomial contrast of Theorem 1 in a setting in which every step is hand-checkable. This reduction is pedagogical; the load-bearing open-set claim is the two-component REE argument above.

## 4.2 Applications 2 and 3: PET and Overconfidence (conjecture-level)

Two further primitives plausibly fit Definition 1. I describe them at the conjecture level: I defer a rigorous derivation that each primitive reduces to a single moving entry of $\Sigma_{\mathrm{sig}}^{\mathrm{perc}}$ in a fully-worked equilibrium setup. Specifically, I do not cite a published equation in which either primitive takes the single-moving-entry form; the matching below is a plausible reduced form, not a worked theorem.

**Application 2: Partial-equilibrium thinking (PET).** In the linear-normal framework of Bastianello and Fontanier (2025), a PET-cursed agent with parameter $\chi \in [0, 1]$ partially ignores how aggregate signals feed into the equilibrium price. In a price-reading specification, I conjecture that this attenuates the perceived informativeness of the aggregate price signal by a factor $(1 - \chi)$, with the attenuation entering the price-signal precision entry of the perceived inverse covariance as a factor $(1 - \chi)^2$. This would make PET single-moving-entry with moving indices $(i^*, j^*) = (n, n)$ (the price-precision entry) and $g(\chi) = (1 - \chi)^2$, fitting the same algebraic template as cursedness (Application 1). Under this conjecture, Theorem 1 applies and PET-welfare is rational-not-polynomial in $\chi$.

The conjecture requires a rigorous derivation I do not provide: in particular, the PET primitive in Bastianello and Fontanier (2025) attenuates the loading the trader perceives on the aggregate signal, and whether this attenuation cleanly reduces to a single moving entry of $\Sigma_{\mathrm{sig}}^{\mathrm{perc}}$ (as opposed to moving both a diagonal entry and an off-diagonal entry simultaneously) depends on details of the PET specification that I have not worked through in this paper.

Application 2 is therefore conjecture-level.

**Application 3: Managerial overconfidence about own-signal precision.** A manager who overestimates the precision of her private signal $y_A$ by a factor $1+\psi$ (with $\psi > 0$) inflates one perceived precision while leaving the prior and cross-covariances alone. Translating to the perceived signal covariance, this multiplies the own-signal precision entry by $(1+\psi)$, which in the inverse acts on a single diagonal entry of $\Sigma_{\text{sig}}^{\text{perc}}$: $(\Sigma_{\text{sig}}^{\text{perc}})_{11} = 1/\tau_{\theta A} + 1/[(1+\psi)\tau_A]$ (or $= 1/\tau_{\theta A} + 1/[(1-\phi)^2\tau_A]$ under a reparameterization $\psi = 1/(1-\phi)^2 - 1$ with $\phi \in [0,1)$). This satisfies Definition 1 with moving indices $(1,1)$ and $g(\phi) = (1-\phi)^2$. Under this reduced-form primitive, Theorem 1 applies.

Application 3 is also conjecture-level in the sense that I have not tied the reduced-form primitive above to a specific published overconfidence model (for example, Daniel et al., 1998; Odean, 1998); the reduced form is an illustrative instantiation of the template, not a derived consequence of a specific behavioral primitive. It is a placeholder showing that the class contains primitives other than cursedness.

**Status of Applications 2 and 3.** I flag Applications 2 and 3 as conjecture-level and do not rely on them for any downstream claim in the paper. Their role is to illustrate the template's scope: the class of single-moving-entry EL primitives plausibly contains several distinct behavioral stories, so the value of Theorem 1 is not purely its bearing on cursedness.

# 5 Welfare Loss and Comparative Statics at the DL Counterpart

This section records two downstream consequences of the polynomial DL decomposition (8) for Application 1 (cursed managerial disclosure). Both results live on the DL side; they characterize the polynomial-counterpart welfare rather than the EL welfare Theorem 1 shows is rational. Because DL is the tower-property polynomial counterpart rather than a behavioral model, these results should be read as properties of the closed-form approximation, not as predictions about EL-cursed managers per se.

## 5.1 Closed-form welfare loss at the DL counterpart

Fix $\kappa_T \in [0,1]$ and let $\tau^0 := \tau^*(0;\kappa_T)$ denote the planner's DL-optimum at the rational Bayesian benchmark. Let $W^*(\kappa_M;\kappa_T) := \min_\tau W_{\text{DL}}(\tau;\kappa_M,\kappa_T)$ and $L(\kappa_M;\kappa_T,c_0) := W_{\text{DL}}(\tau^0;\kappa_M,\kappa_T) - W^*(\kappa_M;\kappa_T)$.

**Proposition 1** (Closed-form welfare loss, conditional on A2)**.** *Conditional on Assumption 1 (crowding-out sign on a compact box, verified numerically on the canonical grid) and on the interior-benchmark regularity of Lemma 3,*

$$L(\kappa_M; \kappa_T, c_0) = \kappa_M^2 \, W_{\text{diff}}(\tau^0; \kappa_T) - \int_0^{\kappa_M} 2s \, W_{\text{diff}}(\tau^*(s; \kappa_T); \kappa_T) \, ds. \tag{10}$$

*$L$ is nonnegative, vanishes at $\kappa_M = 0$, and is weakly increasing in $\kappa_M$.*

*Proof.* From (8) at $\tau = \tau^0$, $W_{\text{DL}}(\tau^0; \kappa_M, \kappa_T) = W_{\text{DL}}(\tau^0; 0, \kappa_T) + \kappa_M^2 W_{\text{diff}}(\tau^0; \kappa_T)$. Since $\tau^0$ minimizes $W_{\text{DL}}(\cdot; 0, \kappa_T)$, $W^*(0; \kappa_T) = W_{\text{DL}}(\tau^0; 0, \kappa_T)$. The envelope theorem gives $dW^*/d\kappa_M = 2\kappa_M W_{\text{diff}}(\tau^*(\kappa_M; \kappa_T); \kappa_T)$. Integrating from 0 to $\kappa_M$ and subtracting delivers (10). Nonnegativity follows from $L = W_{\text{DL}}(\tau^0; \kappa_M, \kappa_T) - \min_\tau W_{\text{DL}}(\tau; \kappa_M, \kappa_T) \geq 0$; monotonicity in $\kappa_M$ uses $L'(\kappa_M) = 2\kappa_M[W_{\text{diff}}(\tau^0; \kappa_T) - W_{\text{diff}}(\tau^*(\kappa_M); \kappa_T)] \geq 0$, which rests on Assumption 1 (only numerically verified). □

The magnitudes of $L/W^*$ span approximately two orders of magnitude as $c_0$ varies from 0.01 to 0.2: at $\kappa_M = 0.5$, for the symmetric scenario $S_1$ the ratio ranges from 0.76% to 3.42%; for the asymmetric-prior scenario the ratio ranges from 15.52% to 42.08% (Appendix C, `welfare_loss_table.csv`). The spread reflects $c_0$'s role as a mathematical regularizer rather than an empirical cost primitive. The qualitative shape of $L$ (nonnegativity, vanishing at $\kappa_M = 0$, and monotonicity) holds across all reported cells; the level does not, and a calibrated magnitude requires an empirical anchor for $c_0$ that this paper does not supply.

## 5.2 Monotone amplification (conditional)

The polynomial decomposition (8) has a natural comparative-statics reading: cursedness should amplify the planner's optimal disclosure precisions. The argument needs three preconditions, each verified numerically on the canonical grid:

**Assumption 1** (Crowding-out sign on a compact box)**.** *There exists $\bar{\tau} > 0$ such that, on the box $[0, \bar{\tau}]^2$ and for each $X \in \{A, B\}$, $\partial W_{\text{diff}}(\tau; \kappa_T)/\partial \tau_X \leq 0$.*

**Assumption 2** (Intra-coordinate submodularity)**.** *$W_{\text{DL}}(\cdot; \kappa_M, \kappa_T)$ is submodular in $(\tau_A, \tau_B)$ on $[0, \bar{\tau}]^2$.*

**Assumption 3** ($\kappa_T$-monotonicity)**.** *$\partial W_{\text{diff}}/\partial \kappa_T \leq 0$ on $[0, \bar{\tau}]^2 \times [0, 1]$.*

I verify Assumption 1 at all 2205 grid points (file `a2_sweep.csv`); I verify Assumptions 2 and 3 on the canonical grid. Analytical proofs would require a closed-form solution for $\rho^*$ in the two-component setting, which does not admit a clean algebraic expression; the grid verification is a report that the partial-equilibrium sign dominates the feedback correction across the tested scenarios.

**Remark 1** (Monotone amplification, conditional on A2–3)**.** Conditional on Assumptions 1–3 and interiority of $\tau^*(\kappa_M, \kappa_T) \in (0, \bar{\tau})^2$, the interior argmin $\tau^*(\kappa_M, \kappa_T)$ is weakly increasing componentwise in $\kappa_M$ and in $\kappa_T$. The argument differentiates (8): $\partial^2 W_{\mathrm{DL}}/(\partial \tau_X \partial \kappa_M) = 2\kappa_M \, \partial W_{\mathrm{diff}}/\partial \tau_X \leq 0$ under Assumption 1, so $W_{\mathrm{DL}}$ is submodular in $(\tau_X, \kappa_M)$; combined with Assumption 2, Theorem 2.8.1 of Topkis (1998) delivers joint-argmin monotonicity of $\tau^*$ in $\kappa_M$. The $\kappa_T$ direction uses Assumption 3.

Because the three preconditions are numerical rather than analytical, Remark 1 is not a theorem. I state it as the comparative-static reading available to a modeler who accepts the numerical assumptions on a given parameter region. Proposition 1's monotonicity claim inherits the same conditionality through Assumption 1.

**Diagnostic on disclosure-rank reversal.** A complementary question for the DL counterpart is whether cursedness can reorder the Bayesian disclosure ranking $\tau_A^* > \tau_B^*$ when $\tau_{\theta A} < \tau_{\theta B}$. Appendix B records the necessary-and-sufficient reversal inequality at the interior benchmark and reports a corner-version numerical sweep (864 configurations, zero violations at the corner $\tau = (0, 0)$, minimum slack $7 \times 10^{-4}$). The corner test is strictly weaker than the interior test, so the sweep is evidence of non-reversal in a region, not a proof. I treat disclosure-rank reversal as a DL-internal diagnostic supporting the numerical-amplification direction, not as a main-text result.

# 6 Discussion

## 6.1 What the paper delivers, and what it does not

The paper delivers three things. First, a structural rational-versus-polynomial characterization (Theorem 1) covering any single-moving-entry EL primitive in the linear-Gaussian Kalman class. Second, a fully-derived application to cursed managerial disclosure in the Goldstein and Yang (2019) environment with an Eyster et al. (2019)-consistent manager (Corollary 1), including explicit non-divisibility at the canonical scenario and numerical magnitudes of the DL-EL gap. Third, conditional downstream results at the DL polynomial counterpart (Proposition 1; Remark 1), flagged as contingent on numerical assumptions.

The paper does not deliver: (i) a worked derivation of the PET or overconfidence applications from published primitives (Applications 2 and 3 are conjecture-level); (ii) an analytical proof of Assumptions 1–3 (the crowding-out direction holds in partial equilibrium but the REE-feedback contribution resists closed-form sign); (iii) an interior-benchmark test of the reversal inequality (the sweep is at a corner, which is strictly weaker); (iv) an empirically anchored magnitude for the welfare loss (the $c_0$ sensitivity band is two orders of magnitude, which reflects the regularizer's status as a mathematical device rather than an economic primitive).

## 6.2 Limitations

The DL counterpart has no behavioral micro-foundation; it is the tower-property polynomial counterpart of an EL primitive, mathematically well-defined for the purpose of comparison but not a competing behavioral model. Readers should interpret the DL-EL contrast as characterizing the algebraic cost of insisting on an ERV-grounded belief-stage primitive, not as a comparison between two behavioral theories. The closed-form welfare loss (Proposition 1) and monotone amplification (Remark 1) are properties of the DL polynomial approximation, conditional on Assumption 1 and its numerical siblings. Translating these properties back to EL-cursed managers goes through the bounded-error approximation of Table 1: the direction survives to EL as long as the error bound is smaller than the effect, which is the case on the tested scenarios but is not proved in general.

Three further limitations. First, the single-moving-entry template is narrow by design: primitives moving the cross-covariance or two independent entries of the perceived signal covariance lie outside. Second, the open-set $\mathcal{U}$ in Corollary 1 is established by continuity around $S_1$; I do not characterize $\partial\mathcal{U}$, and I have not proved the natural conjecture that $\mathcal{U}^c$ is cut out by algebraic constraints on polynomial division (hence Zariski-closed with empty interior). Third, the functional-form scope is CARA-Gaussian with squared-error loss. The DL tower-property identity uses this linear-quadratic-Gaussian structure; the EL rational-versus-polynomial structure uses only the Kalman-inverse-covariance algebra and survives more general linear-quadratic-Gaussian settings.

Uniqueness and $C^1$ regularity of the REE fixed point (Lemma 1) and the interior-benchmark regularity (Lemma 3) are established numerically on the canonical grid. I defer global analytical arguments.

## 6.3 Relation to existing results

Eyster et al. (2019) fix the belief-stage cursedness primitive that Application 1 specializes; their Sections 2–3 give the cursed-trader formulation that the EL manager transcribes. Goldstein and Yang (2019) give the $\kappa_M = 0$ rational Bayesian disclosure framework; Corollary 1 rules out the polynomial-in-$\kappa_M$ extension that a modeler might hope for. Bastianello and Fontanier (2025) give an EL primitive for traders that Application 2 conjectures fits the template. de Clippel and Zhang (2022) analyze non-Bayesian receivers in a persuasion setting whose primitive lies outside the Kalman-filter class; the theorem does not speak to that setting. The theorem also does not speak to primitives that move the cross-covariance $\Sigma^{\text{perc}}_{\theta,\text{sig}}$, such as an agent mis-estimating how signals map to latents.

# 7 Conclusion

When a scalar bias enters the perceived signal covariance of a linear-Gaussian Kalman problem through a single entry, rationally with a linear denominator, the true-measure welfare function is rational with squared linear denominator and generically not polynomial in the bias. The corresponding decision-level action rule (any polynomial-in-bias mixture of correct $\mathbb{P}$-projections) yields a polynomial welfare expression via the tower property. This structural rational-versus-polynomial characterization (Theorem 1) applies to cursed managerial disclosure in the Goldstein-Yang-Eyster-Rabin-Vayanos tradition (Application 1), and plausibly to partial-equilibrium thinking and managerial overconfidence (Applications 2 and 3, conjecture-level). Modelers seeking polynomial welfare decompositions in this class of behavioral-Kalman primitives cannot have them; the polynomial form is the tower-property counterpart available only at the decision stage.

# References


Snehal Banerjee, Jesse Davis, and Naveen Gondhi. Incentivizing effort and informing investment: The dual role of stock prices. *Review of Financial Studies*, 2025. doi: 10.1093/rfs/hhag008. Forthcoming.

Snehal Banerjee, Jesse Davis, and Naveen Gondhi. Information provision and the curse of knowledge. *Management Science*, 2026. doi: 10.1287/mnsc.2023.01839. Forthcoming.

Francesca Bastianello and Paul Fontanier. Expectations and learning from prices. *Review of Economic Studies*, 92(3):1341–1374, 2025. doi: 10.1093/restud/rdae059.

Kent Daniel, David Hirshleifer, and Avanidhar Subrahmanyam. Investor psychology and security market under- and overreactions. *Journal of Finance*, 53(6):1839–1885, 1998. doi: 10.1111/0022-1082.00077.

Geoffroy de Clippel and Xu Zhang. Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642, 2022. doi: 10.1086/720464.

Erik Eyster, Matthew Rabin, and Dimitri Vayanos. Financial markets where traders neglect the informational content of prices. *Journal of Finance*, 74(1):371–399, 2019. doi: 10.1111/jofi.12729.

Itay Goldstein and Liyan Yang. Good disclosure, bad disclosure. *Journal of Financial Economics*, 131(1):118–138, 2019. doi: 10.1016/j.jfineco.2018.08.004.

Terrance Odean. Volume, volatility, price, and profit when all traders are above average. *Journal of Finance*, 53(6):1887–1934, 1998. doi: 10.1111/0022-1082.00078.

Franz Ostrizek and Elia Sartori. Acquisition, (mis)use and dissemination of information: The blessing of cursedness and transparency. *SSRN Working Paper*, 2024. SSRN 3971571.

Donald M. Topkis. *Supermodularity and Complementarity.* Princeton University Press, 1998.

# A Proof of Lemma 1

Fix $(\tau_A, \tau_B, \tau_{\theta A}, \tau_{\theta B}, \alpha, \tau_z, a)$ and $\kappa_T \in [0,1]$. Standard CARA-Gaussian REE derivations for a two-component fundamental give a self-map $T : \mathbb{R}^2_+ \to \mathbb{R}^2_+$ such that a linear noisy-REE with price loading $\rho$ satisfies $\rho = T(\rho; \tau, \kappa_T)$. The self-map is

$$T(\rho; \tau, \kappa_T) = \frac{1}{a}\, \Lambda(\rho; \kappa_T)^{-1} \boldsymbol{\mu}(\tau),$$

where $\Lambda$ is the cursed-trader information matrix in (4) and $\boldsymbol{\mu}(\tau) \in \mathbb{R}^2_+$ collects the prior-and-public-signal precisions the informed traders' demand function assigns to the two fundamentals. $\Lambda$ is positive definite because the additive term $(1-\kappa_T)^2 \tau_z \rho\rho^\top$ is positive semi-definite and the diagonal dominates.

**Explicit bound for Brouwer.** We show $T$ maps the closed ball $\mathcal{B}_R := \{\rho \in \mathbb{R}^2_+ : \|\rho\|_2 \leq R\}$ into itself for $R$ sufficiently large. For any $\rho \in \mathbb{R}^2_+$, the diagonal of $\Lambda(\rho; \kappa_T)$ is bounded below by $\underline{\lambda} := \min(\tau_{\theta A} + \tau_A, \tau_{\theta B} + \tau_B) > 0$, so the smallest eigenvalue of $\Lambda$ satisfies $\lambda_{\min}(\Lambda(\rho; \kappa_T)) \geq \underline{\lambda}$ uniformly in $\rho$ (since the added rank-one terms $\alpha \mathbf{1}\mathbf{1}^\top + (1-\kappa_T)^2 \tau_z \rho\rho^\top$ are positive semidefinite). Consequently $\|\Lambda^{-1}\|_2 \leq 1/\underline{\lambda}$, and

$$\|T(\rho; \tau, \kappa_T)\|_2 \leq \frac{1}{a}\, \|\Lambda^{-1}\|_2\, \|\boldsymbol{\mu}(\tau)\|_2 \leq \frac{\|\boldsymbol{\mu}(\tau)\|_2}{a\, \underline{\lambda}} =: R^*.$$

Pick any $R \geq R^*$. Then $T(\mathcal{B}_R) \subset \mathcal{B}_R$. The set $\mathcal{B}_R \cap \mathbb{R}^2_+$ is compact, convex, and nonempty; $T$ is continuous on it (the map $\rho \mapsto \Lambda(\rho; \kappa_T)^{-1}$ is smooth wherever $\Lambda$ is invertible, which holds throughout by positive-definiteness). Brouwer's fixed-point theorem delivers $\rho^* \in \mathcal{B}_R \cap \mathbb{R}^2_+$ with $T(\rho^*; \tau, \kappa_T) = \rho^*$.

**Local $C^1$ regularity via IFT.** At any fixed point $\rho^*$, consider the residual $F(\rho; \tau, \kappa_T) := \rho - T(\rho; \tau, \kappa_T)$. The Jacobian $D_\rho F(\rho^*) = I - D_\rho T(\rho^*)$. Because $T$ is smooth in $\rho$ on an open neighborhood of $\rho^*$ (the entries of $\Lambda^{-1}$ are rational functions of $\rho$ with nonvanishing denominator), $F$ is $C^1$ jointly in $(\rho; \tau, \kappa_T)$. At the canonical symmetric scenario $S_1$, direct computation yields $\det(I - D_\rho T(\rho^*)) \neq 0$ (verified by evaluation in `prop2_2component_verify.py`). By the implicit function theorem, there exists an open neighborhood $\mathcal{V} \subset \mathbb{R}^6_+ \times [0, 1]$ of $(\tau^{S_1}, \kappa_T^{S_1})$ on which a unique $C^1$ function $(\tau, \kappa_T) \mapsto \rho^*(\tau; \kappa_T)$ satisfies $F(\rho^*(\tau; \kappa_T); \tau, \kappa_T) = 0$. This local $C^1$ regularity is sufficient for the open-set extension of Corollary 1 in the proof of Theorem 1: the map from primitives to $(N, D)$ is $C^1$ (hence continuous) on $\mathcal{V}$, and the nonzero margin $|N(u_*)|$ at $S_1$ extends to an open neighborhood.

Numerical verification on the canonical grid confirms $\det(D_\rho F) \neq 0$ at every evaluated point, so the local IFT argument covers every scenario we use. □

# B Auxiliary Lemmas

**Lemma 2** (Crowding-out sign). *On every parameter configuration in the committed file `signcheck_grid.csv` at which the REE solver converges, $\partial A_A / \partial \tau_A < 0$, where $A_A$ denotes the induced informational loading of $\theta_A$ in the manager's Bayesian posterior. The crowding-out sign is strict.*

Lemma 2 formalizes the claim that the two-component structure is load-bearing: if the loading matrix did not exhibit crowding-out between direct disclosure and price-based learning, the planner's problem would reduce to two uncoupled one-dimensional problems and the decomposition would be degenerate. I verify the crowding-out sign numerically on the grid. An analytical proof from primitives is available under additional regularity assumptions but I do not pursue it here.

**Lemma 3** (Interior Bayesian benchmark). *At $\kappa_M = 0$ with $\kappa_T$ fixed, if $c_X(0, \tau_B; \kappa_T) < 1/\tau_{\theta X}^2$ on the canonical grid for each $X \in \{A, B\}$, the joint minimization of $W_{\mathrm{DL}}(\cdot; 0, \kappa_T)$ over $\mathbb{R}^2_+$ admits a strictly interior critical point $\tau^0 \in \mathbb{R}^2_{++}$, unique on the canonical grid.*

*Proof sketch.* The FOC of (8) at $\kappa_M = 0$ is $1/(\tau_{\theta X} + \tau_X)^2 = c_X(\tau; 0) + c_0 \tau_X$. Under the hypothesis, the right side at $\tau_X = 0$ is below the left side, and the right side tends to infinity as $\tau_X \to \infty$ while the left side stays bounded. The intermediate value theorem and monotonicity under Assumption 1 deliver a strictly interior solution; uniqueness follows from strict monotonicity of both sides on the canonical grid. □

**Lemma 4** (Envelope regularity)**.** *Under strict second-order conditions at the interior optimum of* $W_{\text{DL}}(\cdot;\kappa_M,\kappa_T)$, $\tau^*(\kappa_M;\kappa_T)$ *is* $C^1$ *in* $\kappa_M \in [0,1]$.

*Proof.* Apply the implicit function theorem to the FOC with nonsingular Hessian. The Hessian is verified numerically to be strictly positive definite on the canonical grid. □

## Reversal inequality at the DL counterpart

The following characterization supports the disclosure-rank-reversal diagnostic flagged in Section 5.2. It is a DL-internal result: the object characterized is the polynomial counterpart, not the EL welfare.

**Proposition 2** (Reversal inequality at DL)**.** *Fix* $\kappa_T \in [0,1]$. *Let* $\tau^*(\kappa_M)$ *denote the interior joint argmin of* $W_{\text{DL}}(\cdot;\kappa_M,\kappa_T)$ *under Assumption 1, and* $c_X^*(\kappa_M) := -\partial W_{\text{diff}}/\partial\tau_X|_{\tau=\tau^*(\kappa_M)}$. *Let* $\Delta(\kappa_M) := \tau_A^*(\kappa_M) - \tau_B^*(\kappa_M)$.

(i) $\Delta(1) > 0$ *if and only if* $\tau_{\theta A} < \tau_{\theta B}$.

(ii) *Interior reversal* $\Delta(0) < 0$ *holds if and only if*

$$c_A^*(0) - c_B^*(0) > \frac{1}{(\tau_{\theta A} + \tau_A^*(0))^2} - \frac{1}{(\tau_{\theta B} + \tau_B^*(0))^2}. \tag{Rev}$$

(iii) *If* (Rev) *holds and* $\tau_{\theta A} < \tau_{\theta B}$, *reversal occurs at some* $\kappa_M \in (0,1)$.

*Proof.* The FOC of (8) in $\tau_X$ is $1/(\tau_{\theta X} + \tau_X^*)^2 = (1-\kappa_M^2)c_X^*(\kappa_M) + c_0\tau_X^*$. At $\kappa_M = 1$ the first right-side term vanishes, leaving $1/(\tau_{\theta X} + \tau_X^*)^2 = c_0\tau_X^*$, strictly decreasing in $\tau_{\theta X}$; this gives (i). Differencing the FOCs across $X$ at $\kappa_M = 0$ yields $c_0\Delta(0) = [1/(\tau_{\theta A} + \tau_A^*)^2 - 1/(\tau_{\theta B} + \tau_B^*)^2] - [c_A^*(0) - c_B^*(0)]$, so $\Delta(0) < 0$ iff (Rev). Continuity of $\tau^*(\kappa_M)$ on $[0,1]$ (Lemma 4) and the intermediate value theorem applied to $\Delta(\kappa_M)$ deliver (iii). □

The committed script `reversal_fast.py` sweeps 864 configurations of $(\tau_{\theta A}, \tau_{\theta B}, \alpha, \tau_z, a, c_0)$ with $\tau_{\theta A} < \tau_{\theta B}$ and evaluates (Rev) at the corner $(\tau_A, \tau_B) = (0,0)$ rather than at the interior benchmark $\tau^*(0)$. Across all 864 configurations the corner inequality fails with minimum slack $7 \times 10^{-4}$. The corner test is strictly weaker than the interior-benchmark test at $\tau^*(0)$: corner failure is consistent with interior non-reversal but does not prove it, so the sweep is evidence of non-reversal in a region, not a proof. I defer the interior-benchmark sweep due to optimizer cost at extreme parameter asymmetry.

# C Numerical Scripts and Supporting Tables

## C.1 Scripts

The committed computational artifacts are in `code/explore/`. Each file runs without additional arguments and writes its output to `output/stage2/` or `output/stage3a/`.

- `stage3a_core.py`: REE solver and welfare-component routines; supports $\kappa_T$. Invoked by every other script.
- `prop2_2component_verify.py`: two-component REE verification of Corollary 1 (Application 1 of Theorem 1). Solves the REE fixed point, computes the cursed-Kalman weights in the $3 \times 3$ perceived signal covariance, and reports $\mathrm{MSE}_{\mathrm{EL}}(\kappa_M)$, the DL baseline prediction, and $f_{\mathrm{EL}}(\kappa_M; \tau, \kappa_T)$ at scenarios $S_1, S_2, S_3$. Output: `prop2_2component_verify.txt`.
- `welfare_loss_table.py`: computes $L/W^*$ across scenarios, $\kappa_M$, and $c_0$. Output: `welfare_loss_table.csv` (45 rows).
- `a2_sweep.py`: verifies Assumption 1 on a 2205-point grid over $[0, 5]^2$ at $\kappa_T = 0$. Output: `a2_sweep.csv`.
- `two_sided_cursedness.py`: extends Assumption 1 to $\kappa_T \in \{0.3, 0.5\}$ on a coarser grid. Output: `two_sided_cursedness.csv`.
- `cost_robustness.py`: checks monotone argmin under linear and log disclosure costs. Output: `cost_robustness.csv`.
- `reversal_fast.py`: sweeps 864 configurations and evaluates the reversal inequality at the corner $\tau = (0, 0)$. Output: `reversal_fast.csv`.

## C.2 Supporting statistics

Table 2: Summary of committed numerical artifacts.

| File | Rows | Scope | Violations | De |
|---|---|---|---|---|
| `a2_sweep.csv` | 2205 | $[0, 5]^2$, 5 scenarios, $\kappa_T = 0$ | 0 | As |
| `two_sided_cursedness.csv` | · | $\kappa_T \in \{0.3, 0.5\}$ | 0 on grid | A2 |
| `reversal_fast.csv` | 864 | $\tau = (0, 0)$ corner | 0, min slack $7 \times 10^{-4}$ | (R |
| `welfare_loss_table.csv` | 45 | $\kappa_M \in \{0.3, 0.5, 0.6\}$, $c_0 \in \{0.01, 0.05, 0.2\}$ | n/a | $L/$ |
| `prop2_2component_verify.txt` | — | $S_1, S_2, S_3$, $\kappa_M \in \{0, 0.25, 0.5, 0.75, 1\}$ | n/a | Co |
| `cost_robustness.csv` | · | quadratic / linear / log cost | 0 on grid | Mo |

### C.3 Extended verification of Corollary 1

Table 1 in Section 4.1 reports the three interior $\kappa_M$ values $\{0.25, 0.5, 0.75\}$ for scenarios $S_1$ and $S_3$. The committed text file `prop2_2component_verify.txt` also reports $S_2$ (asymmetric prior precisions $\tau_{\theta A} = 0.5, \tau_{\theta B} = 2, \alpha = \tau_z = a = 1,\ \tau_A = \tau_B = 0.5$), with non-zero interior gaps of the same qualitative structure. At each scenario, the gap vanishes exactly at $\kappa_M \in \{0, 1\}$ and is nonzero at every interior $\kappa_M$ tested. The rational-versus-polynomial structure of Corollary 1 is confirmed at every tested interior point in every tested scenario.