

Disagreement in Asset Prices with Endogenous Beliefs

An Information-Theoretical Framework and Applications



Ignacio Muñoz-Alonso

A thesis submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

EDHEC Business School

February 2026

Dissertation Committee

| | | |
|--------------------|-----------------------|-------------------------|
| Chair: | Nikolaos Tassaromatis | EDHEC Business School |
| Supervisor: | Abraham Lioui | EDHEC Business School |
| External Examiner: | Savitar Sundaresan | Imperial College London |
| Member: | Mirko Rubín | EDHEC Business School |

© 2026 Ignacio Muñoz-Alonso

All rights reserved.

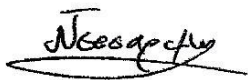
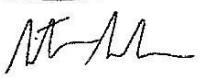




APPROVAL FORM

CANDIDATE'S NAME: Ignacio MUNOZ-ALONSO

TITLE OF DISSERTATION: "Disagreement in Asset Pricing with Endogenous Beliefs:
An Information Theoretical Framework"

APPROVAL:

| | | |
|------------------------------|-------------------|---|
| <u>Nikolaos TESSAROMATIS</u> | Chair |  |
| <u>Savitar SUNDARESAN</u> | External Examiner |  |
| <u>Abraham LIQUI</u> | Supervisor |  |
| <u>Mirco RUBIN</u> | Member |  |

DATE: 6 February 2026

To Anaïs, the light of my days, for her unconditional support and her asymptotic patience. This dissertation would not have been the same without her, and neither would the person who wrote it.

To my two sons, Sacha and Edu, in the hope that they will always preserve their curiosity and appreciation of science and knowledge.

In loving memory of my brother Sacha.

Acknowledgements

This thesis marks the end of a long intellectual journey, but also the culmination of many years of encouragement, patience, and generosity from those who accompanied me along the way.

First and foremost, I would like to express my gratitude to my supervisor, Professor Abraham Lioui, for his guidance, rigor, and encouragement throughout this process. I am especially grateful for the freedom to pursue ideas that were sometimes unconventional, but always sincerely motivated.

I would also like to express my sincere gratitude to Professor Savitar Sundaresan of Imperial College, who served as external examiner. I am especially thankful for his careful engagement with this work, and for the generosity and support reflected in his comments. His thoughtful observations and encouraging perspective were greatly appreciated, and it has been a true honour to receive them.

I am also deeply thankful to the members of my dissertation committee for their time, attention, and willingness to engage with this research. Their comments, questions, and perspectives helped sharpen the argument and improve the final manuscript.

I would like to thank EDHEC Business School for providing the academic environment in which this thesis was written. I am grateful to the faculty and staff, specially Brigitte Bogaerts-Chevillotte and Julie Smadja who, in different ways, contributed to making this path possible.

Finally, I remain grateful for the opportunity to have pursued a question that has fascinated me for years. Whatever the limitations of this thesis, it was written with conviction, curiosity, and the hope of contributing, however modestly, to a better understanding of economics and financial markets.

Contents

| | |
|---|-----------|
| Acknowledgements | 3 |
| I Disagreement in Asset Pricing with Endogenous Beliefs: An Information Theoretical Framework | 8 |
| Abstract | 9 |
| 1 Motivation, positioning and contribution of this work | 11 |
| 2 Notation | 16 |
| 3 Contemporary Theories of Disagreement in Asset Pricing | 23 |
| 3.1 Analytical models of disagreement | 23 |
| 3.2 Computational and high-dimensional models | 26 |
| 3.3 Noise traders and the limits of full rationality | 27 |
| 3.4 Positioning of the present framework | 28 |
| 4 Entropy in Finance: The Quest for Patterns | 28 |
| 4.1 Preliminaries: information, entropy, and the notion of patterns . . . | 29 |
| 4.2 Pattern detection: dependence, predictability, and divergence | 31 |
| 4.3 Uses of entropy in finance | 32 |
| 5 An Entropy-Based Stochastic Discount Factor. A foundational construction before introducing risk preferences | 36 |
| 5.1 Derivation of the Exogenous SDF (see Appendix A.1) | 37 |
| 5.2 Subjective Valuation in an Objective Price System | 44 |
| 5.3 Belief Distortion, No-Arbitrage, and the Valuation Gap | 47 |
| 6 The Value of Information in the Construction of Preferences | 52 |
| 6.1 Information-Based Utility Functions | 53 |
| 6.2 Risk Preferences and the Information Structure of Beliefs | 63 |
| 6.3 Optimal Endogenous Belief Distributions | 67 |
| 6.4 Optimal Beliefs and Subjective Entropy: A Numerical Illustration . | 69 |
| 7 An Endogenous InSDF: Pricing with Informational Disagreement | 73 |
| 7.1 Equilibrium concept (partial equilibrium) | 73 |
| 7.2 Summary derivation of the Endogenous InSDF | 75 |
| 7.3 Sources of Risk: Consumption Growth as the ρ -Channel | 79 |

| | | |
|-----------|---|------------|
| 7.4 | Discussion and Economic Implications | 81 |
| 7.5 | Summary and Conclusion | 84 |
| 8 | Conclusion: Information, Belief Distortion, and Asset Pricing Equilibrium | 87 |
| | | |
| II | Disagreement in Asset Pricing with Endogenous Be- liefs: Empirical Implications and Applications | 90 |
| | Abstract | 91 |
| 9 | Introduction | 92 |
| 10 | A Disagreement Binomial Option-Pricing Model | 95 |
| 10.1 | CRR Pricing with Disagreement: the κ -Tilt | 96 |
| 10.2 | Visualizing Belief-Tilted Option Pricing: The Role of κ | 101 |
| 10.3 | Comparison between the κ -Tilted and CRR Models | 104 |
| 10.4 | Disagreement, Arbitrage and Market Activity | 106 |
| 10.5 | Concluding Remarks | 108 |
| 11 | The Information CAPM | 110 |
| 11.1 | The InCAPM: Disagreement and Informational Risk | 111 |
| 11.2 | From Exogenous Benchmarks to Endogenous Aggregation | 116 |
| 11.3 | Discussion and Sensitivity Analysis | 118 |
| 11.4 | The Arrow–Debreu Benchmark: Full Agreement and Zero Informa- tional Beta | 121 |
| 11.5 | Empirical Implications | 124 |
| 11.6 | Concluding Remarks | 127 |
| 12 | A Plausible Entropy-Based Resolution to the Ellsberg Paradox | 128 |
| 12.1 | The Ellsberg Paradox and Competing Explanations | 128 |
| 12.2 | Ambiguity Premium | 130 |
| 12.3 | Utility Comparison and Resolution | 134 |
| 12.4 | Calibration of the Ambiguity Premium Δ | 139 |
| 12.5 | Ambiguity Premiums in the Information Stochastic Discount Factor | 143 |
| 12.6 | Concluding Remarks | 146 |
| 13 | Belief Convergence as a Microfoundation for Informational Effi- ciency | 148 |
| 13.1 | Framework and Price-Consistent Belief Dynamics | 149 |
| 13.2 | No-Signal Dynamics, Conservation, and Disagreement Contraction . | 150 |

| | | |
|-----------|---|------------|
| 13.3 | From Disagreement Convergence to Informationally Efficient Markets | 153 |
| 13.4 | The Role of Risk Aversion in the Convergence Process | 154 |
| 13.5 | Symmetry, Maximum Entropy, and the Special Uniform Benchmark | 157 |
| 13.6 | Summary and Conclusion | 158 |
| 14 | Final Remarks | 160 |
| 15 | Extensions and Future Work | 162 |
| | Appendices | 165 |
| | Appendices | 165 |
| A | Appendices to Part I: The Model | 165 |
| A.1 | Derivation of the Entropy-based SDF | 165 |
| A.2 | Verification of kernel properties | 167 |
| A.3 | Decomposition of a disagreement price premium | 168 |
| A.4 | Derivation of Entropy-Augmented Utility | 169 |
| A.5 | Derivation of Utility-Maximizing Subjective Probabilities | 170 |
| A.6 | Full Derivation of the Endogenous InSDF | 171 |
| B | Appendices to Part II: Applications | 173 |
| B.1 | From the Common Reference Measure $\bar{\nu}$ to the κ -Tilted Binomial Measure and SDF | 173 |
| B.2 | Derivation of the Information CAPM (InCAPM) | 176 |
| B.3 | KL-Minimizing Aggregate Benchmark under Heterogeneous Beliefs | 178 |
| B.4 | Link to the Entropy-Augmented Utility of Paper 1 | 178 |
| B.5 | Closed-Form Derivation of the Ambiguity Premium $\Delta(\alpha, \beta)$ | 179 |
| B.6 | Equivalence Between the Ambiguity-Based and Information-Based SDFs | 180 |
| B.7 | Maximum Entropy, Symmetry of Beliefs, and Informational Efficiency | 180 |
| | Bibliography | 184 |

Part I

Disagreement in Asset Pricing with Endogenous Beliefs: An Information Theoretical Framework

Abstract

*“In all chaos there is a cosmos,
in all disorder a secret order.”*

— *C. G. Jung*

“Entropy is the price of structure.”

— *Ilya Prigogine and Isabelle Stengers*

This first part develops a micro-founded asset-pricing framework for incomplete markets in which subjective beliefs are chosen optimally. Agents select probability measures to maximize an entropy-weighted objective, so preferences discipline how probability mass is allocated across states and how uncertainty is encoded. Heterogeneity in these optimal belief tilts generates disagreement even under a shared no-arbitrage benchmark and enters prices through an information-based stochastic discount factor (InSDF). In equilibrium, disagreement is priced as an informational premium: assets that perform poorly in states where belief wedges are large must offer higher expected returns.

A central aim of this thesis is to address the unresolved problem of measuring disagreement by combining objective objects, such as probability distributions, with subjective elements, such as preferences, thereby yielding a quantifiable metric for characterizing disagreement.

Relative to the complete-markets benchmark with a common prior, the model explains persistent belief heterogeneity without relying on exogenous private signals. Risk-averse agents choose less concentrated (higher-entropy) beliefs, while more risk-tolerant agents tilt probability mass toward favorable outcomes. A central object is the disagreement wedge κ , a closed-form, state-dependent likelihood ratio between agents’ optimal beliefs that maps directly into the pricing kernel.

Finally, the paper characterizes conditions under which belief disagreement converges. In the absence of new information, entropy-based discipline shrinks belief wedges and recovers the benchmark pricing kernel. Overall, the framework links preferences, subjective probabilities, and asset prices in a tractable no-arbitrage setting and interprets disagreement as a priced informational friction.

Keywords: disagreement, endogenous beliefs, asset pricing, entropy, ambiguity, informational efficiency

JEL Classification: D53, D81, D84, G12

1 Motivation, positioning and contribution of this work

Classical asset-pricing models typically rely on representative-agent aggregation or rational expectations, and therefore abstract from disagreement as a driver of prices, trading volume, and volatility. This abstraction is not innocuous. Under the benchmark logic of common priors and common knowledge of rationality, purely speculative trade is hard to sustain: once information is correctly incorporated into prices, trading based only on “different opinions” should vanish (the no-trade benchmark; Milgrom and Stokey (1982)). Conversely, in rational-expectations settings with fully revealing prices, informational trading incentives collapse unless some friction prevents complete revelation, as in the classic impossibility insight of Grossman and Stiglitz (1980).

Real markets sit precisely in the gap between these benchmarks. Investors often hold divergent views, and the persistence of volume suggests that disagreement is not merely a transient computational error but a structural feature of environments with incomplete model specification. This paper develops a theoretical framework that formalizes such disagreement as a disciplined form of belief distortion anchored to a common no-arbitrage benchmark, and studies its implications for asset prices.

The central claim is that disagreement can be priced. More precisely, belief heterogeneity—when anchored to a common no-arbitrage benchmark—induces state-dependent distortions in subjective probabilities. These distortions admit an informational interpretation: agents impose different perceived entropy levels on the same uncertainty, reflecting heterogeneous risk attitudes and informational discipline. In this setting, returns compensate not only for fundamental risk but also for exposure to disagreement-intensive states, i.e. states in which valuation wedges across beliefs are large. The benchmark plays a dual role throughout: it rules out arbitrage by anchoring prices to a common linear operator, and it provides a common reference object relative to which disagreement is defined (agents must disagree *about* something).

This paper lies at the intersection of six strands of research.

Disagreement and asset prices. Beginning with Harrison and Kreps (1978) and Miller (1977), a large literature shows that heterogeneous beliefs can sustain trading volume and affect valuations. More recently, Hong and Sraer (2016) link disagreement to risk premia through “speculative betas.” Our contribution is to provide a closed-form informational kernel that isolates the pricing impact of belief divergence without relying on leverage or short-sale constraints. A key departure is that beliefs are not imposed exogenously: they are endogenized from preferences.

No-trade benchmarks and rational-expectations limits. The no-trade benchmark formalizes why disagreement is often ruled out in classical models: under common priors and common knowledge of rationality, purely speculative trade cannot be sustained (Milgrom and Stokey (1982)). At the same time, rational-expectations equilibria face the opposite knife-edge: if prices fully reveal information, incentives to acquire and trade on information vanish, so trading volume requires frictions that prevent complete revelation (Grossman and Stiglitz (1980)). Our framework does not “solve” these benchmarks; rather, it identifies a tractable departure—disciplined belief distortion anchored to a shared no-arbitrage operator—that restores a role for disagreement in state pricing while retaining the representative-agent benchmark as a limiting case when belief wedges collapse.

Heterogeneous beliefs in no-arbitrage frameworks. Recent models (e.g. Xu (2017), Bordalo et al. (2022)) show that agents can disagree while respecting no-arbitrage when pricing is disciplined by a common benchmark operator. Our *Information Stochastic Discount Factor* (InSDF) provides an explicit construction of such a kernel—and a measurable disagreement premium—rooted in informational distance rather than ad hoc preference heterogeneity.

Information theory in economics. Building on the insight that Kullback–Leibler divergence (KLD) measures the cost of deviating from a reference model (Sims (2003), Hansen and Sargent (2008)), we treat entropy as an economically meaningful object. We also draw on an axiomatic motivation for KL-based disagreement

measures: Zanardo (2017) provides a formal characterization of disagreement functionals based on natural invariance and separability requirements, and shows how likelihood-ratio-based statistics—with KLD as a canonical limit case—arise as disciplined measures of belief divergence. Unlike robustness or rational-inattention models, we use KLD to connect belief tilts directly to state prices, thereby turning an information cost into a pricing wedge.

Incomplete information and learning. Standard models emphasize learning and convergence to fundamentals. In contrast, Eyster et al. (2019) show that traders may mislearn or ignore what prices reveal, allowing belief divergence to persist. Our framework similarly allows disagreement to be sustained and priced through informational frictions, even under a shared no-arbitrage benchmark.

Utility-based SDFs. Traditional stochastic discount factors arise from marginal utility (Lucas Jr (1978)), while entropy-adjusted preferences (e.g. Hansen and Maenhout (2006)) import KL penalties as taste parameters. We reverse the order: we first construct an information-distortion kernel under a pricing anchor, and then micro-found it by endogenizing beliefs through entropy-weighted utility. This separation clarifies which components of pricing are informational versus preference-driven.

By synthesizing these literatures, the paper provides a tractable link between belief heterogeneity and state pricing, offers a measurable disagreement premium, and prepares the ground for equilibrium applications that integrate information theory with classical asset pricing.

A central objective of this thesis is to address a longstanding unresolved issue in the disagreement literature: how disagreement should be measured in a way that is both economically meaningful and mathematically disciplined. Existing approaches often emphasize either objective components, such as differences in probability distributions, or subjective components, such as heterogeneous preferences, without fully integrating the two. The framework developed here seeks to combine both dimensions. By linking probabilistic beliefs to preference-based

distortions, it produces a quantifiable metric through which disagreement can be formally characterized and compared across agents.

The main contributions of this paper are fourfold:

1. We derive a closed-form informational kernel by minimizing Kullback–Leibler divergence subject to a pricing anchor. This yields an entropy-based change of measure, which we term the Information SDF.
2. We show that the InSDF preserves standard no-arbitrage properties while accommodating heterogeneous beliefs. Disagreement is anchored by a benchmark restriction shared across agents (e.g. a risk-free or risk-neutral pricing anchor), so markets remain arbitrage-free despite belief dispersion.
3. We endogenize beliefs through entropy-weighted utility maximization. This produces optimal subjective probabilities in closed form and an endogenous disagreement wedge, linking risk attitudes to systematic belief tilts.
4. We characterize how disagreement maps into prices through state reweighting: the model isolates an informational premium associated with exposure to states in which belief ratios (and thus valuation wedges) are large.

The rest of the thesis is organized as follows.

Part I develops the theoretical framework. *Section 2* reviews entropy and Kullback–Leibler divergence as tools for characterizing informational distortion and disagreement. *Section 3* derives an exogenous information-based stochastic discount factor by minimizing Kullback–Leibler divergence subject to a pricing constraint. *Section 4* introduces entropy-weighted utility and shows how optimal subjective beliefs arise from an expected-utility problem with informational discipline. *Section 5* embeds this mechanism into an intertemporal asset-pricing setting and derives the endogenous Information SDF (InSDF), in which disagreement enters the pricing kernel through a state-dependent likelihood ratio between agents’ optimal beliefs.

Part II develops the main applications and implications of the framework. *Section 6* applies the model to a disagreement-based binomial option-pricing setting

as an extension of the Cox–Ross–Rubinstein benchmark. *Section 7* introduces the Information CAPM (InCAPM), where disagreement becomes an additional source of systematic risk through an informational beta. *Section 8* studies ambiguity and proposes an entropy-based resolution of the Ellsberg paradox. *Section 9* examines belief convergence and shows how informational efficiency can emerge endogenously in the absence of new information. *Section 10* concludes and discusses extensions and future research.

Taken together, the thesis moves from the theoretical foundations of disagreement under endogenous beliefs to a set of asset-pricing applications in which disagreement, ambiguity, and convergence acquire explicit pricing content.

2 Notation

Sets, indices, timing

- \mathcal{S} : state space; $s \in \mathcal{S}$ denotes a state.
- $t = 0, 1, \dots$: time index.
- i, j : agent indices.

Random variables, prices, returns

- $X_{t+1}(s)$: generic traded payoff at date $t + 1$ in state s .
- $P_t(X_{t+1})$: time- t price of payoff X_{t+1} .
- P_t^A : time- t price of asset A .
- $X_{t+1}^A(s)$: payoff of asset A at $t + 1$ in state s .
- $R_{t+1}^A(s) := X_{t+1}^A(s)/P_t^A$: gross return on asset A .
- $r_{t+1}^A(s) := \log R_{t+1}^A(s)$: log return (only when used).
- $\rho_{t+1}^B(s)$: benchmark claim log return at $t + 1$ in state s (pricing anchor object).
- $\hat{\rho}^B$: target benchmark log return used in the KL-constrained problem.

Probabilities and expectations

- $\bar{\nu}$: benchmark (common) probability measure.
- η_i : subjective belief (probability measure) of agent i .
- η_i^* : optimal (endogenized) belief of agent i .
- $\mathbb{E}^\mu[\cdot], \mathbb{E}_t^\mu[\cdot]$: expectation under measure μ (unconditional / conditional at t).

SDFs, deflators, and Radon–Nikodym (RN) derivatives

- $m_{t+1}^C(s)$: benchmark deflator (benchmark SDF component) at $t + 1$ in state s .

- $m_{t+1}(s)$: generic SDF used for pricing under $\bar{\nu}$.
- $m^{\eta_i|\bar{\nu}}(s) := \frac{d\eta_i}{d\bar{\nu}}(s)$: RN derivative (likelihood ratio) of η_i w.r.t. $\bar{\nu}$.
- $m^{\eta|\bar{\nu}}$: shorthand for the benchmark distortion density (paper-specific; equals $m^{\eta_i|\bar{\nu}}(s)$ state distortion density (when used that way).
- $m^{\eta_i|\eta_j}(s) := \frac{d\eta_i^*}{d\eta_j^*}(s)$: pairwise RN derivative across optimal beliefs.

Disagreement wedge and information objects

- $\kappa(s)$: state-dependent disagreement (information) wedge, defined by $m^{\eta_i|\eta_j}(s) = e^{\kappa(s)}$.
- κ : scalar (state-independent) version of $\kappa(s)$ when specialized.
- $D_{\text{KL}}(\mu \parallel \nu)$: Kullback–Leibler divergence (relative entropy) of μ from ν .
- $H(\mu)$: Shannon entropy of measure μ (when used).

Preferences, outcomes, consumption, and marginal utilities

- $x, x(s)$: generic outcome variable, and outcome associated with state s in a static lottery or reduced-form decision problem.
- $w, w_i(s)$: generic monetary outcome / state-contingent monetary outcome; the payoff-like object entering utility in the static notation.
- $W_i, W_{i,t}$: wealth of agent i in levels, generically and at time t .
- $C_{i,t}, C_{i,t+1}(s)$: consumption of agent i at t and at $t + 1$ in state s .
- $X_{t+1}(s), X_{t+1}^A(s)$: state-contingent payoff of a traded claim, and of asset A , at time $t + 1$ in state s .
- $u(\cdot)$: period utility function.
- $u(x)$: utility of a generic outcome x in reduced-form notation.
- $u'(C)$: marginal utility.
- $u_{i,\eta}^*(s), u_{i,\nu}^*(s)$: equilibrium marginal-utility objects under the relevant belief / benchmark (paper-specific definitions).

- $\mathbb{U}_\gamma(x; P)$: information-adjusted utility of outcome x under belief P .
- γ_i : risk aversion parameter of agent i (e.g. the CRRA coefficient when specialized).
- δ : subjective discount factor.

Convention. Throughout the thesis, x denotes a generic outcome variable, w a monetary or state-contingent outcome entering utility, W wealth in levels, C consumption, and X the payoff of a traded claim or asset. This convention keeps separate the reduced-form decision notation, the utility-relevant monetary outcome, intertemporal consumption, and asset-pricing payoffs. *Convention.* Throughout the thesis, x denotes a generic outcome variable, w denotes a state-contingent monetary outcome, W denotes wealth in levels, C denotes consumption, and X denotes the payoff of a traded claim or asset. This convention is meant to keep the distinction clear between decision-theoretic notation, monetary outcomes, intertemporal consumption, and asset-pricing payoffs.

Multipliers, pricing anchors, and state-price sensitivities

- λ_1 : Lagrange multiplier on the benchmark pricing anchor (KL-constrained distortion problem).
- $\lambda_\eta, \lambda_\nu$: multipliers associated with the agent/benchmark components in the endogenous InSDF construction (as defined in the paper).
- τ_ρ : sensitivity/loading parameter associated with the payoff/return channel $\rho(s)$ (used when introducing the payoff/return channel; paper-specific).
- ξ_i : additional multipliers/normalizers used in entropy-weighted utility (paper-specific).
- Z_i : normalization constant ensuring probabilities sum to one.

Prices, Personal Valuations, and Certainty Equivalents

- $P_t(X_{t+1}) = \mathbb{E}_t^\nu[\hat{m}_{t+1}(s)X_{t+1}(s)]$: common no-arbitrage market price of payoff X_{t+1} .

- $\text{CE}_{i,t}^{\eta_i}(Y) := u_i^{-1}(\mathbb{E}_t^{\eta_i}[u_i(Y)])$: certainty equivalent of payoff Y for agent i .
- $\Delta_{i,t}(X_{t+1}) := V_{i,t}(X_{t+1}) - P_t(X_{t+1})$: valuation wedge between personal valuation and market price of payoff X at time $(t + 1)$

Conventions. All RN derivatives are assumed to exist when used (e.g. absolute continuity / full support), in which case they coincide with probability ratios on \mathcal{S} . The benchmark $\bar{\nu}$ plays a dual role: it anchors no-arbitrage pricing through a common linear operator, and it provides the common reference relative to which disagreement is defined (agents must disagree *about* something).

Model summary: primitives, choice variables, and derived objects

For clarity, the framework developed in this paper can be organized into three layers: primitive objects, choice variables, and derived objects.

1. Primitive objects. These are taken as given.

1. **State space.** A finite state space

$$S = \{s_1, \dots, s_n\}.$$

2. **Benchmark probability measure.** A common benchmark distribution

$$\bar{\nu} \in \Delta(S),$$

which serves as the reference measure for pricing and belief distortion.

3. **Agents.** A set of agents indexed by i, j, \dots , each endowed with heterogeneous preferences.

4. **State-contingent outcomes.** For each agent i , a state-contingent wealth or consumption profile

$$w_i(s) \quad \text{or} \quad c_i(s), \quad s \in S.$$

5. **Preferences.** A utility function

$$U_i(w_i(s), \gamma_i),$$

where γ_i is the coefficient of risk aversion. In the CRRA case,

$$U_i(w, \gamma_i) = \frac{w^{1-\gamma_i} - 1}{1 - \gamma_i}.$$

6. **Information-discipline parameter.** A strictly positive parameter

$$\alpha_i > 0,$$

governing the intensity or cost of belief distortion in the entropy-weighted criterion.

7. **Time preference.** An intertemporal discount factor

$$\delta_i \in (0, 1).$$

8. **Benchmark pricing object.** Either:

- an external pricing anchor \hat{r}^B attached to a benchmark claim with state log return $r_{t+1}^B(s)$ in the exogenous construction, or
- a benchmark pricing kernel $\hat{m}_{t+1}(s)$ in the endogenous construction.

9. **Macro risk channel.** An exogenous aggregate state variable

$$\rho_{t+1}(s),$$

interpreted in the endogenous model as log consumption growth or, more generally, log marginal-utility growth.

10. **Traded assets.** State-contingent payoffs

$$X_{t+1}^A(s),$$

with prices P_t^A and returns

$$R_{t+1}^A(s) = \frac{X_{t+1}^A(s)}{P_t^A}.$$

2. **Choice variables.** These are chosen by agents within the model.

1. **Subjective beliefs.** Each agent chooses a subjective probability distribution

$$\eta_i(\cdot) \in \Delta(S), \quad \Delta(S) = \left\{ p \in \mathbb{R}^n : \sum_{s \in S} p(s) = 1, p(s) \geq 0 \forall s \in S \right\}.$$

2. **Optimal subjective beliefs.** Beliefs are selected as solutions to an entropy-weighted optimization problem, yielding

$$\eta_i^*(\cdot) \in \Delta(S).$$

3. **Derived objects.** These are implied by the primitives and the agents' optimal choices.

1. **Partition functions.** For each agent i ,

$$Z_i = \sum_{u \in S} \bar{v}(u) \exp\left(\frac{\tau_i}{u_i^*(u)}\right),$$

where $\tau_i := \theta_i/\alpha_i$ is the effective tilt parameter.

2. **Optimal belief distortions.** The optimal subjective belief of agent i takes the form

$$\eta_i^*(s) = \frac{\bar{v}(s) \exp\left(\frac{\tau_i}{u_i^*(s)}\right)}{Z_i}.$$

3. **Disagreement wedge.** The benchmark-invariant belief-ratio kernel between agents i and j is

$$m^{\eta_i|\eta_j}(s) = \frac{\eta_i^*(s)}{\eta_j^*(s)} = \exp(\kappa(s)).$$

4. **State-dependent disagreement index.**

$$\kappa(s) = \log\left(\frac{Z_j}{Z_i}\right) + \frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)}.$$

5. **Exogenous benchmark distortion kernel.** In the exogenous construction,

$$m^{\eta_i|\bar{\nu}}(s) = \frac{d\eta_i}{d\bar{\nu}}(s).$$

6. **Endogenous Information-based SDF.** The full pricing kernel is

$$m_{t+1}(s) = \frac{\delta_j}{\delta_i} \exp(\kappa(s)) \exp(-\lambda_1 \rho_{t+1}(s)).$$

7. **Euler pricing restriction.** For any traded asset A ,

$$1 = \mathbb{E}_t^{\bar{\nu}}[m_{t+1}(s) R_{t+1}^A(s)].$$

This decomposition makes explicit the logic of the paper: the primitives define the economic environment, agents choose subjective beliefs within that environment, and the resulting disagreement wedge and InSDF are derived equilibrium valuation objects.

3 Contemporary Theories of Disagreement in Asset Pricing

Modern asset-pricing theory has progressively moved away from the strong representative-agent benchmark in which all investors share the same information set, the same model of the economy, and the same conditional expectations. In the rational-expectations tradition, disagreement should vanish once public information is incorporated into prices. Hansen and Sargent describe this benchmark as a *communism of expectations*: rationality, common knowledge, and full-information efficiency jointly imply that agents converge on a single probability law and therefore on a common valuation rule (Hansen and Sargent, 2010). In that environment, persistent disagreement is difficult to reconcile with both no-arbitrage and rational behavior.

A large contemporary literature challenges this benchmark by showing that disagreement may persist even in disciplined, rational environments. The key insight is that common prices do not necessarily imply common beliefs. Once one departs from the assumption that all investors process information identically, the door opens to heterogeneous beliefs, divergent valuations, trading volume, and endogenous volatility. These theories can be organized into two broad groups: *analytical models*, which retain relatively tractable mathematical structures, and *computational models*, which allow richer forms of heterogeneity at the cost of closed form.

3.1 Analytical models of disagreement

A first analytical class is formed by *heterogeneous-belief equilibrium models*. The classical contribution is Harrison and Kreps (1978), who show that speculative trade can arise in equilibrium when agents disagree about future states, even if they are individually rational. The crucial point is that differences in expectations are sufficient to generate trading motives and speculative components in prices. Related work by Miller (1977) emphasizes that disagreement, especially when combined with

short-sale constraints, can produce overvaluation. In these models, disagreement is analytically transparent and economically powerful, but it is generally taken as exogenous: agents are assumed to disagree, yet the source of that disagreement is not modeled.

A second branch is the literature on *rational beliefs dispersion*. These models reject the idea that rationality requires unanimity. Even with common public information, agents may rationally disagree because they entertain different internally consistent models, different priors, or different views about persistence, regime changes, or the stationarity of the data-generating process. This idea appears in work such as Kurz (1994), and later in empirical and theoretical contributions including Lux (2009), Lux (2012), Alfarano et al. (2005), Bhamra et al. (2019), and Das et al. (2020). The common theme is that disagreement need not reflect irrationality; it may instead be the equilibrium outcome of rational agents operating under model uncertainty. These models are especially important because they provide a bridge between rational expectations and persistent belief heterogeneity.

A third analytical mechanism is *Bayesian learning*. Bayesian learning models preserve rational updating, but they abandon the strong implication that rational agents must immediately converge to a common probability measure. When agents start from different priors, observe noisy signals, or learn gradually from prices and from each other, divergence may persist for long periods. Contributions such as Barberis et al. (1998), Hong and Stein (1999), and Branch and Evans (2006) show that learning itself can generate rich asset-pricing dynamics. In these models, disagreement is not static but evolves over time as agents update their subjective probability assessments. This is especially relevant for finance because it permits temporary or persistent valuation gaps without abandoning disciplined probabilistic updating.

A fourth strand is *rational inattention*. Here, disagreement does not arise because agents are irrational, but because information acquisition and processing are costly. Investors optimally choose not to absorb all available information,

generating heterogeneous posterior beliefs and sluggish adjustment in prices. The foundational references are Sims (2003) and Sims (2006), with further surveys and developments in Maćkowiak et al. (2023). This literature is highly relevant for disagreement because it weakens the accessibility assumption embedded in rational expectations: public information may exist and still not be equally processed by all agents. Rational inattention can therefore generate underreaction, delayed price discovery, portfolio under-diversification, and cross-sectional heterogeneity without requiring a departure from optimization.

Closely related is the literature on *gradual information diffusion*. In Hong and Stein (1999), some investors receive or process information earlier than others, so prices adjust with delay. Hong and Stein (2007) show that common public signals, such as earnings announcements, generate abnormal trading volume because investors interpret the same information differently. This literature gives disagreement a concrete empirical dimension: it is not merely a hidden feature of beliefs, but something that appears in turnover, momentum, and delayed price reactions.

A final analytical benchmark is provided by incomplete-markets macro-finance models such as *Bewley–Aiyagari–Huggett* economies. These models are not disagreement models in the narrow sense, but they are central because they formalize heterogeneous agents under incomplete markets and idiosyncratic risk. In Huggett (1993), Aiyagari (1994), and the broader Bewley tradition, agents differ in wealth, income histories, and borrowing constraints, while aggregate prices emerge from incomplete insurance and self-insurance decisions. These models typically do not make subjective belief disagreement the main primitive, but they establish the analytical foundations for heterogeneous-agent economies in which market incompleteness is structurally central. For that reason, they are an important precursor for models, like the present one, that connect incomplete markets with endogenous subjective probabilities.

3.2 Computational and high-dimensional models

A second large class of theories moves beyond low-dimensional analytical settings and studies disagreement in richer environments through computation. This includes *heterogeneous-agent models* (HAMs) and *agent-based models* (ABMs). Their starting point is that real markets contain many agents with heterogeneous beliefs, preferences, constraints, heuristics, and interaction networks. Once this heterogeneity becomes sufficiently rich, closed-form equilibrium is generally unavailable, and computational methods become necessary.

Foundational contributions include Day and Huang (1990), Chiarella (1992), De Grauwe and Dewachter (1993), Brock et al. (1997), Brock and Hommes (2001), and Hommes (2001). In these models, prices and beliefs co-evolve endogenously. Investors may switch forecasting rules, adapt to past performance, imitate others, or revise their beliefs through social interaction and reinforcement. The central contribution of this literature is to show that disagreement can itself become an endogenous state variable of the system. Asset-price volatility, bubbles, crashes, and regime shifts may then emerge from the interaction of heterogeneous agents rather than from exogenous shocks alone (Hommes, 2006, 2013).

Agent-based finance pushes this logic further. In ABMs, equilibrium is not imposed analytically but emerges from decentralized interaction among many heterogeneous agents. This framework is especially useful when one wants to allow simultaneously for bounded rationality, institutional constraints, social networks, heterogeneous balance sheets, adaptive expectations, and incomplete markets. Reviews such as Axtell and Farmer (2022) and computational surveys around Bewley-Huggett-Aiyagari style economies such as Kirkby (2019) illustrate how far this literature has progressed. More recent work, including Farmer (2002), Farmer and Foley (2009), and related contributions, shows how complex market dynamics can be generated by realistic trading ecologies rather than by a representative agent with rational expectations.

The strength of HAMS and ABMs is realism. They allow disagreement to arise not only from heterogeneous priors, but also from endogenous interaction, market frictions, leverage, learning, and network effects. They are therefore particularly well suited to explaining stylized facts such as clustered volatility, excess trading volume, momentum, and occasional instability. Their weakness is the opposite of that of analytical models: because disagreement is embedded in a high-dimensional computational system, it is often difficult to isolate a single transparent pricing wedge or a closed-form equilibrium object.

3.3 Noise traders and the limits of full rationality

A parallel literature explains disagreement through departures from full rationality. In *noise trader* models, some investors trade on misperceptions, biased interpretations, or spurious signals rather than on correct probabilistic inference. Classic references include De Long et al. (1990) and Shleifer and Summers (1990). More recent work, such as Giglio et al. (2021) and Bhamra et al. (2022), expands this logic by examining how psychological distance, belief distortions, and behavioral biases shape subjective expectations. This literature is useful because it explains why disagreement may persist even when arbitrageurs are present: noise need not cancel instantaneously, and limits to arbitrage may allow distortions to survive long enough to matter for prices.

The broader implication is that disagreement in financial markets may arise from multiple sources simultaneously: different priors, costly information processing, gradual diffusion of public information, bounded rationality, psychological distortions, and endogenous interaction across agents. This multiplicity of mechanisms is precisely why a single representative-agent model struggles to explain observed trading volume, excess volatility, and persistent valuation dispersion.

3.4 Positioning of the present framework

The literature reviewed above establishes that disagreement is real, empirically relevant, and theoretically compatible with no-arbitrage once one departs from the strongest form of rational expectations. However, existing approaches often face a trade-off. Analytical models preserve tractability but usually impose heterogeneous beliefs exogenously. Computational models endogenize heterogeneity more realistically but often lose a transparent pricing representation. The present thesis is positioned between these two poles. It keeps the analytical discipline of a no-arbitrage pricing framework, but it endogenizes beliefs through a utility-based informational criterion. In that sense, it is closer in spirit to rational-beliefs and Bayesian-learning models than to pure noise-trader accounts, while also sharing with incomplete-markets heterogeneous-agent models the view that subjectivity is a structural response to environments in which a unique objective probability law is not available.

The contribution is therefore not merely to add another disagreement model, but to provide a micro-founded and tractable bridge between incomplete markets, endogenous subjective beliefs, and asset pricing. Disagreement does not appear as an arbitrary wedge; it is derived from the interaction between risk preferences and belief formation. This is precisely the gap left open by much of the contemporary literature.

4 Entropy in Finance: The Quest for Patterns

This section introduces entropy as the central informational concept of the thesis and clarifies why it is a natural language for studying disagreement in finance. The guiding idea is that financial markets are not only spaces of risk allocation but also spaces of interpretation: agents differ in the extent to which they believe that prices, states, or returns contain exploitable structure. Entropy provides a disciplined way to represent that dimension. In this thesis it is not used as a primitive taste, nor

merely as a descriptive statistic, but as a measure of the informational geometry of beliefs: high entropy corresponds to symmetry, dispersion, and skepticism about patterns, whereas low entropy corresponds to concentration and the perception that relevant structure exists. The section proceeds by first reviewing the basic information-theoretic objects needed later in the paper, then showing how they map naturally into financial questions of predictability, disagreement, and model specification. This discussion also prepares the analytical role entropy will play in the subsequent sections, where it becomes the bridge linking subjective beliefs, risk preferences, and pricing under a common no-arbitrage benchmark.

4.1 Preliminaries: information, entropy, and the notion of patterns

The concept of information entropy, introduced by Shannon (Shannon, 1948), provides a rigorous framework for quantifying uncertainty and informational content in probabilistic models.¹ It has since become a cornerstone of information theory; see Cover and Thomas (2006), Gray (2011), and Golan et al. (1997).

This thesis uses entropy for a specific purpose: to formalize *beliefs about patterns*. In financial markets, patterns are rarely observable objects; rather, they are conjectures about structure in data—regularities, dependence, or predictability—that investors may believe exist (or not exist). Entropy offers a disciplined way to describe this epistemic stance: *high entropy corresponds to symmetry and minimal structure*, whereas *low entropy corresponds to concentration and potentially exploitable regularity*.

Self-information (surprise). The information content of an event with probability $p_i \in (0, 1]$ is

$$I(p_i) = \log\left(\frac{1}{p_i}\right) = -\log p_i \geq 0. \quad (1)$$

¹Entropy was first developed in thermodynamics in the nineteenth century (Clausius, Boltzmann, Gibbs). While the mathematical forms are related, the interpretation here is informational rather than physical.

The logarithm ensures additivity for independent events: the information of a joint realization is the sum of the information of its components. In economic contexts, this additivity matches the log-likelihood structure of belief updating.

Entropy (expected surprise). For a discrete distribution $p = (p_1, \dots, p_n)$, entropy is expected self-information:

$$H(p) := \sum_{i=1}^n p_i I(p_i) = - \sum_{i=1}^n p_i \log p_i \geq 0, \quad (0 \log 0 := 0). \quad (2)$$

For a density f on a continuous support, the analogous quantity is differential entropy, $H(f) := - \int f(x) \log f(x) dx$, which can be negative and is not invariant to rescaling. The discrete expression (2) is the main object used below, because belief-tilting is most transparent in finite-state form and because the thesis interprets entropy as a property of *belief vectors*.

Units (bits vs. nats). Shannon often uses base-2 logarithms (bits). Throughout this thesis we use natural logs (nats), which align with the exponential structures that appear in utility and pricing kernels:

$$H_{\text{bit}} = \frac{H_{\text{nat}}}{\ln 2}, \quad 1 \text{ nat} = \frac{1}{\ln 2} \approx 1.443 \text{ bits.}$$

Entropy as a benchmark for (a lack of) patterns. Entropy is maximized by the uniform distribution on a finite support. A fair die therefore has maximum entropy:

$$H = \log 6 = \begin{cases} \ln 6 \approx 1.792 \text{ nats,} \\ \log_2 6 \approx 2.585 \text{ bits.} \end{cases}$$

Biasing the die concentrates probability mass, reduces entropy, and makes outcomes more predictable. This intuition is central for the thesis: *beliefs that assign concentrated probability mass encode perceived structure*, whereas beliefs that remain diffuse encode skepticism about predictability (symmetry).

4.2 Pattern detection: dependence, predictability, and divergence

To connect entropy to patterns formally, one needs measures that capture dependence and predictability rather than mere dispersion. Let X and Y be random variables (or, in time series, let X be a future realization and Y a set of past observations or signals). Three standard objects are:

Conditional entropy. The remaining uncertainty about X after observing Y is

$$H(X | Y) := - \sum_x \sum_y p(x, y) \log p(x | y) = H(X, Y) - H(Y).$$

When $H(X | Y)$ is low relative to $H(X)$, Y is informative about X .

Entropy rate (time-series predictability). For a stochastic process $\{X_t\}$, the entropy rate measures average uncertainty per period after conditioning on the past:

$$H_{\text{rate}} := \lim_{n \rightarrow \infty} \frac{H(X_1, \dots, X_n)}{n}.$$

In an i.i.d. process, $H_{\text{rate}} = H(X_1)$; systematic dependence reduces H_{rate} and therefore signals structure.

Mutual information. Mutual information measures the reduction in uncertainty about X due to observing Y :

$$I(X; Y) := H(X) + H(Y) - H(X, Y) = H(X) - H(X | Y) = H(Y) - H(Y | X).$$

If $I(X; Y) > 0$, the variables share information; dependence (structure) is present.

Mutual information as KL divergence (patterns as divergence from independence). A key identity links mutual information to relative entropy:

$$I(X; Y) = D_{\text{KL}}(P_{XY} \| P_X P_Y).$$

This equivalence is conceptually important for the thesis: *patterns can be framed as divergence from symmetry or independence*. When a joint law differs from the product of marginals, the departure is measured by a KL divergence.

KL divergence (belief disagreement). Given two belief models p and q on the same finite support, the Kullback–Leibler divergence is

$$D_{\text{KL}}(p\|q) = \sum_{i=1}^n p_i \log\left(\frac{p_i}{q_i}\right). \quad (3)$$

The likelihood ratio p_i/q_i identifies where beliefs disagree; the divergence aggregates these discrepancies into an informational cost. This is the formal object used later to construct belief distortions and informational wedges in pricing kernels.

Zanardo (2017) provides a rigorous axiomatic treatment of Kullback–Leibler relative entropy. In particular, the author develops a mathematically disciplined framework in which disagreement can be represented through KL divergence, an approach closely aligned with the one adopted in this paper.

Interpretation for this thesis. The empirical question “are there patterns in financial markets?” can be posed as a question about dependence (mutual information or entropy rate). The theoretical question “do agents believe there are patterns?” becomes a question about how they allocate probability mass across states and how their beliefs diverge from a benchmark or from each other. The thesis uses KL divergence and entropy-weighted belief choice to connect these questions within a no-arbitrage pricing system.

4.3 Uses of entropy in finance

Entropy has become increasingly prominent in finance as a measure of uncertainty, dispersion, and complexity (Campbell, 2014, 2017). In the context of this thesis, its appeal is structural: entropy provides a unified language for (i) information constraints, (ii) disciplined belief distortions, and (iii) the diagnosis of patterns or

regularities in market data.

Information as an economic constraint. The economic relevance of entropy was emphasized early by Marschak (Marschak, 1960, 1971, 1974). Sims formalized rational inattention, where agents face costs proportional to Shannon information when processing signals (Sims, 1998, 2003, 2006). In these models, entropy governs how finely agents can resolve the state of the world.

Belief distortion and robustness. Relative entropy appears as a disciplined way to model departures from a reference probability model. In robust-control frameworks, agents choose distorted beliefs subject to an entropy penalty (Hansen et al., 2001; Hansen and Sargent, 2011); entropy therefore measures the informational cost of model misspecification. Related approaches introduce entropy directly into preference specifications, yielding entropy-adjusted valuation kernels (Hansen and Maenhout, 2006). More recently, Bhamra et al. (2022) propose *psychological entropy* linking subjective distance to belief formation.

Entropy and the stochastic discount factor. Several contributions connect entropy to the properties of pricing kernels. Alvarez and Jermann (2005) decompose the SDF into permanent and transitory components and study the conditional entropy of the permanent component, linking lower conditional entropy to positive bond risk premia. Backus et al. (2014) distinguish between short- and long-horizon entropy, relating transitory uncertainty to return volatility. Ghosh et al. (2019) propose an entropy-based pricing kernel (I-SDF) and estimate expected returns via relative-entropy restrictions, reporting improved out-of-sample performance relative to standard factor models. Peng (2005) model optimal allocation of limited informational capacity across multiple entropy sources, connecting entropy minimization to equilibrium pricing.

Option-implied beliefs and maximum-entropy estimation. Entropy methods are widely used to recover risk-neutral densities from options. Gulko (1999,

2002) pioneer maximum-entropy estimation of implied distributions; subsequent work extends the approach by incorporating higher moments (Ardakani, 2022), cross-entropy estimators (Branger, 2004), and applications in incomplete markets (Yu, 2020). In this line, entropy is a practical tool for disciplined inference when state prices are observed indirectly.

Entropy as a diagnostic for patterns and market regimes. A large empirical literature uses entropy to quantify regularity, regime changes, and market efficiency. Olbryś and Majewska (2022) show that market entropy declines in crises, consistent with increased return regularity. (Liu et al., 2020b) document that entropy-based metrics capture shifts in trading behavior, while Ormos and Zibriczky (2014) argue that entropy measures can outperform beta in explaining equity premia. Related approaches include permutation entropy as an efficiency diagnostic (Zunino et al., 2009, 2010), evidence that low entropy predicts crashes (Risso, 2008), and entropy-based classification of market regimes (Patra and Hiremath, 2022; Liu et al., 2020a).

Portfolio choice and risk measurement. Entropy offers alternatives to mean–variance methods. Generalized cross-entropy accounts for estimation uncertainty in mean–covariance inputs (Bera and Park, 2008); entropy-based portfolio selection can improve diversification relative to traditional approaches (Yu and Lee, 2011; Bhattacharyya et al., 2009). More recent work incorporates Rényi entropy into dynamic allocation (MacLean et al., 2022) and uses entropic value-at-risk (EVaR) as a downside risk measure (Imran, 2023). Evidence suggests entropy-based diversification can improve short-horizon risk-adjusted performance (Pola, 2016; Mercurio et al., 2020).

Information asymmetries and allocation. Finally, entropy has been used to study allocation under informational frictions. Gârleanu and Pedersen (2021) use entropy differentials to analyze investor allocation between active and passive

management. A broad theoretical treatment of entropy methods in allocation, risk measures, and pricing is provided by Borwein and Zhu (2017), while Veldkamp (2011) surveys the economic value of information with entropy as a central organizing principle.

Across these applications, entropy plays two roles that are central for this thesis. First, it quantifies informational constraints and disciplined belief distortions. Second, it provides a benchmark for *patterns*: low entropy signals concentration and potential regularity, while high entropy corresponds to symmetry and the absence of exploitable structure. The framework developed in the next sections builds on these two roles to model disagreement as a priced, preference-driven belief distortion under no-arbitrage.

5 An Entropy-Based Stochastic Discount Factor. A foundational construction before introducing risk preferences

This section develops a closed-form exogenous stochastic discount factor (SDF) derived from informational constraints. We begin with two probability distributions: an agent-specific belief distribution η_i and a benchmark distribution $\bar{\nu}$, linked via the Kullback–Leibler divergence (KLD). Minimizing the KLD subject to a pricing constraint yields a reweighting kernel with properties analogous to a SDF.

The KLD quantifies the informational cost of distorting one belief distribution into another—here interpreted as the cost for agent i to adopt the pricing benchmark $\bar{\nu}$ instead of their own subjective beliefs η_i . Although both distributions are treated as exogenous in this section, the structure anticipates the endogenous multi-agent framework developed later, where belief distortions arise from utility maximization under informational frictions and become genuinely cross-sectional objects (i.e. comparisons of η_i versus η_j).

At this stage, $\bar{\nu}$ may be interpreted flexibly: it could represent a competing belief, an externally provided reference measure, the physical (data-generating) distribution, or even a risk-neutral probability measure. These interpretations will be developed further in subsequent sections as the kernel is embedded in a more general asset-pricing framework.

While not yet preference-based, this construction lays the conceptual foundation for the utility-adjusted framework that follows, in which belief formation is endogenized and belief *dispersion across agents* becomes an additional source of valuation wedges and trade. For this reason we refer to the present kernel as an *exogenous* stochastic discount factor, to distinguish it from the endogenous formulation developed in Section 5.

Benchmark distortion (one-agent object). At this stage, it is important

to be explicit about what the present construction is—and what it is not. The minimization problem below confronts a single agent’s subjective belief distribution η_i against a benchmark $\bar{\nu}$ under a common pricing anchor \hat{r}^B applied to a fixed reference benchmark claim with state log return $r_{t+1}^B(s)$. The object $r_{t+1}^B(s)$ plays a dual role: it is (i) the *no-arbitrage anchor statistic* that disciplines admissible belief distortions through the restriction $\mathbb{E}_{\eta_i}[r_{t+1}^B] = \hat{r}^B$, and (ii) the *common payoff mapping* about which agents can disagree through different state weights—i.e., disagreement is defined relative to the same benchmark claim. In that precise sense, the kernel $m^{n_i|\bar{\nu}}$ is a *benchmark distortion* (a one-agent object): it quantifies how agent i must reweight the benchmark reference measure in order to satisfy the external no-arbitrage anchor. This is not yet “true disagreement” in the literal cross-sectional sense of η_i versus η_j ; rather, it is a disciplined way to represent *belief distortion relative to a shared reference*.

This clarification matters conceptually because once $\bar{\nu}$ is common to all agents, it cancels out of any *pairwise* comparison: disagreement across agents is ultimately encoded in $\log(\eta_i/\eta_j)$ rather than in $\log(\eta_i/\bar{\nu})$. We therefore keep $\bar{\nu}$ here as an admissible common coordinate system for pricing, while introducing below the benchmark-invariant objects that capture genuine disagreement.

5.1 Derivation of the Exogenous SDF (see Appendix A.1)

A foundational construction before introducing risk preferences

Let $S = \{s_1, \dots, s_n\}$ denote a finite state space. Consider a traded claim that delivers a state-contingent terminal payoff $A_T(s)$ at date T if state $s \in S$ is realized, and let P_t denote its price at date $t < T$. The associated *gross return* of this claim in state s is

$$R(s) := \frac{A_T(s)}{P_t}.$$

We also introduce the generic *log-return convention*

$$\rho(s) := \ln R(s) = \ln\left(\frac{A_T(s)}{P_t}\right). \quad (4)$$

This definition fixes notation throughout: $A_T(s)$ denotes payoffs, $R(s)$ denotes gross returns, and $\rho(s)$ denotes log returns.

Pricing anchor in r^B . To impose a common no-arbitrage discipline (as in the slides), we now fix a *reference benchmark claim* with state log return

$$r_{t+1}^B(s) := \log R_{t+1}^B(s),$$

and we require that distorted beliefs reproduce a common target expected benchmark log return \hat{r}^B . Importantly, $r_{t+1}^B(s)$ refers to this fixed reference benchmark claim (not to the generic asset being valued later).

Let $\bar{\nu} = (\bar{\nu}_1, \dots, \bar{\nu}_n)$ and $\eta_i = (\eta_{i,1}, \dots, \eta_{i,n})$ be two probability distributions over S , representing a benchmark reference distribution and the subjective beliefs of agent i , respectively. Here $\eta_i(s)$ is the probability that agent i assigns to state s , while $\bar{\nu}(s)$ is an external standard (e.g., a historical distribution, another agent's beliefs, a statistical prior, the physical distribution, or a risk-neutral measure) against which informational distortions are evaluated.

The divergence between η_i and $\bar{\nu}$ is measured by the KLD:

$$D_{\text{KL}}(\eta_i \parallel \bar{\nu}) = \sum_{s \in S} \eta_i(s) \ln\left(\frac{\eta_i(s)}{\bar{\nu}(s)}\right). \quad (5)$$

The agent chooses η_i to minimize this informational cost subject to two constraints: (i) probability coherence and (ii) consistency with a *pricing anchor* \hat{r}^B on the

benchmark claim:

$$\min_{\eta_i(\cdot) \geq 0} \sum_{s \in S} \eta_i(s) \ln \left(\frac{\eta_i(s)}{\bar{\nu}(s)} \right) \quad (6)$$

$$\text{s.t.} \quad \sum_{s \in S} \eta_i(s) = 1, \quad (7)$$

$$\sum_{s \in S} \eta_i(s) r_{t+1}^B(s) = \hat{r}^B. \quad (8)$$

The first constraint ensures that η_i is a valid probability distribution. The second enforces consistency with a target expected *benchmark log return* \hat{r}^B common across agents on the reference benchmark claim. Depending on context, \hat{r}^B may coincide with the log risk-free return (yielding a risk-neutral restriction) or with the benchmark return implied by replication.

Crucially, \hat{r}^B is an *external pricing requirement* common to all agents (a no-arbitrage anchor), not a belief-driven forecast: agents may disagree about state probabilities, but their distorted beliefs must reproduce the same anchor.

Clarification (no circularity, no implicit risk neutrality, no arbitrage).

The statistic $r_{t+1}^B(s)$ entering the KL-tilt is defined for a fixed *reference benchmark* claim, so the kernel is pinned down by the common reference pair $(\bar{\nu}, \hat{r}^B)$ and does not depend on the price of the payoff being valued—hence no circularity. The resulting $m^{\eta_i|\bar{\nu}}$ is an informational change of measure rather than a marginal-utility SDF, so it does not presume risk neutrality; preference curvature enters only through certainty equivalents (and in the endogenous construction later). Finally, because the benchmark is common and the kernel is positive and normalized, the implied pricing operator remains arbitrage-free even under heterogeneous beliefs.

Form the Lagrangian with multipliers λ_0 and λ_1 for the two constraints,

$$\mathcal{L}(\eta_i; \lambda_0, \lambda_1) = \sum_{s \in S} \eta_i(s) \ln \left(\frac{\eta_i(s)}{\bar{\nu}(s)} \right) + \lambda_0 \left(\sum_{s \in S} \eta_i(s) - 1 \right) + \lambda_1 \left(\sum_{s \in S} \eta_i(s) r_{t+1}^B(s) - \hat{r}^B \right).$$

The first-order condition with respect to $\eta_i(s)$ yields

$$\ln\left(\frac{\eta_i(s)}{\bar{\nu}(s)}\right) + 1 + \lambda_0 + \lambda_1 r_{t+1}^B(s) = 0,$$

and therefore an exponential tilt of the benchmark distribution:

$$\eta_i(s) = \frac{\bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}, \quad Z(\lambda_1) = \sum_{u \in S} \bar{\nu}(u) e^{-\lambda_1 r_{t+1}^B(u)}, \quad (9)$$

where $Z(\lambda_1)$ is the normalizing partition function. The multiplier λ_1 is chosen to satisfy the pricing constraint in (8).

Define the exponentially tilted benchmark

$$\bar{\nu}_{\lambda_1}(s) = \frac{\bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}, \quad \text{so that} \quad \eta_i(s) = \bar{\nu}_{\lambda_1}(s).$$

Let

$$f(\lambda_1) := \sum_{s \in S} \bar{\nu}_{\lambda_1}(s) r_{t+1}^B(s) = \frac{\sum_{s \in S} \bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)} r_{t+1}^B(s)}{Z(\lambda_1)}.$$

Then the pricing anchor is equivalently $f(\lambda_1) = \hat{r}^B$. Moreover,

$$f'(\lambda_1) = -\text{Var}_{\bar{\nu}_{\lambda_1}}(r_{t+1}^B) < 0,$$

so the solution for λ_1 is unique whenever $\text{Var}_{\bar{\nu}_{\lambda_1}}(r_{t+1}^B) > 0$.

The implied likelihood ratio (Radon–Nikodym derivative). The likelihood ratio of the tilted belief relative to the benchmark defines the exogenous SDF:

$$m^{\eta_i|\bar{\nu}}(s) = \frac{\eta_i(s)}{\bar{\nu}(s)} = \frac{e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)} = \frac{e^{-\lambda_1 r_{t+1}^B(s)}}{\sum_{u \in S} \bar{\nu}(u) e^{-\lambda_1 r_{t+1}^B(u)}}. \quad (10)$$

Appendix A.2 shows that this kernel satisfies four defining properties of a valid SDF:

1. **Positivity:** $m^{\eta_i|\bar{\nu}}(s) > 0$ for all $s \in S$.

2. **Normalization:** $\sum_{s \in S} \bar{\nu}(s) m^{\eta_i|\bar{\nu}}(s) = 1$.

3. **Pricing consistency (anchor discipline):** $\sum_{s \in S} \bar{\nu}(s) m^{\eta_i|\bar{\nu}}(s) r_{t+1}^B(s) = \hat{r}^B$.

4. **Uniqueness:** the solution for λ_1 is unique since $f'(\lambda_1) < 0$ whenever $\text{Var}_{\bar{\nu}_{\lambda_1}}(r_{t+1}^B) > 0$.

The kernel quantifies how much agent i 's belief differs from the benchmark in each state. Intuitively, $m^{\eta_i|\bar{\nu}}(s) > 1$ (resp. < 1) means that agent i overweights (resp. underweights) state s relative to the benchmark. The multiplier λ_1 acts as a *shadow price of belief distortion*: it governs how sharply the agent must tilt probability mass toward low- r^B states (and away from high- r^B states) in order to satisfy the common pricing anchor \hat{r}^B .

Thus, while the kernel mimics the functional shape of a risk-averse, utility-based SDF, here it arises purely from an informational distortion discipline plus a pricing restriction. For this reason, we refer to it as the *exogenous* SDF in this section.

From benchmark distortion to true disagreement: relative tilts and pairwise divergences

Because the benchmark $\bar{\nu}$ is common across agents in the pricing system, it should be understood as a shared *coordinate system* rather than as the locus of disagreement. To formalize this, define for each agent j the benchmark tilt (log-likelihood ratio)

$$\phi_j(s) := \log\left(\frac{\eta_j(s)}{\bar{\nu}(s)}\right), \quad s \in S.$$

This object summarizes how agent j distorts the benchmark state probabilities. However, *true disagreement across agents* is captured by the *relative tilt*

$$\Delta\phi_{ij}(s) := \phi_i(s) - \phi_j(s) = \log\left(\frac{\eta_i(s)}{\eta_j(s)}\right),$$

which is benchmark-invariant: the common reference $\bar{\nu}$ cancels out identically once we compare two agents.

A natural pairwise measure of disagreement is therefore the Kullback–Leibler divergence between agents’ subjective beliefs,

$$D_{\text{KL}}(\eta_i \parallel \eta_j) = \sum_{s \in S} \eta_i(s) \log \left(\frac{\eta_i(s)}{\eta_j(s)} \right),$$

and, if desired, its symmetrized version

$$D_{\text{KL}}(\eta_i \parallel \eta_j) + D_{\text{KL}}(\eta_j \parallel \eta_i).$$

These are the objects that measure disagreement in the literal sense: one agent’s belief is confronted against another’s, without any residual dependence on $\bar{\nu}$.

Moreover, because each η_j is defined through an exponential tilt of the *same* benchmark under the common pricing anchor \hat{r}^B , the ratio of two agents’ belief-distortion kernels is itself benchmark-free:

$$\frac{m^{\eta_i | \bar{\nu}}(s)}{m^{\eta_j | \bar{\nu}}(s)} = \frac{\eta_i(s) / \bar{\nu}(s)}{\eta_j(s) / \bar{\nu}(s)} = \frac{\eta_i(s)}{\eta_j(s)}.$$

If additional constraints are imposed—e.g., on higher moments or other statistics—then the solution takes the general exponential-family form

$$\eta_i(s) = \frac{\bar{\nu}(s) \exp \left(-\lambda_1 r_{t+1}^B(s) - \sum_{k=2}^m \lambda_k \phi_k(s) \right)}{\sum_{u \in S} \bar{\nu}(u) \exp \left(-\lambda_1 r_{t+1}^B(u) - \sum_{k=2}^m \lambda_k \phi_k(u) \right)}, \quad (11)$$

where each $\phi_k(s)$ is a statistic to be matched (e.g., variance, skewness), and λ_k is the associated Lagrange multiplier. In this way, higher-order features of benchmark states enter as exponential tilts on $\bar{\nu}$.

Variance example. Introducing a second-moment constraint by setting $\phi_2(s) = (r_{t+1}^B(s) - \hat{r}^B)^2$ forces the distorted beliefs to reproduce a target benchmark return variance. The multiplier λ_2 then governs the dispersion of beliefs: larger $|\lambda_2|$

increases the tilt toward states with extreme benchmark realizations. Higher-order constraints (e.g., skewness, kurtosis) are incorporated analogously by choosing appropriate ϕ_k .

| Multiplier | Economic interpretation |
|-------------|--|
| λ_1 | Shadow cost of matching the mean (anchor \hat{r}^B) |
| λ_2 | Shadow cost of matching variance (volatility/dispersion) |
| λ_3 | Shadow cost of matching skewness (asymmetry) |
| λ_4 | Shadow cost of matching kurtosis (tail thickness) |
| \vdots | Higher-order features as additional informational costs |

Economically, the vector of multipliers $(\lambda_1, \lambda_2, \lambda_3, \dots)$ decomposes informational frictions into distinct *risk dimensions*: mean (pricing discipline), volatility, asymmetry, and tails. In later sections we endogenize these distortions by linking them to risk attitudes, at which point the multipliers acquire a preference-based interpretation and κ becomes an equilibrium measure of belief disagreement.

At this stage, the SDF is best understood as a *change of measure*: it distorts the benchmark beliefs $\bar{\nu}$ just enough to satisfy the pricing constraint \hat{r}^B , without yet invoking preferences or equilibrium. Because beliefs are fixed, the kernel reflects an *informational adjustment*, not a utility-based valuation.

Despite these adjustments, all agents remain anchored to the same benchmark expected return constraint $\hat{r}^B = \sum_{s \in S} \eta_i(s) r_{t+1}^B(s)$. This condition ensures no-arbitrage across differing beliefs and is enforced through the exponential form of the exogenous kernel:

$$m^{\eta_i | \bar{\nu}}(s) = \frac{\exp(-\lambda_1 r_{t+1}^B(s))}{\sum_{u \in S} \bar{\nu}(u) \exp(-\lambda_1 r_{t+1}^B(u))}.$$

5.2 Subjective Valuation in an Objective Price System

A potential confusion arises at this stage: the construction of an entropy-based kernel may appear indistinguishable from standard risk-neutral pricing. It is therefore essential to highlight the distinction. In our framework, market prices are pinned down by a common benchmark distribution $\bar{\nu}$, which in this section we interpret as the *risk-neutral measure*. This yields an arbitrage-free pricing operator identical in form to risk-neutral valuation, and therefore $P(A)$ represents the unique, agreed price of any payoff. However, the novelty lies in the fact that *valuations are not exhausted by these common prices*. Agents distort $\bar{\nu}$ through informational distortions (and later, through informational preferences), forming subjective measures η_i that generate idiosyncratic certainty equivalents. This wedge—between objective prices and subjective values—is the engine of trade even in a frictionless and arbitrage-free environment.

Formally, let $\mathbf{A} = (A_1, \dots, A_n)$ be a vector of state-contingent payoffs. Let $\mathbf{R} = (R_1, \dots, R_n)$ denote the corresponding *gross return* vector, and let $\boldsymbol{\rho} = (\rho_1, \dots, \rho_n)$ denote the *log-return* vector with $\rho(s) = \ln R(s)$. All agents observe the same return structure $R(s)$ (equivalently, $\rho(s)$) and share access to the risk-neutral distribution $\bar{\nu}$, which serves as the common prior for pricing. From these primitives, we define the risk-neutral pricing operator as:

$$P(A) = \sum_{s \in S} \bar{\nu}(s) A(s) R^{-1}(s). \quad (12)$$

The expression $P(A)$ corresponds to the replication cost of A and is common across all agents. The associated state-price density is

$$\pi(s) = \bar{\nu}(s) \cdot R^{-1}(s),$$

which is positive and publicly observable. The market pricing functional $\Pi[A] := \sum_s \pi(s)A(s)$ is therefore objective: it satisfies no-arbitrage and is equally accepted

by all.

Yet while prices are common, valuations are not. Each agent i forms a personal belief η_i by optimally tilting the benchmark $\bar{\nu}$ according to the informational distortion criterion above under the common pricing anchor. In this exogenous section, the corresponding belief takes the exponential form (with tilt pinned down by the anchor on the benchmark claim):

$$\eta_i(s) = \frac{\bar{\nu}(s) \exp(-\lambda_i r_{t+1}^B(s))}{\sum_{u \in S} \bar{\nu}(u) \exp(-\lambda_i r_{t+1}^B(u))}, \quad (13)$$

where λ_i summarizes the degree of benchmark-relative belief distortion consistent with the common anchor. The associated information-based SDF component is:

$$m^{\eta_i|\bar{\nu}}(s) := \frac{\eta_i(s)}{\bar{\nu}(s)} = \frac{\exp(-\lambda_i r_{t+1}^B(s))}{\sum_{u \in S} \bar{\nu}(u) \exp(-\lambda_i r_{t+1}^B(u))}. \quad (14)$$

Using this kernel, the agent's subjective valuation of A , conditional on her belief η_i , can be re-expressed under the common risk-neutral measure $\bar{\nu}$ as:

$$\begin{aligned} V^{\eta_i|\bar{\nu}}(A) &= \sum_{s \in S} \eta_i(s) A(s) R^{-1}(s) \\ &= \sum_{s \in S} \bar{\nu}(s) m^{\eta_i|\bar{\nu}}(s) A(s) R^{-1}(s). \end{aligned} \quad (15)$$

Although the valuation operator is written as a linear expectation with state-dependent weights, this does not introduce concavity in payoffs; risk aversion enters only through the certainty-equivalent (CE) (or utility-based) mapping, not through state-dependent discounting.

In this regard, this valuation $V^{\eta_i|\bar{\nu}}(A)$ is directly linked to the CE:

$$u_i(\text{CE}_i(A)) = \sum_{s \in S} \eta_i(s) u_i(A(s) R^{-1}(s)), \quad (16)$$

so that $\text{CE}_i(A)$ represents the *personal fair payoff* that agent i considers equivalent to holding the risky asset A under her distorted belief η_i . In other words, the

informational valuation $V^{\eta_i|\bar{\nu}}(A)$ is driven by—and can be interpreted through—the certainty equivalent under η_i .

In general, $\text{CE}_i(A)$ does not coincide exactly with $V^{\eta_i|\bar{\nu}}(A)$. A second-order Arrow–Pratt expansion yields:

$$\text{CE}_i(A) \approx V^{\eta_i|\bar{\nu}}(A) - \frac{1}{2} \text{RA}_i(x^*) \text{Var}_{\eta_i}(AR^{-1}), \quad (17)$$

so the CE is approximated by the informational valuation minus a risk premium. If the agent is risk-neutral ($u_i(x) = x$), the curvature term vanishes and the equality is exact:

$$\text{CE}_i(A) = V^{\eta_i|\bar{\nu}}(A).$$

Moreover, if beliefs are homogeneous ($\eta_i = \bar{\nu}$), then

$$\text{CE}_i(A) = V^{\eta_i|\bar{\nu}}(A) = P(A),$$

so prices, valuations, and certainty equivalents coincide.

This motivates a natural measure of the valuation–pricing gap:

$$\Delta_i(A) := \text{CE}_i(A) - P(A),$$

which vanishes only under homogeneous beliefs but is otherwise strictly belief- and preference-dependent. It provides a precise metric of how belief dispersion and risk attitudes fuel trade. The sign of $\Delta_i(A)$ determines trading behavior:

$$\Delta_i(A) \begin{cases} > 0 & \Rightarrow \text{Asset underpriced} \Rightarrow \text{buy,} \\ = 0 & \Rightarrow \text{Fairly priced} \Rightarrow \text{no trade,} \\ < 0 & \Rightarrow \text{Asset overpriced} \Rightarrow \text{sell.} \end{cases}$$

The term

$$\frac{1}{2} \text{RA}_i(x^*) \text{Var}_{\eta_i}(AR^{-1})$$

in (17) is the familiar Arrow–Pratt risk premium. In our framework, it plays a dual role: first, as the adjustment that separates the certainty equivalent from the informational valuation $V^{\eta_i|\bar{\nu}}(A)$; second, as the bridge toward a preference-adjusted kernel. While the InSDF introduced so far is purely informational and exogenous, the Arrow–Pratt term reveals how risk attitudes will enter the picture. In Sections 4 and 5, we build directly on this insight: by endogenizing beliefs through entropy-weighted utility, we show how the stochastic discount factor acquires both an informational and a preference-driven component.

In sum, (12) defines a unique, objective, and arbitrage-free risk-neutral price for every payoff, accepted by all agents. Valuations, however, remain subjective and belief-dependent. This divergence—captured by the tilt parameter λ_i and the induced kernel $m^{\eta_i|\bar{\nu}}$ —is what fuels trade. Agreement on prices does not imply agreement on value.

5.3 Belief Distortion, No-Arbitrage, and the Valuation Gap

A central insight of our framework is that *no-arbitrage does not require belief agreement*. As long as agents are disciplined by a common benchmark anchor \hat{r}^B on the reference benchmark claim, they may hold heterogeneous beliefs without violating pricing consistency. What guarantees this is not agreement on state probabilities, but agreement on the pricing discipline embedded in the common benchmark.

To capture this heterogeneity, let us consider agents i and j , who face the same benchmark $\bar{\nu}$ but form distinct subjective beliefs η_i and η_j . Each belief arises from a separate entropy minimization problem with its own tilt parameter (λ_i and λ_j), so their valuations need not coincide.

The classical representative-agent framework corresponds to the knife-edge case

$$\eta_i = \eta_j = \bar{\nu},$$

so that all agents share the benchmark measure. In this scenario, disagreement vanishes:

$$\Delta_i(A) = \text{CE}_i(A) - P(A) = 0 \quad \text{for all } i,$$

and therefore all valuations coincide with the common price $P(A)$. By contrast, once we allow $\eta_i \neq \eta_j$, heterogeneity in informational beliefs generates market activity even under a common, arbitrage-free pricing system.

To see how disagreement operates, recall the anchor discipline:

$$\sum_{s \in \mathcal{S}} \eta_i(s) r_{t+1}^B(s) = \hat{r}^B \quad \iff \quad \sum_{s \in \mathcal{S}} \bar{\nu}(s) m^{\eta_i|\bar{\nu}}(s) r_{t+1}^B(s) = \hat{r}^B.$$

Here, $\bar{\nu}$ is the common benchmark, while $m^{\eta_i|\bar{\nu}}(s)$ encodes agent i 's informational tilt through λ_i . Crucially, the kernel adjusts beliefs so that every agent's distorted measure η_i reproduces the same benchmark anchor \hat{r}^B . In this sense, disagreement enters through the *shape* of the likelihood-ratio kernel, not the existence of a common arbitrage-free pricing operator.

The divergence from the benchmark valuation is captured by the *benchmark valuation gap*. Let the subjective valuation of A under belief η_i be

$$V^{\eta_i|\bar{\nu}}(A) = \sum_{s \in \mathcal{S}} \bar{\nu}(s) m^{\eta_i|\bar{\nu}}(s) A(s) R^{-1}(s),$$

and recall that the objective price is $P(A) = \sum_{s \in \mathcal{S}} \bar{\nu}(s) A(s) R^{-1}(s)$. We define the benchmark valuation gap as

$$\Delta_{i|\bar{\nu}}^V(A) := V^{\eta_i|\bar{\nu}}(A) - P(A).$$

This valuation wedge can be written as a disagreement premium relative to the benchmark price; its decomposition and covariance form are reported in Appendix A.3.

This quantity isolates the *belief-driven* component of disagreement relative to

the common pricing benchmark.

For trading, however, the economically relevant object is the agent's *certainty equivalent*. Recall

$$u_i(\text{CE}_i(A)) = \sum_{s \in S} \eta_i(s) u_i(A(s) R^{-1}(s)), \quad \Delta_i(A) := \text{CE}_i(A) - P(A),$$

which is the primitive driver of trade: it compares the common price against the agent's personal, preference-adjusted evaluation of the payoff.

Risk-neutral special case. If agent i is risk-neutral, $u_i(x) = x$, then the certainty equivalent coincides with the informational valuation,

$$\text{CE}_i(A) = \sum_{s \in S} \eta_i(s) A(s) R^{-1}(s) = V^{\eta_i | \bar{\nu}}(A),$$

and therefore

$$\Delta_i(A) = \Delta_{i|\bar{\nu}}^V(A).$$

In general (under curvature), $\text{CE}_i(A)$ differs from $V^{\eta_i | \bar{\nu}}(A)$ by a risk adjustment, but the sign and dispersion of $\Delta_i(A)$ remain the operative channel through which disagreement generates trading activity.

While the benchmark valuation gap highlights the gap between an individual agent and the common reference, trade stems from disagreement *across agents*. For two agents i and j , define the pairwise disagreement wedge by

$$\Delta_{ij}(A) := \text{CE}_i(A) - \text{CE}_j(A).$$

From certainty-equivalent wedges to trading volume

The previous definitions make it possible to connect disagreement to a concrete, observable implication: trading volume. Because prices are common while certainty equivalents are heterogeneous, the primitive driver of trade is the cross-sectional

dispersion of the certainty-equivalent wedges

$$\Delta_i(A) = \text{CE}_i(A) - P(A),$$

and equivalently the pairwise wedges

$$\Delta_{ij}(A) = \text{CE}_i(A) - \text{CE}_j(A).$$

To make the mapping explicit, consider a reduced-form demand rule in which agent i takes a position proportional to her wedge:

$$x_i(A) = \tau_i \Delta_i(A), \tag{18}$$

where $\tau_i > 0$ is a (possibly agent-specific) scale parameter capturing risk tolerance, trading intensity, or balance-sheet capacity. This rule is deliberately minimal: it only requires that demand increases when the agent views the asset as underpriced relative to her personal certainty equivalent.

Define trading volume in the asset as the aggregate absolute reallocation:

$$\mathcal{V}(A) := \frac{1}{2} \sum_i |x_i(A)| = \frac{1}{2} \sum_i \tau_i |\Delta_i(A)|. \tag{19}$$

Hence, even in a frictionless, arbitrage-free market with a unique objective price $P(A)$, volume is strictly positive whenever the cross-sectional distribution of certainty-equivalent wedges is non-degenerate. In particular, if $\Delta_i(A)$ is dispersed across agents, then $\mathcal{V}(A)$ increases mechanically.

Equivalently, one may express volume in terms of pairwise disagreement: for a population of agents, the average absolute pairwise wedge

$$\bar{\Delta}(A) := \mathbb{E}_{i,j} [|\Delta_{ij}(A)|] = \mathbb{E}_{i,j} [|\text{CE}_i(A) - \text{CE}_j(A)|]$$

is a direct measure of the dispersion that sustains trade. Through (18)–(19), higher

dispersion in certainty equivalents maps into higher trading volume.

Finally, because certainty equivalents are computed under subjective beliefs η_i (and later under entropy-weighted preferences), the same mapping links volume to belief dispersion. In the present exogenous setup, dispersion in $\{\eta_i\}$ is summarized by dispersion in the tilt parameters $\{\lambda_i\}$, while in the endogenous model it will be pinned down jointly by risk preferences and informational frictions. In both cases, the wedge $\Delta_i(A)$ provides the operational bridge from disagreement to market activity.

From certainty equivalents to the SDF (marginal valuation). Define the indifference price $\pi_{i,t}(\varepsilon; X)$ for a small position εX_{t+1} by

$$u_i(C_{i,t}) = \mathbb{E}_t^{\eta_i} \left[\delta u_i \left(C_{i,t+1} + \varepsilon X_{t+1} - R_t^f \pi_{i,t}(\varepsilon; X) \right) \right].$$

Differentiating at $\varepsilon = 0$ yields the marginal price

$$\pi'_{i,t}(0; X) = \mathbb{E}_t^{\eta_i} [m_{t+1}^i X_{t+1}], \quad m_{t+1}^i(s) := \delta \frac{u'_i(C_{i,t+1}(s))}{u'_i(C_{i,t})} \frac{1}{R_t^f}.$$

Hence the SDF is the *marginal* certainty-equivalent price. Under a benchmark operator $(\bar{\nu}, \hat{m}_{t+1})$, the same valuation can be written as

$$\pi'_{i,t}(0; X) = \mathbb{E}_t^{\bar{\nu}} [\hat{m}_{t+1}(s) m^{n_i|\bar{\nu}}(s) X_{t+1}(s)], \quad m^{n_i|\bar{\nu}}(s) := \frac{d\eta_i}{d\bar{\nu}}(s).$$

6 The Value of Information in the Construction of Preferences

This section studies how subjective beliefs and preferences interact once information is treated as an object with economic value. The guiding question is: *how do risk preferences shape the way information, defined as $I(p) := -\log p$, enter utility?* Under complete information, beliefs can be treated as separable from preferences as a basic rationality benchmark. Under uncertainty, however, probabilities are not merely inputs: they may be *selected* or *distorted* in a preference-consistent way, so that an agent’s risk attitude determines whether informational surprise is valued, penalized, or effectively ignored. This motivates entropy-weighted utility specifications, which embed informational content into preferences while maintaining a clear conceptual distinction between *beliefs* (probability assignments) and *preferences* (ranking of outcomes).

The discussion builds on several related strands in the literature on utility, information, and subjectivity. Some approaches attempt to define entropy directly over utility functions by normalizing state utility values and interpreting them probabilistically (Abbas, 2006; Harańczyk et al., 2008; Dionisio et al., 2008). While mathematically convenient, this strategy is conceptually problematic: utility is not inherently probabilistic, and normalization makes the resulting “entropy of utility” depend on the absolute scale of utilities and fail to respect the affine-invariance of expected-utility theory. In effect, these treatments blur the distinction between beliefs and preferences by forcing a structural analogy between them. The approach developed here avoids this conflation: entropy is used to characterize informational properties of *belief vectors*, not to recast preferences as probability-like objects.

A second and influential line of work introduces entropy as an informational *constraint*. In particular, Rouge and El Karoui (2000) model agents who maximize expected utility under subjective beliefs subject to an information cost measured by relative entropy from a reference (typically objective) measure, implying an optimal exponential tilt linked to marginal utility. This thesis shares the core idea

that entropy can discipline belief formation and that a unique “true” probability need not be imposed. The key difference is conceptual and structural: rather than treating relative entropy as an external constraint around an objective reference, the present framework derives belief dispersion endogenously from risk attitudes and embeds the informational term as a primitive component of the preference functional, thereby enabling a fully subjective account of belief heterogeneity.

The interpretation of entropy as a measure of belief dispersion is also consistent with the information-value results of Cabrales et al. (2013), who show that more informative signals (in the sense of Blackwell dominance) reduce posterior entropy and increase expected utility. While their analysis is static, it supports the use of entropy as a natural statistic of informational structure. Here, the same intuition is extended to a market-based environment: belief dispersion interacts with heterogeneous risk preferences and disagreement to affect valuation and, ultimately, state prices.

Finally, Friedman and Sandow (2016) provide a rationale for boundedly rational choice rules by deriving behavior from entropy maximization under utility constraints. Their view of entropy as a rationality principle resonates with the present emphasis on entropy as an organizing structure of subjectivity. The mechanism differs, however: their entropy is over *actions*, whereas the framework developed here applies entropy to *beliefs* and makes explicit how risk aversion shapes the endogenous probability assignments that enter asset valuation.

6.1 Information-Based Utility Functions

The connection between informational content and subjective preferences can be formally established through the concept of *weighted entropy*, an idea first developed by Belis and Guiasu (Guiasu, 1971). They proposed a generalization of Shannon’s entropy by incorporating weights assigned to different events or symbols, reflecting their relevance or qualitative attributes. Weighted entropy modifies the classical Shannon entropy formula by multiplying each term in the sum by a weight

factor:

$$H_\varphi^w = - \sum_s \varphi(x_s) p(x_s) \log p(x_s). \quad (20)$$

Here, $\varphi(x_s)$ is a non-negative weighting function, in our case assigning a utility value to outcomes x_s . Unlike classical Shannon entropy, this formulation allows entropy to reflect not only the dispersion of beliefs but also the subjective importance of each possible outcome.

This connects with the idea of a subjective probability weighting function that describes distortions in the perception of subjective probability. In decision-making under risk, agents often over- or under-weight objective probabilities based on the perceived importance of outcomes (Akrenius, 2020; Yang and Qiu, 2014; Cavagnaro et al., 2013).

In our framework, the weighting function is naturally provided by preferences themselves. Let

$$U : \mathbb{R}_+ \times \Gamma \rightarrow \mathbb{R}, \quad (w, \gamma) \mapsto U(w, \gamma),$$

denote a baseline (von Neumann–Morgenstern) utility index, where w_s is the state- s economic outcome (wealth, consumption, or any payoff-relevant primitive) and $\gamma \in \Gamma \subset \mathbb{R}$ is the risk-preference parameter (e.g. a CRRA curvature coefficient). Throughout, $U(\cdot, \gamma)$ is assumed to satisfy the usual regularity conditions used in expected-utility asset pricing: it is strictly increasing in w , continuous, and twice continuously differentiable on \mathbb{R}_{++} , with curvature governed by γ (typically $U_{ww}(w, \gamma) < 0$ for $\gamma > 0$, with the risk-neutral case recovered at $\gamma = 0$). Given a subjective belief vector $p = (p_s)_{s \in S} \in \Delta(S)$, we treat utility as the weight applied to probabilities and define the *utility-weighted entropy* statistic

$$H_\gamma^U(p) := - \sum_s p_s U(w_s, \gamma) \log p_s, \quad p \in \Delta(S). \quad (21)$$

Here p_s is the subjective probability assigned to state s . The object $H_\gamma^U(p)$ is not

a primitive taste for uncertainty; it is a summary statistic of the *belief distribution* p , weighted by the agent's state utility profile $\{U(w_s, \gamma)\}_{s \in S}$.

From utility-weighted entropy to an entropy-augmented criterion. The utility-weighted entropy statistic $H_\gamma^U(p)$ becomes economically operative once it is embedded in a criterion functional defined over beliefs. We therefore consider the entropy-augmented objective

$$J_\gamma(p) := \sum_s p_s U(w_s, \gamma) - \lambda(\gamma) H_\gamma^U(p), \quad p \in \Delta(S), \quad \lambda(\gamma) = \alpha \operatorname{sign}(\gamma). \quad (22)$$

which is derived in Appendix A.4 (“Derivation of Entropy-Augmented Utility”). The parameter $\alpha > 0$ is the *information-discipline scale*: it governs the strength of the entropy-based discipline applied to belief dispersion, and corresponds to the informational penalty/intensity introduced in the KL-based motivation of the appendix. Risk preferences enter through the sign of the entropy term. In particular, we parameterize the regime dependence by restricting

$$\lambda(\gamma) \in \{-\alpha, 0, +\alpha\},$$

so that $\lambda(\gamma)$ is not an independent free parameter but is pinned down by α up to the sign implied by the risk regime. Equivalently, writing $\lambda(\gamma) = \alpha \cdot \operatorname{sign}(\gamma)$ yields

$$\lambda(\gamma) = \begin{cases} -\alpha, & \gamma < 0, \\ 0, & \gamma = 0, \\ +\alpha, & \gamma > 0, \end{cases} \quad \alpha > 0.$$

Substituting the definition of $H_\gamma^U(p)$ into $J_\gamma(p)$ gives the convenient statewise form

$$J_\gamma(p) = \sum_{s \in S} p_s U(w_s, \gamma) + \lambda(\gamma) \sum_{s \in S} p_s U(w_s, \gamma) \log p_s = \sum_{s \in S} p_s U(w_s, \gamma) \left(1 + \lambda(\gamma) \log p_s\right).$$

Hence the entropy-augmented representation is not a separate preference argument but an equivalent way of writing a utility criterion in which beliefs are disciplined through the utility-weighted entropy term, with α setting the overall strength of discipline and $\lambda(\gamma)$ encoding how the sign of that discipline depends on the risk regime.

Information-based (statewise) representation. The entropy-augmented criterion can be written in statewise form as

$$J_\gamma(p) = \sum_{s \in S} p_s U(w_s, \gamma) \left(1 + \lambda(\gamma) \log p_s\right), \quad p \in \Delta(S), \quad (23)$$

so that the statewise contribution to the criterion is

$$p_s U(w_s, \gamma) \left(1 + \lambda(\gamma) \log p_s\right), \quad p_s \in (0, 1].$$

This distinction is important. The probability term does not enter as a primitive argument of utility itself; rather, it enters through the *belief criterion* $J_\gamma(p)$, which evaluates alternative subjective belief distributions by combining baseline utility with an entropy-based informational discipline. Here w_s denotes the economic outcome in state s , γ is the coefficient of risk aversion, and p_s is the subjective probability assigned to that state. The curvature of $U(w_s, \gamma)$ —as measured by the Arrow–Pratt index—determines how sensitive the agent is to payoff risk, while $\lambda(\gamma)$ determines how the entropy term is signed across risk regimes. Thus, information affects the criterion through the geometry of beliefs, not by turning entropy into a primitive taste variable. This anticipates the endogenous link between risk aversion and belief formation developed in the next subsection.

Three canonical regimes. Because $\alpha > 0$ is fixed and only the sign of $\lambda(\gamma)$ changes with the risk regime, (23) yields three canonical cases:

- **Risk-loving** ($\gamma < 0$, so $\lambda(\gamma) = -\alpha < 0$): for each state, the statewise

contribution becomes

$$p_s U(w_s, \gamma) (1 - \alpha \log p_s).$$

Surprise is desirable and so increases the statewise contribution associated with low-probability, high-reward states.

For each state, the relevant multiplicative information term is

$$1 - \alpha \log p, \quad 0 < p \leq 1.$$

Since $1 - \alpha \log p > 0$, the sign of the contribution remains determined by the baseline utility $u(x)$. Hence convexity in payoffs is still governed by the curvature of u ; the informational term simply rescales that curvature by a positive factor. But convexity by itself is a necessary, not sufficient, condition. The reason is that $1 - \alpha \log p$ is always positive and increasing in surprise, whereas the alternative candidate $1 + \alpha \log p$ turns negative once probability drops below $e^{-1/\alpha}$:

$$1 - \alpha \log p > 0 \quad \forall p \in (0, 1], \quad 1 + \alpha \log p \begin{cases} > 0 & p > e^{-1/\alpha}, \\ < 0 & p < e^{-1/\alpha}. \end{cases}$$

This sign change would reverse the direction of valuation in rare states—precisely where thrill-seeking agents are meant to value surprise, not despise it. Therefore, $1 + \alpha \log p$ is incompatible with coherent risk-seeking behavior.

The informational term therefore *magnifies* the emotional impact of rare events without turning bad outcomes into good ones. For risk-loving agents, rare disasters remain painful—indeed more so—but the amplified contribution of rare windfalls can more than compensate. Risk-loving behavior therefore stems from a greater *trade-off* between magnified upside and magnified downside, not from a taste for catastrophe. That is what precisely defines a

risk-loving attitude.

- **Risk-neutral** ($\gamma = 0$, so $\lambda(\gamma) = 0$):

$$J_0(p) = \sum_{s \in S} p_s U(w_s, 0).$$

Beliefs do not alter the criterion beyond their standard role as probability weights. The agent cares only about expected payoffs.

- **Risk-averse** ($\gamma > 0$, so $\lambda(\gamma) = +\alpha > 0$): for each state, the statewise contribution becomes

$$p_s U(w_s, \gamma) \left(1 + \alpha \log p_s\right).$$

Surprise is penalized and so lowers the contribution of low-probability states to the criterion; agents therefore favor more dispersed, robust belief structures.

Entropy is not a taste, but an epistemological stance. A potential confusion is to interpret the appearance of entropy in (21)–(22) as if agents had a direct preference for “more” or “less” uncertainty. That is *not* the interpretation in this thesis. Entropy is not a primitive argument of felicity like wealth or consumption; it is a summary statistic of the *belief distribution* the agent entertains (and, in the next subsection, chooses). In this sense, entropy measures an epistemological stance: high-entropy beliefs correspond to a skeptical posture in which the agent refrains from imposing sharp probabilistic structure (“no patterns”), while low-entropy beliefs correspond to a more concentrated posture in which the agent assigns large mass to a subset of states (“patterns exist”). Risk preferences matter because they reshape the state utility profile and therefore the incentives to concentrate or disperse subjective probabilities. The comparative-statics results below should therefore be read as statements about *optimal belief geometry* and perceived informational structure, not as evidence that agents intrinsically “like uncertainty”.

Entropy is maximized under maximal dispersion, and the informational term vanishes at certainty (degenerate beliefs) $(p_s) = (0, 1)$ or $(1, 0)$.²

We can now express the entropy-augmented criterion directly in expectation form, without introducing probability as a primitive argument of utility. Starting from the utility-weighted entropy statistic

$$H_\gamma^U(p) := - \sum_i p_i U(w_i, \gamma) \log p_i,$$

the belief criterion is

$$J_\gamma(p) = \mathbb{E}_p[U(w, \gamma)] - \lambda(\gamma) H_\gamma^U(p), \quad \lambda(\gamma) = \alpha \operatorname{sign}(\gamma) \in \{-\alpha, 0, +\alpha\}, \quad \alpha > 0. \quad (24)$$

Equivalently,

$$\begin{aligned} J_\gamma(p) &= \sum_i p_i U(w_i, \gamma) + \lambda(\gamma) \sum_i p_i U(w_i, \gamma) \log p_i \\ &= \sum_i p_i U(w_i, \gamma) \left(1 + \lambda(\gamma) \log p_i\right). \end{aligned} \quad (25)$$

This decomposition shows that entropy does not enter as a separate taste component, but through a belief-dependent adjustment to the baseline expected utility. Because the probability vector p is subjective, the entropy term is itself endogenously shaped by the agent's beliefs. Entropy is therefore no longer an objective measure of dispersion, but a preference-disciplined statistic of the way the agent organizes uncertainty when markets are incomplete and the true model is not known.

In this context, entropy is not the result of passively measuring the amplitude of market patterns, but rather a reflection of the agent's subjective belief in the existence of such patterns. This *subjective entropy* captures the extent to which the individual sees structure—or randomness—in the world, and how much weight is assigned to rare versus likely outcomes.

²Using L'Hôpital's rule, it can be shown that $0 \log 0 = 0$.

Only the second term in (24) deviates from standard von Neumann expected utility and captures how the subjective probability structure modifies valuation. We refer to this component as the *incremental entropy adjustment*: the informational contribution to the belief criterion induced by the geometry of beliefs. For risk-averse agents ($\gamma > 0$), $\lambda(\gamma) = +\alpha$, so the entropy term acts as a penalty: more diffuse beliefs (higher entropy) reduce the criterion. This aligns with the psychological idea that unlikely but favorable outcomes are not fully trusted, so surprise lowers valuation even when payoffs themselves remain unchanged. In this framework, entropy represents the criterion cost of epistemic uncertainty.

More precisely, in the risk-averse case,

$$\alpha H_\gamma^U(p) = \mathbb{E}_p[U(w, \gamma)] - J_\gamma(p).$$

Thus the entropy term measures the gap between baseline expected utility and the entropy-disciplined belief criterion.

Conversely, for risk-seeking individuals ($\gamma < 0$), $\lambda(\gamma) = -\alpha$, so the entropy term becomes a bonus: belief structures that place more weight on rare, high-payoff states increase the criterion. In that case, surprise is valuable, and the agent is willing to sustain more concentrated, pattern-seeking beliefs.

All three risk regimes can therefore be summarized in a single expression:

$$J_\gamma(p) = \mathbb{E}_p[U(w, \gamma)] - \lambda(\gamma) H_\gamma^U(p), \quad \lambda(\gamma) = \alpha \operatorname{sign}(\gamma).$$

Here $\mathbb{E}_p[U(w, \gamma)]$ is the baseline expected utility, $H_\gamma^U(p)$ is the utility-weighted entropy, and $\lambda(\gamma)$ governs how risk preferences transform entropy into either a penalty ($\gamma > 0$), a neutral adjustment ($\gamma = 0$), or a bonus ($\gamma < 0$).

In all cases, since the belief structure p is under the agent's control, the entropy term becomes a reflection of subjective epistemic posture. Beliefs are no longer passive inputs—they are molded by preferences, and in turn they reshape the criterion used to evaluate uncertainty.

It follows that higher belief concentration (lower entropy) is less penalized for risk-averse individuals, while flatter beliefs (higher entropy) increase the penalty and push the criterion further below its conventional expected-utility counterpart. Entropy thus acts as an endogenous informational cost, reflecting the agent’s discomfort with epistemic dispersion rather than risk alone.

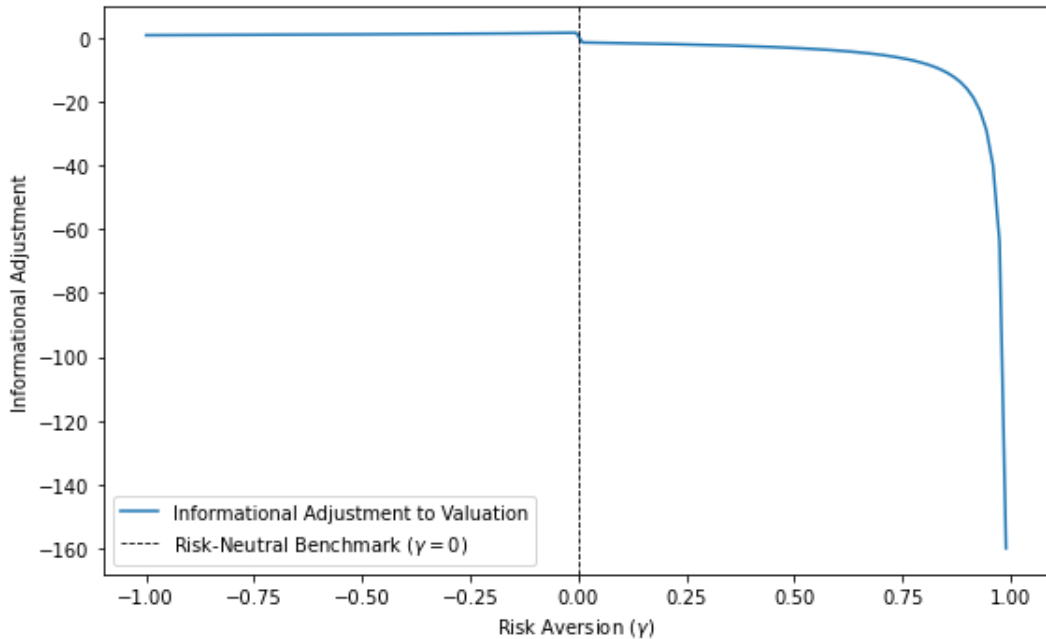


Figure 1: Illustrative role of utility-weighted entropy in valuation across risk-aversion regimes: holding the payoff structure fixed, the figure shows how the belief distribution, summarized by entropy, modifies the valuation criterion under different values of γ . The sign change at risk neutrality separates regimes in which more diffuse beliefs lower valuation from those in which they raise it.

The plot should therefore be read as a representation of how different *belief structures* affect the entropy-augmented criterion across risk regimes, not as evidence that entropy itself is a primitive taste variable. It shows that the entropy adjustment contributes negatively to the criterion for risk-averse agents ($\gamma > 0$) and positively for risk-loving agents ($\gamma < 0$), with a nonlinear, asymmetric profile. This pattern captures not only conventional risk attitudes but a deeper form of *surprise aversion*: for $\gamma > 0$, higher entropy (more diffuse beliefs, more potential for surprise) is costly; when beliefs are more concentrated (low entropy), the penalty shrinks. Conversely, risk-loving agents benefit more when beliefs are diffuse in informational

terms because rare, high-payoff states receive greater effective weight; as beliefs concentrate, the bonus fades. Thus, entropy shapes the criterion through *belief structures* (how probabilities are distributed) rather than through outcomes alone.

This asymmetry suggests that entropy operates as an *informational analog to loss aversion*. In prospect theory, losses relative to a reference point are weighted more heavily than equivalent gains, creating a kinked value function. In our framework, it is not outcomes but rather *beliefs* that generate the asymmetry. For risk-averse agents ($\gamma > 0$), diffuse beliefs (higher entropy) act like informational “losses,” reducing the criterion even when payoffs are unchanged. For risk-loving agents ($\gamma < 0$), the same entropy term is interpreted as an informational “gain,” rewarding exposure to surprising, low-probability events. The chart above captures this duality: entropy penalizes risk-averse regimes but enhances risk-loving ones, producing a behavioral kink in the utility–information trade-off. In this sense, *surprise aversion* plays the same structural role as loss aversion, but in the informational domain.

The departure from standard expected utility is captured by the utility-weighted entropy term:

$$\begin{aligned} H_\gamma^U(p) &= - \sum_i p_i U(w_i, \gamma) \log p_i \\ &= - \mathbb{E}_p[U(w, \gamma) \log p] \\ &= \mathbb{E}_p[U(w, \gamma)] \cdot H(p) - \text{Cov}_p(U(w, \gamma), \log p). \end{aligned}$$

This decomposition shows that subjective (utility-weighted) entropy depends not only on average dispersion, $H(p)$, but also on the alignment between state utilities and their assigned probabilities. A positive covariance means high-utility states are perceived as relatively unlikely (large negative $\log p$), amplifying the entropy penalty; a negative covariance indicates favorable states are expected, damping the informational cost.

Notation bridge. When working with a single generic agent we write $p \in \Delta(S)$ for the belief vector. In the multi-agent notation, this corresponds to setting $p = \eta_i$. Likewise, the state utility weight $U(w_s, \gamma)$ corresponds to $U_i(s)$, the information-discipline parameter α corresponds to α_i , and the simplex multiplier θ corresponds to θ_i . For convenience we also use the composite tilt parameter

$$\tau_i := \frac{\theta_i}{\alpha_i},$$

so that only the ratio θ_i/α_i governs the strength of belief concentration.

6.2 Risk Preferences and the Information Structure of Beliefs

To understand how beliefs and preferences interact under uncertainty, we now examine how risk aversion shapes the entropy-augmented belief criterion. We model baseline preferences using an isoelastic utility function:

$$U(w, \gamma) = \frac{w^{1-\gamma} - 1}{1 - \gamma},$$

with the usual log limit as $\gamma \rightarrow 1$.

To keep notation explicit, we specialize to the *risk-averse* regime ($\gamma > 0$), for which $\lambda(\gamma) = +\alpha$ with $\alpha > 0$. In that case, the statewise contribution to the criterion is

$$c_\gamma(w, p) := pU(w, \gamma)(1 + \alpha \log p), \quad p \in (0, 1]. \quad (26)$$

Equivalently, for a full belief vector $p \in \Delta(S)$, the corresponding criterion is

$$J_\gamma(p) = \sum_{s \in S} p_s U(w_s, \gamma)(1 + \alpha \log p_s).$$

Here $p \in (0, 1]$ denotes the subjective probability assigned to the state associated with outcome w , so $1 + \alpha \log p$ is the risk-averse information weight applied state-by-

state inside the belief criterion. The baseline utility $U(w, \gamma)$ continues to describe preferences over outcomes only; the probability term enters through the geometry of the belief distribution and therefore should be interpreted as an epistemic, not hedonic, component of the criterion.

The factor $1 + \alpha \log p$ captures informational surprise under the risk-averse regime. Differentiating the statewise contribution with respect to γ ,

$$\frac{\partial c_\gamma(w, p)}{\partial \gamma} = p \frac{\partial U}{\partial \gamma} (1 + \alpha \log p).$$

Since $\partial U / \partial \gamma \leq 0$ under CRRA (for $w \neq 1$), the sign depends on $(1 + \alpha \log p)$, which changes at $p = e^{-1/\alpha}$. This state-dependence motivates the following characterization.

Lemma 1 (sign of the information-adjusted marginal effect of risk aversion). Fix $w > 0$, $p \in (0, 1)$, and CRRA utility $U(w, \gamma) = \frac{w^{1-\gamma}-1}{1-\gamma}$ (with the usual log limit at $\gamma = 1$). For the statewise contribution

$$c_\gamma(w, p) = p U(w, \gamma) (1 + \alpha \log p),$$

we have

$$\frac{\partial c_\gamma(w, p)}{\partial \gamma} = p \frac{\partial U}{\partial \gamma} (1 + \alpha \log p), \quad \frac{\partial U}{\partial \gamma} \leq 0,$$

with equality only if $w = 1$ (or in the continuous extension at $\gamma = 1$). Hence

$$\text{sign}\left(\frac{\partial c_\gamma(w, p)}{\partial \gamma}\right) = -\text{sign}(1 + \alpha \log p) = \begin{cases} < 0, & p > e^{-1/\alpha}, \\ = 0, & p = e^{-1/\alpha}, \\ > 0, & p < e^{-1/\alpha}. \end{cases}$$

Proof. Write $a := 1 - \gamma$ and $U(w, \gamma) = \frac{w^a - 1}{a}$. Differentiating with respect to γ gives

$$\frac{\partial U}{\partial \gamma} = -\frac{\partial}{\partial a} \left(\frac{w^a - 1}{a} \right) = \frac{(w^a - 1) - a w^a \ln w}{a^2} = \frac{h(a)}{a^2},$$

where $h(a) := (w^a - 1) - a w^a \ln w$. Then $h(0) = 0$ and

$$h'(a) = -a w^a (\ln w)^2 \leq 0,$$

so $h(a) \leq 0$ for all a , with equality only at $a = 0$. Since $a^2 > 0$, $\partial U / \partial \gamma \leq 0$.

Finally,

$$\frac{\partial c_\gamma(w, p)}{\partial \gamma} = p \frac{\partial U}{\partial \gamma} (1 + \alpha \log p),$$

and since $p > 0$, its sign is the opposite of $\text{sign}(1 + \alpha \log p)$.

Corollary (state vs. aggregate effect). Statewise, increasing γ lowers the contribution $c_\gamma(w, p)$ when $p > e^{-1/\alpha}$ and raises it when $p < e^{-1/\alpha}$. At the aggregate level, the net effect depends on the full belief distribution through the entropy term entering $J_\gamma(p)$.

Knife-edge cases in γ . When $\gamma = 0$ (risk-neutral), the entropy adjustment vanishes because $\lambda(\gamma) = 0$, so the criterion reduces to standard expected utility. As $\gamma \rightarrow 1$, the continuous log limit preserves the sign pattern in Lemma 1.

The following plot shows how the entropy adjustment to the criterion varies with risk aversion γ for different subjective probabilities p , using the wealth rescaling $w = 100e^\gamma$.

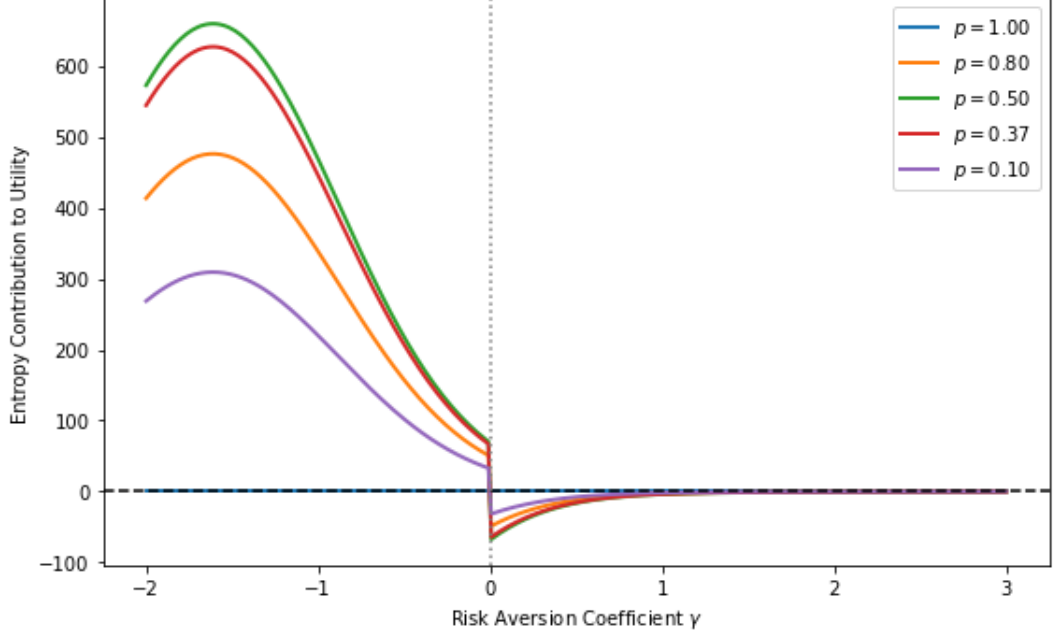


Figure 2: Illustrative informational adjustment across risk-aversion regimes: for a given subjective probability assignment, the figure shows how different belief configurations—summarized by utility-weighted entropy—affect valuation under different values of γ . Entropy is not interpreted as a primitive taste for uncertainty, but as a statistic of the agent’s subjective belief distribution.

This scaling is used to prevent CRRA utility from collapsing at high γ and to isolate the informational effect across regimes. The plot should not be read as showing that entropy is a primitive taste. Rather, it shows how different *belief structures* alter the entropy-augmented criterion across risk regimes. In particular, the entropy term acts as a bonus for $\gamma < 0$ and as a penalty for $\gamma > 0$, with a sign change at $\gamma = 0$.

The relative penalty is given by

$$\frac{H_{\gamma}^U(p)}{U(w, \gamma)}.$$




Figure 3: Illustrative relative entropy-utility contribution for a constant level of wealth, w : the ratio of utility-weighted entropy to baseline utility summarizes the proportional effect of belief dispersion on valuation, showing how the informational adjustment varies across risk-aversion regimes.

6.3 Optimal Endogenous Belief Distributions

We now derive the optimal belief distribution chosen by the agent when beliefs are decision variables; see Appendix A.5 for complete derivation.

Consider the optimization program (using $U(w_s, \gamma)$ as the state utility weight).

Let

$$\Delta := \left\{ p \in \mathbb{R}^{|S|} : \sum_{s \in S} p_s = 1, \quad p_s \geq 0 \quad \forall s \in S \right\}$$

denote the probability simplex. The agent solves

$$\max_{p \in \Delta} \sum_{s \in S} p_s U(w_s, \gamma) \left(1 + \alpha \log p_s \right), \quad \alpha > 0.$$

The Lagrangian is

$$\mathcal{L}(p, \theta) = \sum_{s \in S} p_s U(w_s, \gamma) \left(1 + \alpha \log p_s \right) - \theta \left(\sum_{s \in S} p_s - 1 \right),$$

where θ is the Lagrange multiplier associated with the probability normalization

constraint $\sum_{s \in S} p_s = 1$.

The first-order condition for each state $s \in S$ yields

$$U(w_s, \gamma) \left(2 + \alpha \log p_s \right) = \theta.$$

Solving for p_s gives the exponential-inverse closed form:

$$\boxed{p_s^* = \frac{1}{Z(\tau)} \exp\left(\frac{\tau}{U(w_s, \gamma)}\right)}. \quad (27)$$

$$Z(\tau) := \sum_u \exp\left(\frac{\tau}{U(w_u, \gamma)}\right), \quad \tau := \frac{\theta}{\alpha},$$

where the state-invariant constants are absorbed into the normalizer (Appendix A.5).

Normalization and the tilt parameter. The constant $Z(\tau)$ is the normalizer (partition function) that enforces $\sum_{s \in S} p_s^* = 1$:

$$Z(\tau) := \sum_{u \in S} \exp\left(\frac{\tau}{U(w_u, \gamma)}\right).$$

We also use the composite tilt parameter

$$\tau := \frac{\theta}{\alpha},$$

so that only the ratio θ/α governs belief concentration: θ is the simplex multiplier, while α scales the information-discipline term and therefore determines how costly it is (in utility units) to concentrate probability mass. As shown in Appendix A.5, state-invariant constants from the FOC are absorbed into $Z(\tau)$.

This exponential-inverse structure has a natural interpretation: probabilities are tilted inversely with state utility. The sign and magnitude of τ capture the direction and intensity of the distortion.

Interpretation: p^* as an endogenous epistemic stance. The distribution p^* is not an exogenous “belief input” and it is not the object of a direct taste for uncertainty. Rather, it is the *utility-maximizing subjective probability assignment* implied by the agent’s preferences under the information discipline encoded by α . In that sense, p^* summarizes the agent’s *epistemic stance* toward the state space: a more concentrated p^* corresponds to a posture in which the agent assigns substantial probability mass to a subset of states (a sharper view about which outcomes are likely, i.e. a stronger perceived structure or “patterns”), whereas a flatter p^* corresponds to a more skeptical posture in which the agent refrains from imposing sharp structure and keeps probability mass more evenly spread. The comparative statics of p^* (through $U(\cdot, \gamma)$ and the effective tilt τ) should therefore be read as statements about *optimal belief geometry* and perceived informational structure, not as evidence that the agent intrinsically “likes” entropy or uncertainty.

6.4 Optimal Beliefs and Subjective Entropy: A Numerical Illustration

This subsection illustrates how preferences shape optimal subjective probabilities and how these distortions give rise to what we call *subjective entropy*. Moving from the closed-form solution in (27) to a concrete simulation makes clear how risk attitudes translate into belief structures that differ systematically across agents. These beliefs are not derived from objective frequencies or statistical estimation, but are endogenously generated from the agent’s criterion under informational discipline. The resulting probability distributions therefore reflect how risk aversion or risk-seeking behavior alters informational judgment, producing either concentrated or dispersed beliefs. Crucially, the entropy of the optimal distribution serves as a measure of the agent’s epistemic posture under uncertainty, and therefore provides the bridge to the endogenous stochastic discount factor developed in the next section.

We consider a stylized payoff vector representing ten possible states of the

world:

$$w_s = \{8, 15, 25, 40, 60, 85, 120, 180, 260, 400\}.$$

This distribution is moderately right-skewed, with mean $\mu \approx 119.3$, standard deviation $\sigma \approx 120.6$, and skewness approximately 1.20. While asset returns are often left-skewed, terminal wealth outcomes are typically right-skewed, since payoffs are bounded below but unbounded above. This structure combines dense intermediate states—resembling volatility clustering—with a dispersed upper tail that captures rare but extreme payoffs. It therefore provides a realistic setting for examining how different levels of risk aversion map into subjective probabilities.

Table 1 reports the utility values across states for several levels of the risk-aversion coefficient γ , under the isoelastic utility function

$$U(w_s, \gamma) = \begin{cases} \log w_s, & \gamma = 1, \\ \frac{w_s^{1-\gamma} - 1}{1 - \gamma}, & \text{otherwise.} \end{cases}$$

Table 1: Isoelastic Utility Values Across Risk-Aversion Levels and Payoff States

| Payoff Levels w_s | | | | | | | | | | |
|---------------------|-------|--------|--------|--------|---------|---------|---------|----------|----------|----------|
| γ | 8 | 15 | 25 | 40 | 60 | 85 | 120 | 180 | 260 | 400 |
| -1 | 31.50 | 112.00 | 312.00 | 799.50 | 1799.50 | 3612.00 | 7199.50 | 16199.50 | 33799.50 | 79999.50 |
| 0 | 7.00 | 14.00 | 24.00 | 39.00 | 59.00 | 84.00 | 119.00 | 179.00 | 259.00 | 399.00 |
| 0.5 | 3.66 | 5.75 | 8.00 | 10.65 | 13.49 | 16.44 | 19.91 | 24.83 | 30.25 | 38.00 |
| 1 | 2.08 | 2.71 | 3.22 | 3.69 | 4.09 | 4.44 | 4.79 | 5.19 | 5.56 | 5.99 |
| 1.5 | 1.34 | 1.47 | 1.57 | 1.65 | 1.71 | 1.75 | 1.78 | 1.81 | 1.83 | 1.85 |
| 2 | 0.88 | 0.93 | 0.96 | 0.98 | 0.98 | 0.99 | 0.99 | 0.99 | 1.00 | 1.00 |

Using these state utilities, the optimal belief distribution p_s^* is determined by the entropy-disciplined maximization problem derived in the previous subsection:

$$p_s^* = \frac{1}{Z(\tau)} \exp\left(\frac{\tau}{U(w_s, \gamma)}\right), \quad Z(\tau) = \sum_u \exp\left(\frac{\tau}{U(w_u, \gamma)}\right).$$

For the logarithmic case ($\gamma = 1$), the normalizing condition requires solving

$$\sum_{s=1}^{10} \exp\left(\frac{\tau^*}{\log w_s}\right) = 1,$$

which yields

$$\tau^* \approx -10.85.$$

The corresponding optimal probabilities are then

$$p_s^* = \frac{\exp\left(\frac{\tau^*}{U(w_s, \gamma)}\right)}{\sum_u \exp\left(\frac{\tau^*}{U(w_u, \gamma)}\right)}.$$

Table 2 reports the values of τ^* , the associated optimal probabilities p_s^* , and the Shannon entropy $H(p) := -\sum_s p_s \log p_s$ for several levels of risk aversion.

Table 2: Optimal Belief Distributions for 10 Payoff States

| γ | τ^* | $H(p)$ | p_1^* | p_2^* | p_3^* | p_4^* | p_5^* | p_6^* | p_7^* | p_8^* | p_9^* | p_{10}^* |
|----------|-----------|--------|---------|---------|---------|---------|---------|---------|---------|---------|---------|------------|
| -1.0 | 211831.00 | 0.76 | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.00% | 0.01% | 3.76% | 30.76% | 65.46% |
| 0.0 | -696.10 | 1.14 | 0.00% | 0.00% | 0.00% | 0.00% | 0.01% | 0.46% | 3.95% | 14.41% | 31.31% | 49.85% |
| 0.5 | -65.22 | 1.58 | 0.00% | 0.03% | 0.21% | 0.83% | 2.46% | 5.67% | 10.84% | 17.90% | 26.40% | 35.65% |
| 1.0 | -10.85 | 2.06 | 0.90% | 2.12% | 3.83% | 5.92% | 8.26% | 10.73% | 13.28% | 15.83% | 18.34% | 20.80% |
| 1.5 | -1.00 | 2.30 | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% |
| 2.0 | -1.00 | 2.30 | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% | 10.00% |

The patterns in Table 2 highlight the central role of preferences in shaping beliefs. Risk-loving agents ($\gamma < 0$) assign sharply concentrated probabilities to extreme high-payoff states, producing low entropy and lottery-like behavior. Risk-neutral agents ($\gamma = 0$) begin to spread beliefs across the upper tail but still overweight extremes. As risk aversion increases ($\gamma > 0$), beliefs flatten progressively, converging toward uniformity and maximizing entropy. In short, the table shows how preferences alone dictate the perceived informational structure: convex utility generates concentrated optimism, while concave utility yields skeptical dispersion.

This numerical illustration also clarifies why entropy in this framework should be understood as an epistemic, rather than hedonic, object. The entropy values in Table 2 do not measure a direct taste for uncertainty. Instead, they summarize the shape of the optimal belief distribution chosen under the criterion developed above. Low entropy means that the agent behaves as if the state space contains structure that can be identified and exploited; high entropy means that the agent behaves as if such structure is absent and therefore assigns probability mass more

symmetrically across states.

In conclusion, the numerical example shows that what we call *subjective entropy* is not a passive statistic but a preference-disciplined, endogenous object: it reflects how risk attitudes shape the perception of patterns, uncertainty, and ultimately prices. Risk-loving agents effectively “see” exploitable structure in the form of extreme states, whereas risk-averse agents deny such structure and move toward maximum-entropy belief assignments as a safeguard against surprise. Thus, the dispersion of subjective probabilities is itself endogenous and preference-driven, rather than exogenously imposed. This is a central contribution of the thesis: entropy becomes a behavioral mirror of risk attitudes, transforming the perception of uncertainty into a structured, decision-relevant object.

In this sense, subjective entropy provides a unified way to connect preferences, beliefs, and asset pricing under uncertainty, and prepares the ground for the endogenous stochastic discount factor developed in the next section.

7 An Endogenous InSDF: Pricing with Informational Disagreement

The exogenous kernel introduced earlier provided a consistent way to map benchmark expectations into subjective expectations under disciplined distortions, but without grounding in optimizing behavior it remained a formal device. In this section we close the loop: the Information-based Stochastic Discount Factor (InSDF) is derived *endogenously* from agents' entropy-weighted preferences. This establishes micro-foundations for the kernel and shows how informational disagreement—rooted in heterogeneous belief formation—becomes an explicit, state-dependent component of the pricing kernel, and therefore a potential source of priced risk *relative* to a common no-arbitrage benchmark.

7.1 Equilibrium concept (partial equilibrium)

This section adopts a *partial-equilibrium* perspective. A common no-arbitrage pricing operator is taken as given, together with the aggregate (macro) state process. In particular, the benchmark distribution $\bar{\nu}$ and the pricing anchor \hat{r}^B (or equivalently a benchmark pricing kernel) are treated as *exogenous inputs* summarizing no-arbitrage restrictions. The purpose of the section is not to derive benchmark returns from a full market-clearing fixed point, but to study how *heterogeneous, preference-driven beliefs* generate disagreement wedges and certainty-equivalent gaps *relative* to the common benchmark.

Given objects. Fix a finite state space S and a benchmark probability measure $\bar{\nu}$. Fix an external pricing anchor \hat{r}^B for a reference (anchor) claim with state log return $r_{t+1}^B(s)$, or equivalently a benchmark pricing kernel $\hat{m}_{t+1}(s)$ that defines an arbitrage-free pricing operator:

$$P_t(X_{t+1}) = \mathbb{E}_t^{\bar{\nu}}[\hat{m}_{t+1}(s) X_{t+1}(s)].$$

In what follows, we work primarily with the benchmark-kernel representation $\hat{m}_{t+1}(s)$; the anchor \hat{r}^B is retained as the exogenous object pinning down the benchmark no-arbitrage environment.

InSDF factorization: a Radon–Nikodym chain-rule decomposition. Assume all relevant measures have full support on S , so that the Radon–Nikodym derivatives below are well defined. No-arbitrage implies a benchmark pricing representation: for any traded payoff X_{t+1} ,

$$P_t(X_{t+1}) = \mathbb{E}_t^{\bar{\nu}}[\hat{m}_{t+1}(s) X_{t+1}(s)].$$

For any other measure η_j , Radon–Nikodym implies that for all integrable Z ,

$$\mathbb{E}_t^{\bar{\nu}}[Z] = \mathbb{E}_t^{\eta_j}\left[\frac{d\bar{\nu}}{d\eta_j}(s) Z\right].$$

Taking $Z = \hat{m}_{t+1}X_{t+1}$ yields

$$P_t(X_{t+1}) = \mathbb{E}_t^{\eta_j}\left[\underbrace{\hat{m}_{t+1}(s) \frac{d\bar{\nu}}{d\eta_j}(s)}_{=: m_{t+1}^{(j)}(s)} X_{t+1}(s)\right].$$

Now use the RN *chain rule* through an intermediate belief η_i :

$$\frac{d\bar{\nu}}{d\eta_j} = \frac{d\bar{\nu}}{d\eta_i} \cdot \frac{d\eta_i}{d\eta_j}.$$

Therefore,

$$m_{t+1}^{(j)}(s) = \hat{m}_{t+1}(s) \frac{d\bar{\nu}}{d\eta_i}(s) \cdot \frac{d\eta_i}{d\eta_j}(s) = \underbrace{\left(\hat{m}_{t+1}(s) \frac{d\bar{\nu}}{d\eta_i}(s)\right)}_{=: m_{t+1}^{\text{bench},i}(s)} \cdot \underbrace{\frac{d\eta_i}{d\eta_j}(s)}_{=: m^{\eta_i|\eta_j}(s)=e^{\kappa(s)}}.$$

Hence the two-layer kernel is not arbitrary:

$$\boxed{m_{t+1}(s) = m_{t+1}^{\text{bench},i}(s) m^{\eta_i|\eta_j}(s) = m_{t+1}^{\text{bench},i}(s) \exp(\kappa(s)), \quad \kappa(s) := \log\left(\frac{d\eta_i}{d\eta_j}(s)\right).}$$

Macro shocks (e.g. consumption growth) are treated as exogenous and enter the benchmark component through $\hat{m}_{t+1}(s)$. We reserve the notation $m_{t+1}^C(s)$ for the explicit consumption-based implementation introduced below.

Agents' choices. Given the benchmark environment, each agent i chooses a subjective belief $\eta_i(\cdot)$ by optimizing an entropy-weighted criterion, yielding a closed-form optimal tilt η_i^* (Section 6, and Appendices A.4 and A.5). Belief choice implies a *benchmark-invariant* disagreement wedge between agents i and j ,

$$m^{\eta_i|\eta_j}(s) := \frac{\eta_i^*(s)}{\eta_j^*(s)} = \exp(\kappa(s)).$$

Economic implication (trade without arbitrage). Prices are common because they are pinned down by the benchmark operator $P_t(\cdot)$, but agents generally form different certainty equivalents under η_i^* . These certainty-equivalent wedges generate trade and volume *despite* the absence of arbitrage; disagreement is priced as a wedge in valuation relative to the benchmark, not as a violation of no-arbitrage.

We refer to this formulation as an *endogenous* SDF because the disagreement component is micro-founded: agents' subjective beliefs are chosen optimally under entropy-weighted preferences, and the resulting belief ratios generate a state-dependent wedge in valuation relative to the common benchmark pricing operator. By contrast, the exogenous kernel in Section 3 is a purely informational change-of-measure device that disciplines distortions around a given benchmark without deriving beliefs from preferences.

7.2 Summary derivation of the Endogenous InSDF

The purpose of this section is to present the key steps, state the resulting kernel compactly, and clarify its economic content—in particular, the distinct sources of risk embedded in the endogenous InSDF. The full derivation is deferred to Appendix A.6.

Let $S = \{s_1, \dots, s_n\}$ be a finite state space and let $\bar{\nu}$ denote a common

benchmark distribution over S (e.g. a market or no-arbitrage reference). To avoid any asset-specific circularity, belief formation is stated at the level of *states* rather than tied to a particular asset payoff: let $w(s)$ denote the agent's state-contingent wealth (or consumption) in state $s \in S$. Assets deliver state-contingent payoffs that contribute to $w(s)$, but the belief-tilt mechanism applies to the state space itself.

State-contingent utility notation. For compactness define

$$u_i^*(s) := u_i(w(s)), \quad s \in S, \quad (28)$$

and analogously $u_j^*(s) := u_j(w(s))$.

All probability vectors lie in the simplex

$$\Delta(S) = \left\{ p \in \mathbb{R}^n \mid \sum_{s \in S} p(s) = 1, p(s) \geq 0 \forall s \right\}.$$

1. Optimal belief formation under entropy-weighted informational preferences.

Each agent i selects a subjective belief $\eta_i(\cdot) \in \Delta(S)$ by maximizing an entropy-weighted objective relative to the benchmark $\bar{\nu}$. This benchmark-relative formulation is the multi-agent counterpart of the entropy-augmented criterion introduced in Section 6; the benchmark $\bar{\nu}$ enters here as the common reference measure around which subjective beliefs are tilted. A convenient representation is:

$$\max_{\eta_i \in \Delta(S)} \sum_{s \in S} \eta_i(s) u_i^*(s) \log \left(\frac{\bar{\nu}(s)}{\eta_i(s)} \right), \quad (29)$$

subject to $\sum_{s \in S} \eta_i(s) = 1$.

The first-order condition yields the exponential form used throughout the

thesis:

$$\eta_i^*(s) = \frac{\bar{\nu}(s) \exp\left(\frac{\tau_i}{u_i^*(s)}\right)}{Z_i}, \quad Z_i = \sum_{u \in S} \bar{\nu}(u) \exp\left(\frac{\tau_i}{u_i^*(u)}\right), \quad (30)$$

where $\tau_i := \theta_i/\alpha_i$ is the effective tilt parameter: θ_i is the simplex (normalization) multiplier and $\alpha_i > 0$ is the information-discipline parameter.

2. Constructing the disagreement wedge via relative belief distortions.

Let η_i^* and η_j^* denote the optimal beliefs of agents i and j , both tilted around the same benchmark $\bar{\nu}$. Define the belief-ratio kernel:

$$m^{\eta_i|\eta_j}(s) := \frac{\eta_i^*(s)}{\eta_j^*(s)} = \frac{Z_j}{Z_i} \exp\left(\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)}\right) \quad (31)$$

Interpretation (endogenous disagreement wedge). Equation (31) is the core object of the thesis: it isolates *disagreement* as a benchmark-invariant likelihood ratio between the two agents' *utility-maximizing* subjective probability measures. Both agents form beliefs by tilting a common benchmark $\bar{\nu}$ under the same information discipline; the resulting normalizers Z_i and Z_j ensure that each η_k^* is a proper probability distribution. When we take the ratio η_i^*/η_j^* , the benchmark $\bar{\nu}$ cancels out, so disagreement depends only on *relative* tilts.

The exponential term shows that the wedge is driven by the agents' state-by-state trade-off between (i) the *effective cost of belief distortion* (captured by the tilt strength τ_k) and (ii) the *utility level* in that state (captured by $u_k^*(s)$). In particular, each agent's contribution $\tau_k/u_k^*(s)$ can be read as a utility-scaled shadow cost of reallocating probability mass across states: states that deliver low utility to agent k carry a larger distortion index and are treated as more "expensive" to overweight under the information discipline. Since $u_k^*(s)$ is shaped by the agent's risk aversion, differences in risk preferences

map directly into differences in optimal belief tilts. Hence $m^{\eta_i|\eta_j}(s)$ quantifies how two agents, facing the same benchmark and the same set of states, end up assigning systematically different state probabilities purely because they solve different (preference-driven) belief-choice problems.

It is useful to write this ratio in exponential form:

$$m^{\eta_i|\eta_j}(s) = \exp(\kappa(s)), \quad \kappa(s) := \log\left(\frac{Z_j}{Z_i}\right) + \frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)}. \quad (32)$$

Thus $\kappa(s)$ is a *state-dependent disagreement wedge* summarizing relative belief reweighting.

Proposition 1 (Collapse to the benchmark kernel). *With a common benchmark \bar{v} and optimal beliefs tilted around \bar{v} , the belief-ratio kernel $m^{\eta_i|\eta_j}(s)$ is identically equal to 1 for all $s \in S$ if and only if $\kappa(s) = 0$ for all $s \in S$. A sufficient knife-edge case is that*

$$\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)} = c \quad \forall s \in S,$$

for some constant c , so that the state dependence cancels after normalization. In particular, the condition $\eta_i^* = \eta_j^*$ implies that the endogenous disagreement wedge disappears.

3. Embedding disagreement in an intertemporal pricing kernel.

Let $\rho_{t+1}(s)$ denote the exogenous macro state variable (log consumption growth or, more generally, log marginal-utility growth), written in the log convention. The endogenous InSDF is:

$$m_{t+1}(s) = \frac{\delta_j}{\delta_i} \exp(\kappa(s)) \exp(-\lambda_1 \rho_{t+1}(s)). \quad (33)$$

Interpretation: $\frac{\delta_j}{\delta_i}$ is the constant relative time-discount factor, $\exp(-\lambda_1 \rho_{t+1}(s))$ is the fundamental-risk component (high in bad states), and $\exp(\kappa(s))$ is

the informational disagreement wedge. Unlike τ_i , which is agent-specific and belief-generating, λ_1 belongs to the common benchmark risk channel and captures the sensitivity of marginal valuation to the macro state variable $\rho_{t+1}(s)$.

4. Pricing and the Euler restriction.

In this partial-equilibrium setting, the Euler equation is not used to determine the benchmark return process; it is a pricing restriction that maps a given payoff X_{t+1}^A into its arbitrage-free price $P_t^A = \mathbb{E}_t^{\bar{\nu}}[m_{t+1} X_{t+1}^A]$, from which the return $R_{t+1}^A = X_{t+1}^A/P_t^A$ is then implied.

No-arbitrage pricing under the benchmark anchor $\bar{\nu}$ takes the familiar form:

$$P_t^A = \mathbb{E}_t^{\bar{\nu}}[m_{t+1}(s) X_{t+1}^A(s)], \quad (34)$$

and defining $R_{t+1}^A(s) := X_{t+1}^A(s)/P_t^A$, equivalently

$$1 = \mathbb{E}_t^{\bar{\nu}}[m_{t+1}(s) R_{t+1}^A(s)] = \mathbb{E}_t^{\bar{\nu}}\left[\frac{\delta_j}{\delta_i} \exp(\kappa(s) - \lambda_1 \rho_{t+1}(s)) R_{t+1}^A(s)\right]. \quad (35)$$

This makes clear that the same $m_{t+1}(s)$ prices all assets, which differ only through their state-contingent payoffs $X_{t+1}^A(s)$ (or implied returns $R_{t+1}^A(s)$).

Role of the common benchmark. Because both agents tilt the same $\bar{\nu}$, comparative valuation depends on *relative* distortions. The benchmark enters only through normalizations (e.g. Z_i, Z_j); the economic wedge is $\kappa(s)$, i.e. the relative tilt across agents.

7.3 Sources of Risk: Consumption Growth as the ρ -Channel

To make the economic content of the endogenous InSDF explicit, we let $\rho_{t+1}(s)$ denote *log consumption growth* (or more generally marginal-utility growth), which acts as the exogenous macro risk driver in this partial-equilibrium environment.

This convention matches the notation discipline adopted throughout: asset returns are denoted by R^A , while ρ is reserved for the macro state variable.

The exogenous construction in Section 3 defined $\rho(s)$ as a state-dependent *log-return index* for a priced payoff $A_T(s)$. In the endogenous formulation, we now impose an explicit economic source of risk by interpreting the same object as *log consumption growth*. This commits the model to a standard consumption-based risk channel while preserving the informational (belief) channel developed above. For notational simplicity, we suppress the time subscript when no confusion arises and write $\rho(s)$ for $\rho_{t+1}(s)$.

Consumption growth convention. Let c_t denote consumption at date t , and let $c_{t+1}(s)$ denote next-period consumption in state $s \in S$. Define the state-contingent log consumption growth rate:

$$\rho(s) := \ln\left(\frac{c_{t+1}(s)}{c_t}\right). \quad (36)$$

Uncertainty enters through cross-state dispersion of $c_{t+1}(s)$ and therefore dispersion of $\rho(s)$.

Benchmark: CRRA SDF as an exponential in ρ . Assume CRRA preferences $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ so that $u'(c) = c^{-\gamma}$. The standard one-agent consumption-based SDF is

$$m_{t+1}^C(s) = \delta \frac{u'(c_{t+1}(s))}{u'(c_t)} = \delta \left(\frac{c_{t+1}(s)}{c_t}\right)^{-\gamma} = \delta \exp(-\gamma \rho(s)).$$

For completeness, the benchmark consumption-based stochastic discount factor

$$m_{t+1}^i(s) = \delta \frac{u'_i(C_{i,t+1}(s))}{u'_i(C_{i,t})}$$

follows from the standard marginal indifference-pricing argument under expected utility. Since this derivation is textbook, it is omitted here.

Hence writing the fundamental risk channel in exponential form is equivalent

to adopting the log-growth convention (36).

Endogenous InSDF with two risk channels. In our framework, the endogenous InSDF incorporates an informational disagreement term in addition to the fundamental consumption-risk term:

$$m(s) = \underbrace{\left[\frac{\delta_j}{\delta_i} \right]}_{\text{time preference}} \cdot \underbrace{\exp(\kappa(s))}_{\text{informational (disagreement) component}} \cdot \underbrace{\exp(-\lambda_1 \rho(s))}_{\text{fundamental component}}. \quad (37)$$

Here τ_i governs belief tilts (and therefore $\kappa(s)$), whereas λ_1 is the loading on the macro state $\rho(s)$ (under CRRA, $\lambda_1 = \gamma$).

Under the CRRA implementation, setting $\lambda_1 = \gamma$ recovers the standard consumption-based term $\exp(-\gamma\rho(s))$.

Equation (37) makes explicit that the model contains two separable sources of risk: (i) *fundamental risk* from stochastic consumption growth, and (ii) *informational risk* from the state-dependent disagreement wedge $\kappa(s)$ embedded in $\exp(\kappa(s))$.

The figure is a visual summary of the mechanism in (37). Holding the benchmark (fundamental) component fixed, the disagreement term $\exp(\kappa(s))$ acts as a relative state-reweighting: states where the belief ratio is large receive higher effective state-price weight, while states where the ratio is small are downweighted. The point is not that agents “prefer entropy,” but that heterogeneous optimal belief tilts generate a state-dependent likelihood ratio, and this likelihood ratio shifts valuation toward disagreement-intensive states.

7.4 Discussion and Economic Implications

The endogenous InSDF implies that prices reflect not only time preference and fundamental risk, but also informational disagreement. The wedge $\kappa(s)$ summarizes the relative distortion between agents’ endogenously formed beliefs and therefore shifts the valuation of states. When $\kappa(s)$ is more dispersed across s , the economy

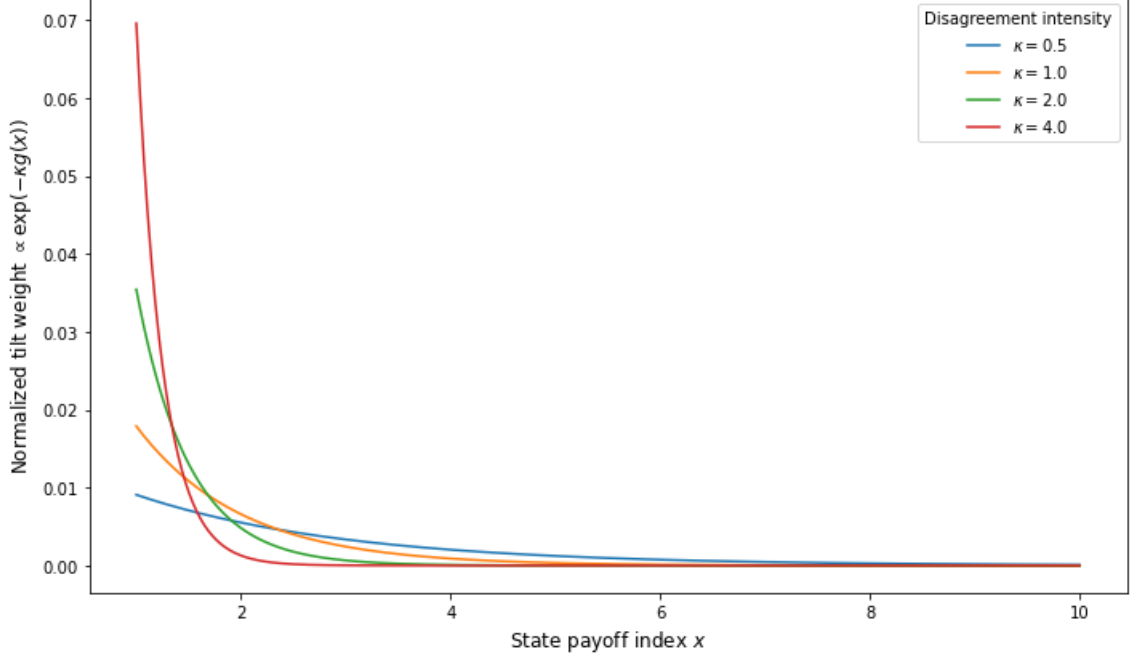


Figure 4: Illustrative impact of disagreement on valuation: holding the benchmark component fixed, the wedge $\exp(\kappa(s))$ reweights state prices and shifts value toward states in which belief ratios are large.

exhibits a larger informational premium: assets that load on high-disagreement states must offer compensation in expected returns.

No-trade benchmark and disagreement. In knife-edge cases where $\kappa(s) \equiv 0$ (full agreement in the sense of Proposition 1), the informational wedge disappears and the pricing kernel reduces to the benchmark fundamental component. This aligns with the no-trade intuition of Milgrom and Stokey (1982): when agents share the same effective beliefs and there is no informational motive to trade, gains-from-trade due purely to belief dispersion vanish. Importantly, this does not eliminate fundamental risk premia; it eliminates the *additional* informational wedge created by disagreement.

Risk-free rate as a special case. For a risk-free payoff, $R_T(s) \equiv R_{rf}$ is constant, and (35) gives:

$$1 = R_{rf} \mathbb{E}^{\bar{\nu}}[m(s)] \implies R_{rf} = \frac{1}{\mathbb{E}^{\bar{\nu}}[m(s)]}. \quad (38)$$

When disagreement vanishes ($\kappa(s) = 0$), the expression reduces to the standard benchmark mapping driven by time preference and fundamentals. Under the log-return convention $r_{rf} := \ln R_{rf}$, one may write $r_{rf} = -\ln(\mathbb{E}^{\bar{\nu}}[m(s)])$, making explicit how state dependence of the kernel maps into the risk-free rate.

Higher-order disagreements. The framework can accommodate richer forms of disagreement by allowing additional state-dependent wedges (e.g. linked to tail states or higher moments) to enter multiplicatively in the pricing kernel. These extensions preserve the same logic: state-contingent valuation arises from fundamentals and from informational wedges that reweight states.

Overall, the endogenous InSDF unifies preferences, beliefs, and prices into a single equilibrium structure. When disagreement is absent, prices reflect the benchmark fundamental component. When beliefs diverge endogenously, agents require compensation not only for fundamental risk but also for exposure to disagreement-intensive states. In this sense, the model interprets an important component of market premia as an *informational premium* arising from endogenously formed belief distortions.

7.5 Summary and Conclusion

This section has shown how the Information-based Stochastic Discount Factor can be derived *endogenously* once beliefs are no longer treated as exogenous primitives but as objects chosen under entropy-weighted informational discipline. This step is central to the paper. The exogenous construction developed earlier established that a benchmark-relative informational kernel can be obtained as a disciplined change of measure under a common pricing anchor. The present section closes that argument by showing that, once preferences are introduced, disagreement itself becomes an equilibrium valuation object.

The key conceptual move is the following. Under a common no-arbitrage benchmark $\bar{\nu}$, agents do not differ because they face different prices; they differ because they choose different subjective probability distortions around the same benchmark. These distortions are governed by the interaction between utility and informational discipline, and therefore depend on the agents' risk preferences. As a result, disagreement is not imposed from outside the model and it is not reduced to a behavioral residual. It is generated within the model as the benchmark-invariant ratio between utility-maximizing subjective beliefs,

$$m^{\eta_i|\eta_j}(s) = \frac{\eta_i^*(s)}{\eta_j^*(s)} = \exp(\kappa(s)),$$

where $\kappa(s)$ summarizes the state-dependent informational wedge between agents.

This yields the endogenous InSDF in multiplicative form:

$$m_{t+1}(s) = \frac{\delta_j}{\delta_i} \exp(\kappa(s)) \exp(-\lambda_1 \rho_{t+1}(s)).$$

The structure is economically transparent. The ratio δ_j/δ_i captures relative time preference, the term $\exp(-\lambda_1 \rho_{t+1}(s))$ captures the benchmark fundamental-risk channel, and the term $\exp(\kappa(s))$ captures the informational disagreement channel. The resulting pricing kernel therefore contains two distinct but complementary

layers: a standard macroeconomic source of priced risk, and an endogenous informational source of priced risk arising from heterogeneous belief formation. This decomposition makes explicit that disagreement does not replace fundamentals; it reweights state prices around them.

Several implications follow immediately. First, no-arbitrage is preserved throughout. Since all agents are anchored to the same benchmark pricing operator, disagreement does not imply inconsistent prices in the arbitrage sense. What it generates are heterogeneous certainty equivalents and valuation wedges relative to a common price. In this way, the model provides a disciplined route from heterogeneous beliefs to trade and volume without abandoning benchmark pricing consistency. Second, disagreement is priced only insofar as it is *state dependent*. If $\kappa(s)$ is constant across states, it cancels through normalization and has no pricing effect. What matters for returns is not the level of disagreement as such, but the way disagreement tilts relative state weights. Third, because the wedge is benchmark-invariant, the economically relevant object is not each agent's benchmark distortion in isolation, but the pairwise ratio between optimal beliefs. This is what makes disagreement a genuinely cross-sectional pricing object.

The section also clarifies the role of market efficiency as a limiting case of the model rather than as a maintained assumption. When belief wedges vanish, so that $\eta_i^* = \eta_j^*$ and therefore $\kappa(s) \equiv 0$, the endogenous InSDF collapses to the benchmark kernel. In that knife-edge case, disagreement ceases to contribute to prices, trade, or informational premia. Efficiency is therefore recovered as the special case in which preference-driven belief distortions no longer differ across agents. This is an important implication of the framework: the model does not assume away disagreement in order to recover benchmark pricing, but rather shows under what conditions benchmark pricing re-emerges from the endogenous convergence of beliefs.

More broadly, the contribution of this section is to connect three objects that are usually kept separate in standard asset-pricing models: subjective beliefs,

informational structure, and state pricing. In the present framework, risk aversion affects not only marginal valuation of payoffs, but also the selection of subjective probabilities themselves; those probabilities generate benchmark-invariant wedges across agents; and those wedges enter the pricing kernel as an additional source of priced risk. The endogenous InSDF is therefore the mechanism through which subjective, preference-driven disagreement acquires asset-pricing content.

For the purposes of Paper 1, this is the main theoretical result. The model shows that disagreement can be formalized as a measurable informational wedge consistent with no-arbitrage, micro-founded in optimal belief formation, and embedded directly in the stochastic discount factor. The next natural step is to study the implications of this object in more specific asset-pricing environments. That is the role of Paper 2, where the disagreement wedge is taken to applications such as binomial pricing, the Information CAPM, and ambiguity-sensitive valuation.

8 Conclusion: Information, Belief Distortion, and Asset Pricing Equilibrium

This first paper developed a framework in which *beliefs are endogenous objects*: agents choose subjective probabilities to maximize an entropy-weighted utility criterion. In incomplete markets, subjective valuation depends not only on payoffs but also on the agent’s epistemic posture—captured by the belief structure itself. Risk preferences (and the informational discipline embedded through the entropy term) determine the optimal belief tilt, so perceived likelihoods are functions of preferences rather than exogenous inputs or learning. The central insight is that risk aversion shapes not only how agents evaluate outcomes, but also which outcomes they deem likely through the utility-maximizing probability p^* .

This structure gives rise to a class of *subjective entropy functions* that quantify belief distortions. Highly risk-averse agents gravitate toward flatter, higher-entropy beliefs, reflecting skepticism about exploitable patterns or arbitrage opportunities. Risk-loving agents, by contrast, concentrate probability mass on favorable states, effectively behaving as if the world is pattern-rich and extreme gains are attainable. These belief distortions generate utility-maximizing probability distributions that differ across agents even in the absence of asymmetric information. In this sense, disagreement and dispersion in markets can emerge purely from preference heterogeneity, providing a foundation for phenomena often attributed to informational gaps.

The Information-based stochastic discount factor (InSDF) introduced earlier was first derived as an informational change of measure: a disciplined transformation of a benchmark distribution into subjective beliefs. The paper then closed the loop by embedding this construction into a pricing framework under a common no-arbitrage operator. Optimal subjective probability measures generated by entropy-weighted utility imply a benchmark-invariant *belief-ratio kernel* (the disagreement wedge), which enters the pricing kernel multiplicatively. The resulting InSDF therefore

captures both (i) time preference and fundamental risk through the benchmark deflator, and (ii) *informational disagreement* through the endogenous likelihood-ratio wedge. It preserves the classical no-arbitrage restriction while carrying endogenous informational content: prices reflect not only expected payoffs and risk aversion, but also the cross-sectional structure of belief heterogeneity.

In sum, disagreement in this framework is not noise but an *endogenous informational asymmetry* produced by preference-driven belief tilts. It both *drives market activity*—trade volume and rebalancing through heterogeneous certainty equivalents—and is *priced* through the InSDF as an informational premium. As the belief-ratio kernel flattens (disagreement fades), the informational wedge collapses and excess returns shrink toward the benchmark; as dispersion in entropy-weighted beliefs widens, premia rise precisely in those states where belief ratios diverge most. This pins observed returns to the equilibrium reconciliation of heterogeneous, preference-shaped beliefs under a common no-arbitrage pricing discipline.

Beyond its theoretical contribution, the framework opens several avenues. First, it provides a micro-founded rationale for persistent disagreement, speculative trading, and excess volatility in incomplete markets. Second, it reframes part of the risk premium as an *informational premium*, suggesting empirical tests based on entropy-weighted measures of belief dispersion and disagreement wedges. Finally, it offers a bridge between classical asset pricing and information economics: markets are not priced by risk and time alone, but also by the informational consistency (or inconsistency) of the beliefs that participants bring to them. In this way, subjective entropy becomes both a diagnostic and a driver of asset-pricing outcomes.

Part II

**Disagreement in Asset Pricing with Endogenous
Beliefs: Empirical Implications and Applications**

Abstract

This paper develops the applied implications of a belief-based asset pricing framework in which agents optimize entropy-weighted utility functions under subjective uncertainty. Building on prior theoretical work, it incorporates informational disagreement into four related settings: the Cox–Ross–Rubinstein (CRR) option pricing model, the Capital Asset Pricing Model (CAPM), expected utility under ambiguity, and the dynamic convergence of beliefs.

First, we introduce a disagreement-tilted binomial pricing model in which agents hold distorted beliefs derived from entropy-penalized utility. A belief-distortion parameter κ modifies the benchmark pricing measure, generating heterogeneous valuations and providing a structural explanation for trading volume and persistent dispersion in expectations. Second, we derive an information-adjusted CAPM (InCAPM) in which the pricing kernel embeds both risk preferences and informational asymmetries. The resulting informational beta captures belief dispersion as a priced factor, leading to a generalized risk premium. Third, we present a minimal resolution of the Ellsberg paradox using an entropy-based adjustment to utility, which generates ambiguity aversion without violating additivity or Bayesian updating. Fourth, we develop a theory of belief convergence showing that, in the absence of new information, disagreement cannot persist indefinitely under common pricing discipline. As informational wedges compress over time, subjective valuations align and market prices converge toward their benchmark no-arbitrage relation, providing a microfoundation for informational efficiency as an endogenous limiting outcome.

Taken together, these results offer a unified framework in which belief heterogeneity, entropy, ambiguity, and convergence are treated as core components of asset pricing rather than as anomalies. Classical models emerge as limiting cases when disagreement vanishes or beliefs converge.

9 Introduction

This paper is the second in a two-part dissertation on disagreement in asset pricing with endogenous beliefs. The first paper develops the theoretical foundation of the framework. There, subjective beliefs are derived endogenously from entropy-weighted preferences, and the information-based stochastic discount factor (InSDF) is obtained as the pricing object associated with optimal belief formation under a common benchmark.

The central result of that first paper is that subjectivity need not be treated as irrational or ad hoc: heterogeneous beliefs can be disciplined within a coherent no-arbitrage structure. This second paper develops the applied and dynamic implications of that theoretical construction. Its purpose is to show how the framework can be embedded in canonical asset-pricing settings, how disagreement modifies their economic interpretation, and how disagreement itself evolves over time when no new information arrives. The guiding idea is simple but consequential: agents do not merely disagree about payoffs, but about the informational structure used to assess them. Once such disagreement is represented through entropy-weighted belief distortions, it becomes possible to incorporate heterogeneous expectations into standard pricing models without abandoning internal consistency, Bayesian discipline, or the no-arbitrage restrictions that organize modern finance.

The paper makes four contributions.

1. **A binomial asset-pricing model with disagreement.** We extend the classical Cox–Ross–Rubinstein framework by allowing agents to hold distinct subjective beliefs around a common benchmark pricing measure. Disagreement is represented as a relative, state-dependent tilt, so that the standard CRR model appears as the limiting case in which informational distortions vanish. This formulation preserves the benchmark pricing structure while showing how heterogeneous beliefs generate divergent valuations, trade, and volume even when arbitrage opportunities are absent.

2. **An Information CAPM (InCAPM).** We derive an asset-pricing relation in which expected returns depend not only on exposure to the market factor, but also on exposure to disagreement-intensive states. The resulting model introduces an informational component into the pricing kernel and yields an informational beta that captures how an asset co-moves with belief dispersion. In this way, disagreement is not treated as noise external to valuation, but as a priced source of risk whenever informational wedges are state dependent and interact with fundamentals.
3. **A minimal resolution of the Ellsberg paradox.** We show that ambiguity aversion can be interpreted within the same informational structure that generates disagreement in asset pricing. By introducing an entropy-based penalty associated with unresolved uncertainty, the model explains why agents may prefer known over unknown probability structures without abandoning additivity or Bayesian updating. The Ellsberg problem is therefore reinterpreted not as a failure of rationality, but as a rational response to informational incompleteness.
4. **A theory of belief convergence and informational efficiency.** We show that, in the absence of new information, disagreement cannot persist indefinitely under common pricing discipline. As informational wedges compress over time, subjective valuations align and market prices converge toward their benchmark no-arbitrage relation. Informational efficiency is therefore not imposed as a primitive assumption, but emerges as the limiting outcome of a system in which disagreement is gradually eroded by learning, common pricing, and the absence of fresh informational asymmetries.

Taken together, these four applications clarify the economic content of the framework developed in the first paper. The binomial model shows how disagreement can be embedded in a canonical arbitrage-based setting; the InCAPM shows how disagreement becomes a priced factor in equilibrium asset returns; the Ellsberg

application shows that ambiguity aversion can be understood as a special case of entropy-sensitive belief formation; and the convergence analysis shows how informational efficiency can arise endogenously as disagreement fades over time. In each case, the classical benchmark is recovered when informational distortions disappear or beliefs converge.

The broader implication is that disagreement is not a peripheral friction layered onto otherwise standard models. It is a disciplined way of representing how agents form probabilities when the structure of uncertainty itself is contested. This perspective preserves the core logic of asset pricing while expanding its scope: prices remain constrained by no-arbitrage, but beliefs need not coincide, and those differences can have systematic valuation consequences both statically and dynamically.

The remainder of the paper is organized as follows. Section 10 develops a binomial asset-pricing model with disagreement as an extension of the CRR framework. Section 11 presents the Information CAPM and studies the role of informational exposure in expected returns. Section 12 develops a minimal entropy-based resolution of the Ellsberg paradox and shows its connection to the disagreement-based pricing kernel. Section 13 studies belief convergence and its implications for endogenous informational efficiency.

10 A Disagreement Binomial Option-Pricing Model

This section develops a binomial option-pricing model under belief disagreement, building on the equilibrium stochastic discount factor derived in Munoz-Alonso (2025). The goal is to embed the information-based disagreement structure into the classical Cox–Ross–Rubinstein (CRR) framework without abandoning its benchmark arbitrage-free logic. The resulting model preserves the binomial architecture of CRR while allowing agents to hold distinct subjective assessments of future states around a common benchmark measure.

A central point is that disagreement does not price assets in isolation. It does not replace the benchmark pricing operator, nor does it generate an independent valuation rule detached from the CRR structure. Rather, disagreement acts as a relative, state-dependent distortion around a common benchmark, altering how agents weight the up and down states while leaving the underlying no-arbitrage discipline intact. In this sense, the model extends CRR rather than overturning it: the classical case is recovered when the disagreement wedge vanishes.

This formulation has two important implications. First, it preserves internal pricing consistency: disagreement changes the relative valuation of states, but it does not generate arbitrage opportunities. Second, it provides a natural mechanism through which heterogeneous beliefs can sustain divergent option valuations, and therefore trade and volume, even when all agents remain anchored to the same benchmark pricing structure.

The section proceeds as follows. Subsection 10.1 briefly recalls the standard CRR model as the benchmark case. Subsection 10.2 introduces the disagreement-based extension and the associated κ -tilted measure. Subsection 10.3 studies the effect of disagreement on call option prices. Subsection 10.4 compares the benchmark and disagreement-based frameworks in a one-step numerical example. Section 10.5 explains why the model remains arbitrage-free and how disagreement can sustain trade volume.

10.1 CRR Pricing with Disagreement: the κ -Tilt

We now embed disagreement into the classical Cox–Ross–Rubinstein (CRR) binomial model. The construction follows the same logic as in Paper 1 (see Munoz-Alonso (2025)), but adapted to the two-state binomial setting. We proceed in two steps. First, we recover the standard CRR benchmark pricing measure as the unique no-arbitrage reference point. Second, we introduce disagreement as a relative, state-dependent tilt around that benchmark. The key point is that disagreement does not price assets in isolation. Rather, it distorts the relative weighting of states around a common benchmark pricing measure.

Binomial environment and notation. Consider a one-period binomial model with a risk-free asset and a risky asset:

$$B_{t+1} = \rho_f B_t, \quad S_{t+1} \in \{uS_t, dS_t\}, \quad u > d > 0, \quad (39)$$

where $\rho_f \equiv 1 + r$ denotes the gross risk-free return. Define the risky asset's state-contingent gross return by

$$\rho(s) := \frac{S_{t+1}(s)}{S_t}, \quad \rho(u) = u, \quad \rho(d) = d. \quad (40)$$

Throughout this section, q^* denotes the CRR benchmark up-state probability, and ν^* the associated benchmark pricing measure, with

$$\nu^*(u) = q^*, \quad \nu^*(d) = 1 - q^*.$$

Agents are indexed by i and j , with tilt intensities τ_i, τ_j and state-contingent utility objects $U_i(s), U_j(s)$, as in Paper 1.

Step 1: Kullback–Leibler (KL) projection recovers CRR. The benchmark pricing measure is obtained as the KL projection of an arbitrary reference belief $\bar{\nu}$

onto the set of measures satisfying the martingale restriction:

$$\nu^* = \arg \min_{\nu} \sum_{s \in \{u,d\}} \nu(s) \log \left(\frac{\nu(s)}{\bar{\nu}(s)} \right) \quad \text{s.t.} \quad \sum_s \nu(s) = 1, \quad \sum_s \nu(s) \rho(s) = \rho_f. \quad (41)$$

The solution is the exponential tilt

$$\nu^*(s) = \frac{\bar{\nu}(s) e^{-\lambda \rho(s)}}{\sum_{s' \in \{u,d\}} \bar{\nu}(s') e^{-\lambda \rho(s')}}. \quad (42)$$

Here $\bar{\nu}$ denotes the common reference belief, while ν^* denotes its no-arbitrage projection in the binomial setting; once the projection is obtained, ν^* becomes the relevant benchmark pricing measure.

In the two-state binomial model, this recovers the standard CRR benchmark. Indeed, if q denotes the up-state probability, the martingale condition implies

$$qu + (1 - q)d = \rho_f,$$

hence

$$q^* = \frac{\rho_f - d}{u - d}. \quad (43)$$

Therefore the benchmark pricing measure is uniquely pinned down by no-arbitrage:

$$\nu^*(u) = q^*, \quad \nu^*(d) = 1 - q^*.$$

Given any one-period payoff X_1 , benchmark pricing is then

$$P_0(X) = \frac{1}{\rho_f} \mathbb{E}_{\nu^*}[X_1]. \quad (44)$$

Equivalently, the Euler restriction can be written as

$$\mathbb{E}_{\nu^*}[m^* R_{X,1}] = 1, \quad m^* = \frac{1}{\rho_f}, \quad R_{X,1} := \frac{X_1}{P_0(X)}. \quad (45)$$

Step 2: disagreement as a state tilt around the benchmark. Once the benchmark ν^* is fixed, disagreement is introduced through a relative state-by-state wedge between agents i and j :

$$\kappa_{ij}(s) := \log\left(\frac{Z_j}{Z_i}\right) + \frac{\tau_i}{U_i(s)} - \frac{\tau_j}{U_j(s)}. \quad (46)$$

This object is the two-state specialization of the relative log-likelihood ratio developed in Paper 1. It induces the agent-pair tilted pricing measure

$$\nu_{ij}^\kappa(s) = \frac{\nu^*(s)e^{\kappa_{ij}(s)}}{\sum_{s' \in \{u,d\}} \nu^*(s')e^{\kappa_{ij}(s')}}. \quad (47)$$

Thus disagreement does not replace the benchmark pricing operator. It modifies state weights around ν^* , preserving the no-arbitrage structure while allowing heterogeneous valuations.

The corresponding disagreement-adjusted up-state probability is

$$q_{\kappa,ij} := \nu_{ij}^\kappa(u) = \frac{q^* e^{\kappa_{ij}(u)}}{q^* e^{\kappa_{ij}(u)} + (1 - q^*) e^{\kappa_{ij}(d)}}. \quad (48)$$

Integrated disagreement-adjusted probability. The constant term $\log(Z_j/Z_i)$ appears in both states and therefore cancels under normalization. What matters for prices is not the absolute level of the wedge, but its state dependence. Therefore the disagreement-adjusted probability can be written directly as

$$q_{\kappa,ij} = \frac{q^* \exp\left(\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)}\right)}{q^* \exp\left(\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)}\right) + (1 - q^*) \exp\left(\frac{\tau_i}{U_i(d)} - \frac{\tau_j}{U_j(d)}\right)} \quad (49)$$

Equation (49) is the key binomial object in the section: it combines the CRR benchmark probability q^* with the disagreement wedge. It shows that prices move because disagreement reweights states *relative* to the benchmark, not because it introduces a separate pricing rule.

See Appendix B.1 for the full derivation of (49).

To isolate the source of the tilt, define the relative state dependence by

$$\Delta\kappa_{ij} := \kappa_{ij}(u) - \kappa_{ij}(d) = \left(\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)} \right) - \left(\frac{\tau_i}{U_i(d)} - \frac{\tau_j}{U_j(d)} \right). \quad (50)$$

despite disagreement, arbitrage is nevertheless ruled out. The reason is that disagreement does not replace the CRR pricing benchmark; it only distorts the subjective valuation of states around it. All agents remain anchored to the same no-arbitrage benchmark, even if they assign different relative weights to the up and down states. Hence disagreement generates heterogeneous valuations, and therefore a motive for trade, without generating riskless profit opportunities.

Reading the tilt. Equation (49) makes clear that disagreement matters only through relative state reweighting. When $\Delta\kappa_{ij} > 0$, agent i tilts relatively more probability mass toward the up state than agent j ; when $\Delta\kappa_{ij} < 0$, agent i tilts relatively more mass toward the down state. Thus disagreement affects prices only insofar as it changes the *relative* valuation of the two states.

The intuition is simple. The benchmark probability q^* carries the no-arbitrage pricing content of the CRR model. Disagreement does not overturn that benchmark. It acts on top of it, amplifying or attenuating the relative importance of the up and down states. What matters economically is not disagreement per se, but whether disagreement varies across states.

Economic reading. The benchmark measure ν^* captures the state prices implied by no-arbitrage. The disagreement term introduces an informational wedge around those state prices. Hence the model remains anchored in standard asset-pricing logic: the benchmark determines the reference valuation, while heterogeneous beliefs distort relative state weights around it.

This is why the framework remains disciplined. If disagreement vanishes, or if it is identical across states, the benchmark alone determines prices and the model collapses back to CRR. Only genuinely state-dependent disagreement changes

valuation.

Equivalent SDF representation. The same pricing rule can be written as a disagreement-adjusted stochastic discount factor:

$$m_{ij}^\kappa(s) = \frac{1}{\rho_f} \frac{d\nu_{ij}^\kappa}{d\nu^*}(s), \quad (51)$$

so that

$$P_0(X) = \mathbb{E}_{\nu^*} [m_{ij}^\kappa(s) X_1(s)] = \frac{1}{\rho_f} \mathbb{E}_{\nu_{ij}^\kappa} [X_1]. \quad (52)$$

Hence disagreement can be represented either as a change of measure $\nu^* \rightarrow \nu_{ij}^\kappa$ or as a wedge in the stochastic discount factor around the benchmark. The pricing implication is the same in either representation.

CRR as a nested special case. The nesting result is easiest to see directly from equation (49). Suppose first that disagreement vanishes. Then

$$\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)} = 0, \quad \frac{\tau_i}{U_i(d)} - \frac{\tau_j}{U_j(d)} = 0,$$

and equation (49) immediately reduces to

$$q_{\kappa,ij} = \frac{q^*}{q^* + (1 - q^*)} = q^*.$$

More generally, the classical CRR model is recovered whenever disagreement is state-invariant:

$$\kappa_{ij}(u) = \kappa_{ij}(d) \implies q_{\kappa,ij} = q^*. \quad (53)$$

To see this, let $\kappa_{ij}(u) = \kappa_{ij}(d) = \bar{\kappa}$. Then equation (48) becomes

$$q_{\kappa,ij} = \frac{q^* e^{\bar{\kappa}}}{q^* e^{\bar{\kappa}} + (1 - q^*) e^{\bar{\kappa}}} = \frac{q^* e^{\bar{\kappa}}}{e^{\bar{\kappa}}} = q^*.$$

Thus the common factor cancels under normalization. In that case,

$$\nu_{ij}^\kappa = \nu^*,$$

and the model collapses exactly to standard CRR pricing. The conclusion is therefore sharp: disagreement affects prices only through *state dependence*. A pure level shift has no pricing effect.

Option pricing under disagreement. Let K denote the strike price of a one-period European call option, with terminal payoffs

$$C_u = \max(uS_0 - K, 0), \quad C_d = \max(dS_0 - K, 0).$$

Its disagreement-adjusted price is

$$C_0^\kappa = \frac{1}{\rho_f} \left[q_{\kappa,ij} C_u + (1 - q_{\kappa,ij}) C_d \right]. \quad (54)$$

In an n -period recombining tree, assuming for simplicity that the disagreement tilt remains constant across steps, the call price becomes

$$C_0^\kappa = \frac{1}{\rho_f^n} \sum_{k=0}^n \binom{n}{k} q_{\kappa,ij}^k (1 - q_{\kappa,ij})^{n-k} \max(S_0 u^k d^{n-k} - K, 0). \quad (55)$$

The tree structure is unchanged; only the relative weighting of terminal nodes is modified by disagreement. This is precisely why the model extends CRR rather than replacing it.

10.2 Visualizing Belief-Tilted Option Pricing: The Role of κ

To illustrate the pricing implications of disagreement, we now visualize how the tilt parameter κ affects European call option values in the binomial model developed above. The goal is not to introduce a new pricing rule, but to show how disagreement modifies valuations by reweighting states around the common CRR benchmark.

In this sense, the figure provides a direct illustration of equation (49): once the benchmark probability q^* is fixed, disagreement changes option prices through relative state tilting.

We consider a recombining binomial tree with the following parameters:

- Initial stock price: $S_0 = 100$,
- Log-step volatility: $\sigma = 0.10$, so that $u = e^\sigma \approx 1.105$ and $d = e^{-\sigma} \approx 0.905$,
- Number of steps: $n = 50$,
- Annual risk-free rate: $r = 5\%$, implying a per-step gross risk-free return $\rho_f = e^{r/n}$,
- Benchmark CRR probability:

$$q^* = \frac{\rho_f - d}{u - d} \approx 0.480,$$

- Disagreement levels: $\kappa \in \{-1.0, -0.5, 0, 0.5, 1.0\}$.

In the numerical exercise, we use the reduced-form tilt

$$q_\kappa = \frac{q^* e^{\kappa \ln u}}{q^* e^{\kappa \ln u} + (1 - q^*) e^{\kappa \ln d}}, \quad (56)$$

which is the two-state specialization of the disagreement-adjusted probability in equation (49). The case $\kappa = 0$ reproduces the benchmark CRR model, while positive and negative values of κ tilt relative probability mass toward the up and down states, respectively.

For each value of κ , we compute European call prices across strike prices ranging from $K = 60$ to $K = 140$. The resulting curves are shown in Figure 5.

Interpretation. Figure 5 shows that disagreement shifts the entire option-pricing schedule in a systematic way. The benchmark curve $\kappa = 0$ corresponds to the standard CRR model. When $\kappa > 0$, the up state receives relatively greater weight,




Figure 5: European call prices under disagreement-tilted binomial probabilities. The benchmark case $\kappa = 0$ reproduces the standard CRR valuation. Positive values of κ tilt relative probability mass toward the up state and increase call prices, while negative values tilt mass toward the down state and lower valuations. The figure shows that disagreement does not replace the CRR benchmark; it modifies option values through relative state reweighting around the common benchmark probability q^* .

so the expected payoff of the call rises and the valuation curve shifts upward. When $\kappa < 0$, the down state receives relatively greater weight, reducing the expected payoff and shifting the curve downward. The effect is economically intuitive: call options are convex claims that benefit from higher weight on favorable states and lose value when those states are discounted.

The figure also clarifies the economic role of disagreement. What changes across curves is not the existence of a pricing anchor, but the subjective valuation of states around that anchor. All agents remain disciplined by the same CRR benchmark, so disagreement does not generate arbitrage opportunities. Instead, it produces heterogeneous but internally coherent valuations, which is precisely the mechanism through which disagreement can sustain trade and volume without violating no-arbitrage. In this way, the figure provides a visual bridge between the theoretical result of Paper 1 and its binomial implementation here: disagreement projects subjective valuations onto a common benchmark pricing structure rather

than replacing that structure altogether.

10.3 Comparison between the κ -Tilted and CRR Models

To illustrate the internal consistency of the model, we consider a one-step European call option and compare the standard CRR valuation with the disagreement-tilted valuation.³ The purpose of the example is not to replace the CRR benchmark, but to show how disagreement modifies subjective valuations around that benchmark while preserving the same no-arbitrage pricing structure.

We use the following parameter values:

- $S_0 = 100$: initial stock price,
- $u = e^{0.10} \approx 1.105$: gross return in the up state,
- $d = e^{-0.10} \approx 0.905$: gross return in the down state,
- $r = 0.05$: one-period risk-free rate, so that $\rho_f = 1 + r = 1.05$,
- $K = 100$: strike price of the European call option,
- $\kappa = 0.3$: disagreement parameter.

The terminal stock prices are therefore

$$S_u = uS_0 \approx 110.5, \quad S_d = dS_0 \approx 90.5,$$

and the corresponding call payoffs are

$$C_u = \max(S_u - K, 0) = 10.5, \quad C_d = \max(S_d - K, 0) = 0.$$

Under the benchmark CRR model, the risk-neutral probability is

$$q^* = \frac{\rho_f - d}{u - d} = \frac{1.05 - 0.905}{1.105 - 0.905} \approx 0.725.$$

³Extending the exercise to more steps is straightforward, although the one-step case makes the role of disagreement particularly transparent.

The benchmark call value is therefore

$$V_0^{\text{CRR}} = \frac{1}{\rho_f} [q^* C_u + (1 - q^*) C_d] = \frac{1}{1.05} [0.725 \cdot 10.5] \approx 7.26. \quad (57)$$

We now introduce disagreement through the reduced-form tilt defined in equation (56). For $\kappa = 0.3$,

$$q_\kappa = \frac{q^* u^\kappa}{q^* u^\kappa + (1 - q^*) d^\kappa} = \frac{0.725 \cdot 1.105^{0.3}}{0.725 \cdot 1.105^{0.3} + 0.275 \cdot 0.905^{0.3}} \approx 0.736.$$

The tilted call value becomes

$$V_0^\kappa = \frac{1}{\rho_f} [q_\kappa C_u + (1 - q_\kappa) C_d] = \frac{1}{1.05} [0.736 \cdot 10.5] \approx 7.38. \quad (58)$$

The comparison is informative. The benchmark probability $q^* \approx 0.725$ yields the CRR price $V_0^{\text{CRR}} \approx 7.26$, while the optimistic tilt $\kappa = 0.3$ raises the effective up-state probability to $q_\kappa \approx 0.736$, producing the slightly higher valuation $V_0^\kappa \approx 7.38$. The difference is modest, but economically meaningful: disagreement shifts valuation through state reweighting, not through any change in the benchmark pricing rule.

The benchmark case is nested immediately. If $\kappa = 0$, then

$$q_\kappa = \frac{q^* \cdot 1}{q^* \cdot 1 + (1 - q^*) \cdot 1} = q^*,$$

so that $V_0^\kappa = V_0^{\text{CRR}}$. Thus the disagreement-tilted model contains the classical CRR model as a special case.

Table 3: One-step call option: disagreement tilts around the CRR benchmark

| Quantity | $\kappa = -0.3$ | CRR ($\kappa = 0$) | $\kappa = 0.3$ |
|---------------------------------|-----------------|----------------------|----------------|
| Up-state probability q | 0.712 | 0.725 | 0.736 |
| Down-state probability $1 - q$ | 0.288 | 0.275 | 0.264 |
| Up-state payoff C_u | 10.500 | 10.500 | 10.500 |
| Down-state payoff C_d | 0.000 | 0.000 | 0.000 |
| Expected payoff $\mathbb{E}[C]$ | 7.479 | 7.613 | 7.729 |
| Discount factor $1/\rho_f$ | 0.952 | 0.952 | 0.952 |
| Call value V_0 | 7.123 | 7.250 | 7.361 |

Table 3 makes the mechanism especially clear. A positive tilt raises the relative weight placed on the up state and therefore increases the option valuation; a negative tilt has the opposite effect. For small values of κ , the resulting valuation changes are approximately symmetric around the benchmark, although the exponential form of the tilt implies that this symmetry need not persist for larger distortions.

The economic interpretation is straightforward. Agents remain anchored to the same CRR benchmark, but they project different subjective valuations onto that benchmark through the tilt q_κ . Hence disagreement does not create arbitrage opportunities. It creates heterogeneous but internally coherent valuations around a common no-arbitrage reference point. This is exactly the mechanism emphasized in Paper 1: disagreement affects subjective state valuation, not the existence of the pricing benchmark itself.

In this sense, the κ -tilted model generalizes CRR without overturning it. It preserves discounting, replication, and no-arbitrage, while adding belief dispersion as a structural source of valuation differences, trade, and volume.

10.4 Disagreement, Arbitrage and Market Activity

The binomial extension developed above makes it possible to separate two distinct objects: the common benchmark pricing operator and the agents' subjective valuation wedges around that benchmark. This distinction is crucial. Disagreement does not by itself create a new arbitrage-free pricing rule. Rather, it modifies

how agents value states relative to a benchmark that remains disciplined by the no-arbitrage logic of the CRR model.

In particular, the benchmark price of any payoff X_1 is still anchored by

$$P_0(X) = \frac{1}{\rho_f} \mathbb{E}_{\nu^*}[X_1], \quad (59)$$

where ν^* is the common CRR benchmark measure. Under disagreement, agents may instead evaluate the same payoff through their tilted measure ν_{ij}^κ , so that their subjective valuation takes the form

$$V_{0,ij}^\kappa(X) = \frac{1}{\rho_f} \mathbb{E}_{\nu_{ij}^\kappa}[X_1]. \quad (60)$$

Equation (60) should not be interpreted as replacing the market pricing rule in equation (59). It represents the agent-pair's subjective valuation of the payoff around the common benchmark. The difference

$$V_{0,ij}^\kappa(X) - P_0(X)$$

is therefore a *valuation wedge*, not an arbitrage spread.

This point is central. Arbitrage would require an inconsistency in the benchmark pricing structure itself. But disagreement, as modeled here, does not break that structure. All agents remain anchored to the same benchmark no-arbitrage operator, even though they may assign different relative weights to the up and down states. Hence disagreement can generate heterogeneous valuations without generating riskless profit opportunities.

The economic implication is straightforward. If one agent's disagreement tilt implies a valuation above the benchmark, that agent has a motive to buy; if another agent's tilt implies a valuation below the benchmark, that agent has a motive to sell. What disagreement generates, therefore, is not arbitrage, but cross-sectional demand dispersion. In this sense, valuation wedges drive trade and volume, while

observed prices remain disciplined by the common benchmark.

This mechanism helps explain why disagreement can be associated with active trading even in the absence of obvious mispricing. Agents need not disagree about the existence of a market price; they disagree about the relative valuation of future states around that price. The market-clearing price remains unique, but the willingness to trade around it reflects heterogeneous beliefs. This interpretation is consistent with empirical evidence linking disagreement to higher turnover, wider bid–ask spreads, and persistent trading activity during periods of interpretive divergence (Hong and Stein, 2007).

The broader lesson is that disagreement should be understood as a source of market activity, not as a violation of arbitrage-free pricing. The κ -tilted model preserves the benchmark discipline of CRR while providing a tractable way to study how heterogeneous beliefs generate valuation dispersion, trade, and liquidity effects.

In the next section, this same logic is carried into a broader asset-pricing setting, where disagreement no longer affects only option valuation in a lattice, but expected returns through systematic exposure to disagreement-intensive states.

10.5 Concluding Remarks

The κ -tilted model extends the classical Cox–Ross–Rubinstein framework by allowing disagreement to distort the relative valuation of states around a common benchmark. In this sense, it does not replace CRR; it generalizes it. When disagreement vanishes, the model collapses exactly to the standard benchmark case.

Its economic contribution is equally clear. Disagreement does not generate arbitrage, because all valuations remain anchored to the same no-arbitrage pricing structure. What disagreement changes is the subjective weighting of states around that benchmark. As a result, agents may rationally arrive at different valuations of the same claim even though the benchmark pricing logic is unchanged.

This distinction is important for understanding market activity. The model explains how heterogeneous beliefs can generate trade, turnover, and liquidity without requiring mispricing or violations of the law of one price. Disagreement is therefore not a breakdown of rational pricing, but a structural source of valuation dispersion within an arbitrage-free framework.

11 The Information CAPM

The traditional Capital Asset Pricing Model (CAPM), developed by Sharpe (1964), Lintner (1965), Mossin (1966), and later formalized in related work by Treynor (1999), provides a parsimonious equilibrium relation between systematic risk and expected return. Its central implication is that only non-diversifiable risk should command a premium in equilibrium. This result, however, is obtained under strong assumptions, among them homogeneous expectations and a common assessment of the probability structure governing future payoffs.

Subsequent extensions, most notably the consumption-based models of Lucas Jr (1978) and Breeden (1979), broadened the economic interpretation of systematic risk by allowing asset returns to co-move with aggregate consumption. Yet these models preserve the same basic discipline of representative pricing under a common belief structure. In that sense, they abstract from one of the central themes of this dissertation: agents may disagree not only about payoffs, but about the probability structure used to evaluate them.

This section develops the Information CAPM (InCAPM), an extension of the CAPM in which disagreement enters the stochastic discount factor as an informational wedge. The key idea is that heterogeneous beliefs do not merely introduce noise around equilibrium. When disagreement is state dependent and interacts with the benchmark market channel, it becomes a priced source of systematic risk. Expected returns may then depend not only on exposure to the market factor, but also on exposure to disagreement-intensive states.

The logic is the same as in the CRR extension of the previous section. There, disagreement distorted subjective valuations around the common no-arbitrage benchmark of the binomial model. Here, the benchmark is no longer the risk-free branch-by-branch anchor, but the market return channel itself, as in the classical CAPM. Disagreement operates around that benchmark in exactly the same structural way, although the derivation is different and the resulting pricing

relation is expressed in equilibrium form rather than through a binomial tree.

The section proceeds as follows. The next subsection introduces the disagreement wedge, the dual-risk stochastic discount factor, and the InCAPM pricing equation itself. The following subsections discuss the economic interpretation of the model, illustrate its implications numerically, and show how the classical benchmark is recovered when disagreement vanishes or becomes state invariant.

11.1 The InCAPM: Disagreement and Informational Risk

This subsection summarizes the assumptions, main equations, and pricing interpretation of the Information CAPM (InCAPM). The full derivation is presented in Appendix B.2. The objective here is to retain the essential algebraic steps while keeping the economic mechanism transparent.

Key differential assumptions.

- Risk-averse investors form subjective beliefs about future states, and these beliefs enter asset pricing through the stochastic discount factor.
- Disagreement across agents is represented as an informational wedge rather than treated as exogenous noise.
- Asset prices remain anchored to a common benchmark market channel, but expected returns may also depend on exposure to disagreement-intensive states.

The last assumption is the crucial extension. In the classical CAPM, the market return is the relevant systematic benchmark. In the InCAPM, that benchmark remains in place, but an additional informational channel appears whenever disagreement is state dependent.

We begin with the disagreement wedge already derived in Paper 1. At horizon T , define

$$\kappa_{t,T}(s) := \frac{\tau_i(t)}{u_i^*(s, T)} - \frac{\tau_j(t)}{u_j^*(s, T)}, \quad (61)$$

where $\tau_i(t)$ and $\tau_j(t)$ are the agents' information-cost multipliers, and $u_i^*(s, T)$ and $u_j^*(s, T)$ are the corresponding state-contingent utility objects. Thus, $\kappa_{t,T}(s)$ measures the relative informational distortion between the two agents across states.

The stochastic discount factor of the InCAPM contains *two* sources of risk:

$$m_{t,T}(s) = \left(\frac{\delta_j}{\delta_i} \right) \exp(\kappa_{t,T}(s)) \exp(-\tau_M \rho_M(s)), \quad (62)$$

where $\rho_M(s)$ denotes the market return channel and τ_M is its associated price-of-risk parameter. The first exponential term, $\exp(\kappa_{t,T}(s))$, is the informational channel; the second, $\exp(-\tau_M \rho_M(s))$, is the benchmark market channel. This is the exact analogue of the CRR section: the benchmark changes, but the principle does not. There, disagreement acted around the CRR benchmark measure; here, it acts around the market return benchmark.

Accordingly, the Euler equation for any asset A takes the form

$$1 = \mathbb{E}_t^{\bar{\nu}} [m_{t,T}(s) R_{t,T}^A(s)] = \mathbb{E}_t^{\bar{\nu}} \left[\left(\frac{\delta_j}{\delta_i} \right) \exp(\kappa_{t,T}(s) - \tau_M \rho_M(s)) R_{t,T}^A(s) \right], \quad (63)$$

where $R_{t,T}^A(s)$ is the return of the asset being priced and expectations are taken under the common benchmark measure $\bar{\nu}$. Equation (63) makes the dual structure of the model explicit: pricing depends jointly on the benchmark market factor and on the informational wedge.

For convenience, isolate the informational component of the kernel:

$$m_t^\kappa(s) := \left(\frac{\delta_j}{\delta_i} \right) \exp(\kappa_{t,T}(s)). \quad (64)$$

It is useful to define the disagreement-linked benchmark payoff

$$\rho_\kappa(s) := -\frac{1}{\tau_M} \left[\log \left(\frac{\delta_j}{\delta_i} \right) + \kappa_{t,T}(s) \right], \quad (65)$$

so that

$$m_t^\kappa(s) = e^{-\tau_M \rho_\kappa(s)}. \quad (66)$$

The object $\rho_\kappa(s)$ is therefore the payoff collinear with the informational component of the stochastic discount factor.

Starting from equation (63), one can apply the covariance decomposition

$$\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + \text{Cov}(X, Y)$$

to isolate the informational term. In valuation space, let τ_A denote the sensitivity parameter associated with asset A 's payoff channel. Then the InCAPM pricing equation can be written as (see Appendix B.2)

$$\boxed{\mathbb{E}_{\bar{\nu}} \left[e^{-\tau_A R_{t,T}^A(s)} \right] = e^{-\tau_A \rho_f} + \beta_{\kappa,A} \left[\mathbb{E}_{\bar{\nu}} \left(e^{-\tau_A \rho_\kappa(s)} \right) - e^{-\tau_A \rho_f} \right]} \quad (67)$$

Equation (67) is the Information CAPM itself. The term $e^{-\tau_A R_{t,T}^A(s)}$ is the asset-specific valuation term, while $e^{-\tau_A \rho_f}$ is its risk-free benchmark counterpart.⁴ Together with the informational beta $\beta_{\kappa,A}$, these terms mark the departure from the classical CAPM: valuation depends not only on the benchmark market channel, but also on the asset's loading on disagreement-intensive states.

Equation (67) is the informational analogue of the Security Market Line. The term on the right-hand side is an *informational premium*, generated by disagreement rather than by exposure to the benchmark market channel alone. The coefficient $\beta_{\kappa,A}$ measures how strongly asset A loads on disagreement-intensive states.

If disagreement vanishes, or becomes constant across states, then $m_t^\kappa(s)$ becomes state-invariant. In that case,

$$\text{Var}(m_t^\kappa) = 0,$$

so the informational premium collapses. The model therefore reduces to its classical benchmark case: the market channel remains, but informational risk disappears.

⁴The parameter τ_A does not represent disagreement. It belongs to the asset payoff channel and therefore appears in both the risky valuation term $e^{-\tau_A R_{t,T}^A(s)}$ and its risk-free counterpart $e^{-\tau_A \rho_f}$. Informational disagreement enters separately through the state-dependent wedge $e^{\kappa_{t,T}(s)}$.

The informational beta is defined by

$$\beta_{\kappa,A} := \frac{\text{Cov}_t^{\bar{v}}\left(\left(\frac{\delta_j}{\delta_i}\right) e^{\kappa_{t,T}(s)}, e^{-\tau_A R_{t,T}^A(s)}\right)}{\text{Var}_t^{\bar{v}}\left(\left(\frac{\delta_j}{\delta_i}\right) e^{\kappa_{t,T}(s)}\right)}. \quad (68)$$

Since δ_j/δ_i is constant across states, it cancels from the covariance ratio. Hence the same beta can be written more simply as

$$\beta_{\kappa,A} := \frac{\text{Cov}_t^{\bar{v}}\left(e^{\kappa_{t,T}(s)}, e^{-\tau_A R_{t,T}^A(s)}\right)}{\text{Var}_t^{\bar{v}}\left(e^{\kappa_{t,T}(s)}\right)}. \quad (69)$$

The full derivation of equations (62)–(69) is presented in Appendix B.2

This beta captures the asset’s exposure to informational disagreement. A high value of $\beta_{\kappa,A}$ indicates that the asset’s valuation is especially sensitive to disagreement across states, and therefore that its expected return is more strongly affected by informational risk.

The covariance entering $\beta_{\kappa,A}$ measures the extent to which asset A ’s valuation co-moves with disagreement across states. A positive covariance means that the asset is highly valued precisely in disagreement-intensive states, so it is exposed to informational risk and commands an informational premium. A negative covariance means that the asset tends to perform well when disagreement is low, so it provides a hedge against informational risk. For example, an asset whose payoff is especially sensitive to episodes of strong divergence in beliefs—such as a speculative technology stock or an option on a highly controversial underlying—would naturally display a positive covariance with the disagreement wedge.

The economic interpretation follows directly from equations (62) and (67). The market return channel $\exp(-\tau_M \rho_M(s))$ plays the benchmark role that, in the CRR section, was played by the binomial pricing anchor. Disagreement does not replace that benchmark. It modifies valuation around it. Thus, the InCAPM applies the same principle as the CRR extension, but with a different benchmark object and a different derivation method: in CRR the anchor was the no-arbitrage binomial measure, while here the anchor is the market return channel of the CAPM

environment.




Figure 6: Informational Security Market Line (InSML). The intercept at R_f is the benchmark anchor of the pricing relation. As disagreement becomes stronger and more state dependent, the price of informational risk rises and the line steepens, while the informational beta $\beta_{\kappa,A}$ determines each asset's position along the line.

Figure 6 provides a schematic representation of the informational component of the pricing relation. The intercept at R_f should be read as the benchmark anchor of the Security Market Line, not as implying that all assets become risk-free when disagreement vanishes. When $\kappa = 0$, the informational wedge disappears, but the standard market-risk component of the CAPM remains. In that case, the InCAPM adds nothing beyond the benchmark CAPM structure; it does not imply that risky assets cease to be risky.

The slope of the InSML is determined by the price of informational risk, just as the slope of the classical Security Market Line is determined by the market premium. In the present framework, that informational price of risk is governed by

the disagreement wedge and its interaction with the benchmark market channel. The informational beta $\beta_{\kappa,A}$ does not determine the slope; rather, it determines an asset's position along the line by measuring its exposure to disagreement-intensive states.

Consequently, disagreement is priced only when it is state dependent. If $\kappa_{t,T}(s)$ is identically zero, or constant across states, the informational component contributes only a level term and the model collapses back to its classical benchmark. When disagreement varies across states, however, it becomes an additional non-diversifiable source of risk.

This interpretation also has empirical relevance. In particular, it suggests a natural way to think about the well-known puzzle that the observed Security Market Line is often flatter than the classical CAPM predicts. Subsection 5 develops this point and relates it to the empirical literature on disagreement, return differentials, and pricing anomalies.

11.2 From Exogenous Benchmarks to Endogenous Aggregation

A delicate issue in disagreement-based asset pricing is the status of the benchmark distribution against which subjective beliefs are evaluated. In many standard applications, benchmark probabilities or pricing kernels are treated as given, either because they are inferred from observed prices or because they arise from a representative-agent equilibrium. In that sense, beliefs are often modeled as deviations from a benchmark that is already in place.

The present framework adopts a different closure. Once heterogeneous beliefs are taken as primitive, it becomes natural to ask whether the benchmark itself should also be formed endogenously rather than imposed from outside. This is the perspective adopted here. Instead of treating the benchmark as exogenous, we allow it to emerge from the aggregation of heterogeneous beliefs. Formally, if agents hold subjective beliefs $p_i(n)$ over states $n \in \mathcal{N}$, and if w_i denotes the relative

influence of agent i in the market, Appendix B.3 shows that the distribution

$$q^*(n) = \sum_i w_i p_i(n)$$

minimizes the weighted average Kullback–Leibler divergence from individual beliefs to a common benchmark. Under this criterion, the belief-weighted average is the natural aggregate benchmark.

The point is not that all classical models treat prices as literally exogenous. Rather, the point is that once disagreement becomes central, it is useful to work with a benchmark that is itself consistent with heterogeneous belief formation. In this framework, prices are interpreted less as fixed signals to which beliefs passively respond, and more as equilibrium objects reflecting the aggregation of divergent probabilistic views.

This does not create a damaging circularity. At the deepest level, the relevant wedge is not simply between benchmark and beliefs in the abstract, but between *pricing* and *valuation*. The benchmark pricing rule is common, while disagreement generates heterogeneous subjective valuations around that benchmark. The apparent loop is closed by the convergence principle developed in Paper 1: in the absence of new information, heterogeneous beliefs cannot drift indefinitely, so the valuation wedge shrinks over time. In the limit, beliefs converge, valuations align, and the benchmark pricing object is no longer distinct from agents' pricing assessments.

This perspective is important for the InCAPM. Disagreement is not treated as residual noise around an otherwise fixed benchmark. It appears as a systematic wedge between common pricing and heterogeneous valuation, and is therefore capable of affecting equilibrium expected returns. Prices do not reveal consensus in a strong sense; they reflect the disciplined aggregation of heterogeneous beliefs under a common market structure.

11.3 Discussion and Sensitivity Analysis

The following numerical example illustrates how belief heterogeneity can translate into pricing differences under the InCAPM. The purpose is not to estimate the model structurally, but to provide intuition for how disagreement, once mapped into an informational wedge, affects valuation and expected returns.

Setup. Consider an economy with three states of nature (s_1, s_2, s_3) and two agents, η and ν , whose subjective belief vectors are

$$p_\eta = [0.7, 0.2, 0.1], \quad p_\nu = [0.1, 0.2, 0.7].$$

Let the asset payoff channel be

$$\rho = [1.0, 1.5, 3.0].$$

Both agents are assumed to share the same state utility specification,

$$u(s) = -e^{-\rho(s)},$$

so that, in absolute value,

$$|u(s)| = [0.3679, 0.2231, 0.0498].$$

Agent weights in the economy are $w_\eta = 0.45$ and $w_\nu = 0.55$. The aggregate benchmark belief is therefore

$$q^* = w_\eta p_\eta + w_\nu p_\nu = [0.37, 0.20, 0.43].$$

Using this benchmark, the informational distances of the two agents are

$$\tau_\eta = D_{\text{KL}}(p_\eta \| q^*) = 0.3004, \quad \tau_\nu = D_{\text{KL}}(p_\nu \| q^*) = 0.2103.$$

Hence the difference in informational intensity is

$$\tau_\eta - \tau_\nu = 0.0901.$$

Informational wedge and beta. Under the simplifying assumption that the state utility term is common across agents, the disagreement wedge can be written as

$$\kappa(s_n) = \frac{\tau_\eta - \tau_\nu}{|u(s_n)|}.$$

Substituting the numerical values yields

$$\kappa(s_n) = \frac{0.0901}{|u(s_n)|} = [0.2451, 0.4041, 1.8111].$$

Thus disagreement is strongest in the state with the lowest utility weight, which here coincides with the highest payoff state.

Let $x(s_n) = \rho(s_n)$ denote the payoff channel. Using the aggregate benchmark weights q^* , we compute

$$\begin{aligned}\bar{\kappa} &= \sum_n q^*(s_n) \kappa(s_n) = 0.9503, \\ \bar{x} &= \sum_n q^*(s_n) x(s_n) = 1.9600, \\ \text{Cov}(\kappa, x) &= \sum_n q^*(s_n) (\kappa(s_n) - \bar{\kappa})(x(s_n) - \bar{x}) = 0.6857, \\ \text{Var}(\kappa) &= \sum_n q^*(s_n) (\kappa(s_n) - \bar{\kappa})^2 = 0.5623.\end{aligned}$$

The resulting informational beta is therefore

$$\beta_\kappa = \frac{\text{Cov}(\kappa, x)}{\text{Var}(\kappa)} = \boxed{1.2195}.$$

A beta above one indicates that the asset payoff is strongly exposed to disagreement-intensive states. In the present example, the high-payoff state is also the state in which the disagreement wedge is largest, so the covariance is positive

and economically meaningful: the asset loads positively on informational risk.

Expected return from the InCAPM. To build intuition, consider the valuation equation in exponential form:

$$\mathbb{E}[\rho_i] = -\frac{1}{\tau} \log \left[e^{-\tau\rho_f} + \beta_{\kappa,\rho} \left(e^{-\tau\rho_\kappa} - e^{-\tau\rho_f} \right) \right].$$

For illustration, let

$$\rho_f = 1.5, \quad \rho_\kappa = 3.0, \quad \tau = 1, \quad \beta_{\kappa,\rho} = 1.2195.$$

Then

$$e^{-1.5} \approx 0.2231, \quad e^{-3.0} \approx 0.0498,$$

so that

$$\mathbb{E}[\rho_i] = -\log [0.2231 + 1.2195(0.0498 - 0.2231)] = -\log(0.0117) \approx \boxed{4.45}.$$

The numerical value is high because the example is deliberately stylized: the asset loads strongly on the state in which disagreement is most intense. The calculation should therefore be read as an illustration of mechanism rather than as a calibrated empirical prediction. Its purpose is to show that when informational beta is high, disagreement can materially alter valuation relative to the classical benchmark.

The following figure provides a visual representation of the informational pricing relation. Assets with higher informational beta are more exposed to disagreement-intensive states, and therefore their expected returns depart more strongly from the benchmark. In this sense, disagreement shifts the Security Market Line by introducing an additional systematic dimension of risk.

The important point is that the benchmark pricing structure is not abandoned. As in the CRR section, disagreement operates around a common benchmark; it

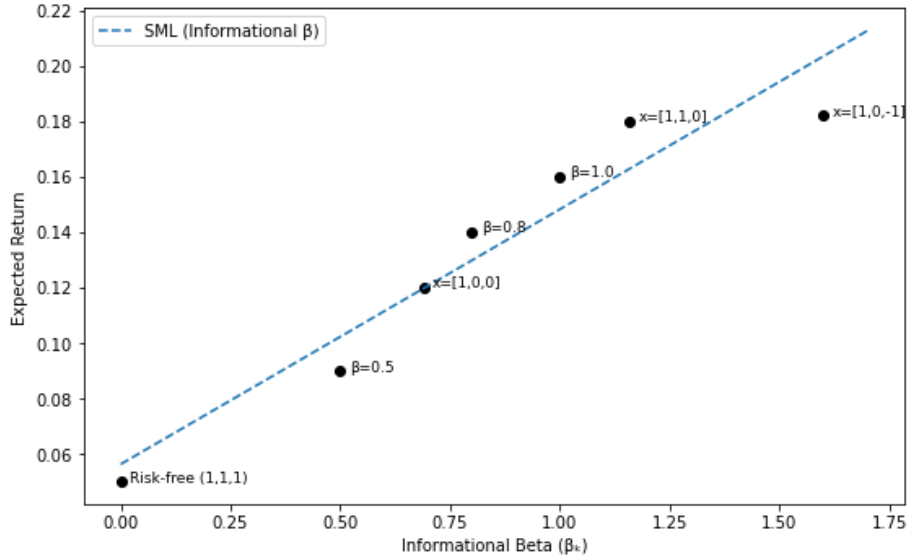


Figure 7: Informational Security Market Line from the stylized numerical example. Each point represents an asset or payoff configuration plotted against its informational beta $\beta_{K,A}$ and expected return. The dashed line is the fitted informational Security Market Line. Assets with higher informational beta are more exposed to disagreement-intensive states and therefore command higher expected returns. The point (1, 1, 1) corresponds to the benchmark risk-free case with zero informational beta, while the remaining portfolios illustrate how heterogeneous state exposure shifts assets along the informational pricing relation.

does not replace it. What changes is the valuation of states around that benchmark. The example therefore illustrates the central mechanism of the InCAPM: belief heterogeneity can affect expected returns through a priced informational wedge, even when the underlying pricing structure remains disciplined by a common market benchmark.

We now turn to the benchmark limiting case in which disagreement disappears. This corresponds to the Arrow–Debreu environment with zero informational beta, where the pricing relation collapses back to the classical benchmark.

11.4 The Arrow–Debreu Benchmark: Full Agreement and Zero Informational Beta

We now consider the limiting case in which all agents agree on probabilities and the asset delivers the same payoff in every state. This is the natural zero-disagreement benchmark of the model. In such an environment, the informational

wedge disappears, the informational beta is zero, and the InCAPM adds nothing beyond the benchmark pricing structure.

Suppose all agents share the same belief distribution, so that

$$p_\eta = p_\nu = q^* = [q_1, q_2, q_3].$$

Let the asset pay identically in every state:

$$R^A = [1, 1, 1].$$

This is the Arrow–Debreu risk-free benchmark, since the return of the asset being priced does not vary across states.

Assume both agents have the same state utility specification. Then the utility term is also constant across states:

$$u(s_n) = -e^{-R^A(s_n)} = -e^{-1} \approx -0.3679.$$

Because beliefs coincide with the common benchmark, the KL-based informational intensities satisfy

$$\tau_\eta = D_{\text{KL}}(p_\eta \| q^*) = 0, \quad \tau_\nu = D_{\text{KL}}(p_\nu \| q^*) = 0.$$

It follows that the disagreement wedge vanishes:

$$\kappa(s_n) = \frac{\tau_\eta}{u^*(s_n)} - \frac{\tau_\nu}{u^*(s_n)} = 0 \quad \text{for all } n.$$

Hence the informational component of the stochastic discount factor becomes constant:

$$m_t^\kappa(s_n) = \left(\frac{\delta_j}{\delta_i} \right) e^{\kappa(s_n)} = \left(\frac{\delta_j}{\delta_i} \right).$$

There is therefore no state-dependent informational distortion. The pricing kernel

contains no disagreement premium, which is exactly the benchmark case studied in Paper 1 under full agreement.

Under the common benchmark belief q^* , the expected payoff is simply

$$\mathbb{E}_{q^*}[R^A] = \sum_n q^*(s_n) R^A(s_n) = 1.$$

If the asset price is $P = 0.945$, the implied gross risk-free return is

$$R_f = \frac{1}{P} \approx 1.058,$$

or equivalently,

$$r_f = R_f - 1 \approx 0.058.$$

Because $\kappa(s_n) = 0$ in every state, the informational covariance term vanishes:

$$\text{Cov}(\kappa, R^A) = 0, \quad \text{Var}(\kappa) = 0.$$

Accordingly, the informational beta is zero:

$$\beta_{\kappa, A} = 0.$$

This benchmark clarifies the limiting logic of the InCAPM. When disagreement is absent, the informational premium disappears because there is no state-dependent wedge to price. The model therefore collapses back to its benchmark pricing structure. This does *not* mean that every risky asset in the economy must earn the risk-free rate. Rather, it means that the InCAPM adds no extra informational component beyond the standard benchmark relation.

In that sense, the figure should be read schematically. The flat line at $\kappa = 0$ represents the disappearance of the *informational* premium, not the disappearance of all other forms of systematic risk. As disagreement becomes positive and state dependent, the informational Security Market Line tilts upward: assets with

higher informational beta command higher expected returns because they are more exposed to disagreement-intensive states. The benchmark case therefore provides the natural reference point from which the pricing effects of informational heterogeneity emerge.

11.5 Empirical Implications

One of the main empirical motivations for the InCAPM is a well-known difficulty of the classical CAPM: the observed Security Market Line (SML) is often flatter than the benchmark model predicts. Rather than treating this as a simple failure of the CAPM, the present framework suggests that part of the gap may reflect variation in the intensity and state dependence of disagreement. When disagreement is weak, absent, or insufficiently state dependent, the informational premium is small and the pricing relation remains close to the flatter empirical benchmark. When disagreement becomes stronger, however, it adds an additional source of priced risk and alters the slope of the relation. In this sense, the model does not imply that zero disagreement drives all returns to the risk-free rate; rather, it shows how disagreement can help explain deviations between the empirical SML and the sharp benchmark predictions of the CAPM.

This perspective is consistent with a broad empirical literature linking disagreement to return differentials and pricing anomalies, whether disagreement is proxied by forecast dispersion, analyst disagreement, or differences in opinion (Diether et al., 2002; Yu, 2011; Banerjee et al., 2009; Boehme et al., 2006; Hong and Stein, 2007). It also helps interpret more recent evidence that the slope of the SML varies substantially across market conditions.

The chart below, based on Wang et al. (An et al., 2024), is useful in this respect. Using CRSP data on NYSE monthly stock returns from January 2000 to December 2022, they show that the slope of the SML varies significantly with the level of active risk-taking in the market. To capture this, they introduce a quantity-based variable q that helps explain differences in return sensitivities to the beta factor

across regimes.

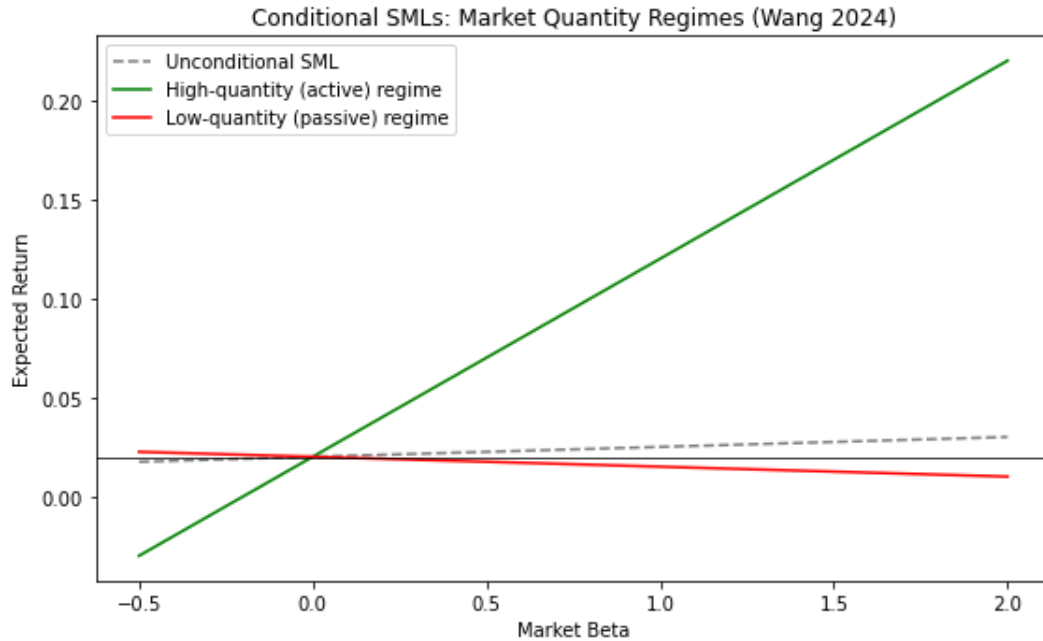


Figure 8: Stylized conditional Security Market Lines inspired by Wang et al. (An et al., 2024). The figure illustrates three market-quantity regimes. The *unconditional* SML is nearly flat, consistent with the familiar empirical finding that the average relation between beta and expected return is weaker than the classical CAPM predicts. The *active* regime refers to high-quantity periods, that is, months in which investors take on more active risk and trading intensity is higher; in these episodes the SML becomes noticeably steeper. The *passive* regime refers to low-quantity periods, in which active risk-taking is muted and market participation is more subdued; in those episodes the SML may become nearly flat or even slightly negative. The chart is schematic and is intended to summarize the qualitative regime variation documented in the data rather than to reproduce estimated coefficients exactly.

From the perspective of the InCAPM, these findings can be interpreted through the lens of disagreement. When disagreement is weak, absent, or not sufficiently state dependent, the informational wedge contributes little to pricing and the SML remains close to the flatter empirical relation. When disagreement becomes stronger, the informational premium increases and the pricing relation steepens. In this sense, the model does not replace the CAPM benchmark; it helps explain why the observed pricing relation may depart from its sharpest predictions.

Three broad empirical regimes can be organized within this framework.

Regime 1: Flat SML (quiet markets or low disagreement). Empirically, the average SML is often found to be nearly flat, a result closely related to the low-risk-premium or zero-beta puzzle. In the InCAPM, this corresponds to periods in which the informational gap κ is small, or in which disagreement is present but not sufficiently state dependent to generate a meaningful premium. In such environments, the informational beta $\beta_{\kappa,A}$ carries little pricing power, and the relation remains close to the flatter observed benchmark.

Regime 2: Steep SML (high activity or strong disagreement). During periods of elevated trading activity or greater cross-sectional divergence in beliefs, the SML steepens. In the present model, this is precisely the situation in which disagreement becomes a priced source of systematic risk. A larger and more state-dependent κ increases the price of informational risk, so assets with higher informational beta command larger premia. The result is a steeper pricing relation, consistent with the idea that disagreement amplifies the compensation required for bearing exposure to disagreement-intensive states.

Regime 3: Inverted or weakly negative SML (asymmetric disagreement). Some empirical episodes reveal a slightly negative SML, particularly in periods of low activity or strong asymmetry in beliefs. The InCAPM can accommodate such cases when disagreement is not only small on average, but skewed across states in a way that changes the sign of the effective informational premium. In that case, assets with lower informational exposure may be relatively favored. This should be interpreted as a conditional possibility of the framework rather than as a universal prediction, but it illustrates the model's ability to organize otherwise puzzling empirical patterns.

Taken together, these regimes suggest that the slope of the Security Market Line need not be constant across time. In the InCAPM, it depends not only on the benchmark market premium, but also on the intensity and state dependence of informational disagreement. This allows the model to reconcile the classical

CAPM benchmark with conditional empirical findings such as those documented in Wang et al. (An et al., 2024).

11.6 Concluding Remarks

The InCAPM extends classical asset pricing by introducing belief heterogeneity as a systematic component of valuation. By relaxing the assumption of homogeneous expectations, it identifies informational disagreement—captured by the wedge κ —as an additional source of priced risk. The resulting informational beta, $\beta_{\kappa,A}$, measures the extent to which an asset is exposed to disagreement-intensive states and therefore provides a disciplined way to connect belief dispersion to expected returns.

The model does not replace the benchmark logic of the CAPM. Rather, it augments it. The market return channel remains the benchmark source of systematic risk, while disagreement adds an informational premium whenever it is sufficiently strong and state dependent. In this sense, the InCAPM applies the same principle developed earlier in the binomial model: disagreement does not generate an autonomous pricing rule, but modifies valuation around a common benchmark.

This perspective is also useful empirically. The framework helps organize well-known deviations from the sharpest CAPM predictions, including the tendency of the observed Security Market Line to appear flatter on average and steeper in periods of heightened market activity or belief dispersion. More broadly, it offers a coherent way to interpret conditional return dynamics in markets where information is fragmented and beliefs do not coincide.

The InCAPM should therefore be understood not as a rejection of classical asset pricing, but as an extension of it to environments in which disagreement is itself an economically relevant and systematically priced force.

12 A Plausible Entropy-Based Resolution to the Ellsberg Paradox

This section revisits the Ellsberg paradox through the lens of information-adjusted utility. We begin by formally presenting the structure of the paradox and surveying leading theoretical responses, including non-additive beliefs, multiple priors, and smooth ambiguity preferences.

We then develop a minimal and tractable resolution based on entropy-penalized utility, showing that ambiguity aversion can arise naturally from a single additional term: the *ambiguity premium* Δ . Unlike traditional approaches, our framework preserves additive probabilities and Bayesian updating, while accounting for second-order uncertainty through an informational channel.

The chapter derives and calibrates the ambiguity premium Δ using a second-order Beta prior, demonstrates its empirical plausibility, and integrates it into the information-sensitive stochastic discount factor (InSDF). The resulting formulation generalizes the classical asset pricing kernel to account for ambiguity in beliefs, yielding intuitive and testable predictions.

12.1 The Ellsberg Paradox and Competing Explanations

The Ellsberg paradox (Ellsberg, 1961) poses a fundamental challenge to Savage's theory of subjective expected utility. In a now-classic experiment, agents systematically prefer to bet on events with known probabilities (for example, a 50/50 urn) rather than on events with ambiguous probabilities (for example, an urn with an unspecified proportion of red and black balls), even when the expected payoff is identical. This behavior contradicts the *Sure-Thing Principle*, according to which preferences should depend only on outcomes and not on the way uncertainty is described. In other words, individuals should maintain consistent preferences regardless of whether the underlying probabilities are known or ambiguous.

The paradox has generated a rich literature exploring how ambiguity—understood as second-order⁵ uncertainty, or more broadly Knightian uncertainty—can be incorporated into decision theory. The most prominent responses include the following.

Non-additive beliefs (Schmeidler, 1989). When agents are unsure how to assign probabilities—that is, under second-order uncertainty—they may represent beliefs by means of non-additive set functions rather than standard additive probabilities. These *capacities* reflect ambiguity rather than risk. Schmeidler models preferences over such non-additive beliefs using the *Choquet integral*, a generalization of expected utility that does not require additivity.

Multiple priors and MaxMin Expected Utility (Gilboa and Schmeidler, 1989). In the MaxMin Expected Utility (MMEU) framework (Gilboa and Schmeidler, 1989), agents do not commit to a single probability distribution. Instead, they consider a set of plausible priors, reflecting ambiguity about the true data-generating process. Ambiguity aversion is captured by a worst-case evaluation rule: the agent computes the minimum expected utility across all admissible priors. MMEU explains the Ellsberg paradox naturally: the ambiguous urn induces a range of plausible beliefs, whereas the risky urn has a unique known prior, making it strictly preferable for ambiguity-averse agents.

Smooth ambiguity preferences (Klibanoff, Marinacci, Mukerji, 2005). This model (Klibanoff et al., 2005) addresses the limitations of traditional subjective probability models, such as Savage’s framework (Savage, 1954), by introducing a two-stage decision process. It separates attitudes toward risk (within a given prior)

⁵A first-order belief refers to a subjective probability distribution $P = (p_1, \dots, p_n)$ that an agent assigns directly to outcomes. Classical expected utility theory, as in Savage (1954), assumes that such a belief exists and is unique, allowing the agent to compute subjective expected utility as

$$\mathbb{E}_P[u(x)] = \sum_i p_i u(x_i).$$

Second-order beliefs reflect uncertainty about the first-order probabilities themselves. Rather than committing to a fixed p , the agent models uncertainty about p as a distribution over distributions. This is particularly appropriate in the Ellsberg setting, where the probability of drawing a red ball, q , is not known, but is instead believed to lie somewhere in $[0, 1]$.

from attitudes toward ambiguity (across priors), allowing a formal representation of ambiguity aversion. Agents act as Bayesian updaters conditional on a prior, but their overall preferences reflect ambiguity aversion through a second-order utility function over expected utilities. This structure makes the model both behaviorally rich and empirically tractable.

Variational preferences (Maccheroni et al., 2006). The framework of Maccheroni et al. (2006) generalizes the MaxMin model by replacing strict worst-case evaluation with smooth, preference-sensitive penalties. Instead of assessing actions solely by their lowest expected utility, agents subtract a cost functional from expected utility under each prior. This cost typically reflects divergence from a reference or most-trusted belief, often measured using relative entropy (KL divergence). Ambiguity aversion is therefore represented by an increasing penalty as beliefs move farther away from the reference prior. The model nests several special cases, including entropy-penalized utility, and provides a flexible and tractable representation of ambiguity-sensitive preferences.

Each of these models relaxes different aspects of standard expected utility theory and introduces increasingly rich mathematical machinery. In contrast, the entropy-adjusted approach developed here preserves the expected utility structure while adding a single, economically interpretable term: an epistemic ambiguity premium associated with unresolved second-order uncertainty. The result is a microfounded and parsimonious account of ambiguity aversion, with testable implications and direct links to asset pricing.

12.2 Ambiguity Premium

Stylized paradox set-up: two urns, two colors. Consider a decision problem involving two urns, each containing white and black balls. The agent is offered a gamble that pays a high payoff x_W if a white ball is drawn, and zero otherwise. The structure of the urns is as follows:

- **Risky urn** (U_R): Contains exactly 50 white and 50 black balls. The probability of drawing white is objectively known:

$$P_R = \left(\frac{1}{2}, \frac{1}{2}\right).$$

- **Ambiguous urn** (U_A): Also contains 100 balls, but the proportion of white is unknown. The agent is told only that the probability of drawing white is some unknown $q \in [0, 1]$. If the agent places a uniform second-order prior on q , then the induced predictive distribution is

$$P_A = \int_0^1 (q, 1 - q) dq = \left(\frac{1}{2}, \frac{1}{2}\right).$$

At the level of first-order predictive probabilities, the two urns therefore look identical. Yet their informational content is fundamentally different. The risky urn U_R involves *risk*: uncertainty over outcomes with known probabilities. The ambiguous urn U_A involves *ambiguity*, or second-order uncertainty: uncertainty not about the outcomes themselves, but about the probability law generating them. The former reflects knowledge of the data-generating process; the latter reflects ignorance about that process.

In the ambiguous urn setup, the decision-maker knows only that the true probability of drawing a white ball, q , lies somewhere in the interval $[0, 1]$, with no reason to privilege any particular value. This is a state of complete ignorance. A natural benchmark in such a setting is provided by the *Maximum Entropy Principle* (Jaynes, 1957), according to which, in the absence of further information, one should adopt the distribution that maximizes entropy subject to the available constraints. In the binary case, entropy is

$$H(p) = -p \log p - (1 - p) \log(1 - p),$$

and it is maximized at $p = \frac{1}{2}$. Hence the predictive distribution associated with

complete ignorance is again

$$P_A = \left(\frac{1}{2}, \frac{1}{2}\right).$$

This point is important. The ambiguous urn does *not* differ from the risky urn in its first-order predictive mean. It differs in the informational structure underlying that prediction. In the risky urn, $(1/2, 1/2)$ is a known and objective probability vector. In the ambiguous urn, $(1/2, 1/2)$ is only the average of a family of possible first-order distributions indexed by the unknown q . The same predictive vector therefore conceals a deeper informational asymmetry.

The entropy-based utility framework developed here captures this distinction through an additional term: the *ambiguity premium* $\Delta > 0$. This premium measures the extra uncertainty associated with unresolved second-order beliefs. It does not arise from the outcomes themselves, which are symmetric in both urns, but from uncertainty about the probability model behind those outcomes. The premium then adjusts utility under ambiguity, either penalizing or rewarding the ambiguous urn depending on the agent's risk preferences. In this way, ambiguity aversion or ambiguity seeking can be generated within an otherwise standard expected utility framework, without abandoning additivity or Bayesian updating.

A key feature of this approach is that the Ellsberg paradox appears as a special case of a more general framework grounded in empirically observed heterogeneity in risk preferences. Ambiguity aversion arises naturally for the large class of risk-averse agents, but the same structure also accommodates ambiguity neutrality and ambiguity seeking. The model therefore does not treat Ellsberg behavior as an isolated anomaly; it embeds it in a broader theory of information-sensitive preferences.

Entropy premium and second-order uncertainty. Because both urns share the same first-order predictive distribution, their ordinary Shannon entropy is the same:

$$H(P_R) = H(P_A) = -\frac{1}{2} \ln \frac{1}{2} - \frac{1}{2} \ln \frac{1}{2} = \ln 2.$$

This equality is not a problem for the theory; it is precisely the source of the paradox. At the first-order level, the risky and ambiguous urns are observationally equivalent. What differs is the unresolved second-order uncertainty behind P_A .

To capture that difference, define the ambiguity premium as the Jensen gap between the entropy of the predictive distribution and the average entropy of the underlying first-order distributions:

$$\Delta := H(\mathbb{E}_q[P(q)]) - \mathbb{E}_q[H(P(q))] > 0. \quad (70)$$

In the Ellsberg setting, $P(q) = (q, 1 - q)$ and the expectation is taken over the second-order prior $q \sim \text{Uniform}[0, 1]$. Since Shannon entropy is strictly concave, Jensen's inequality implies $\Delta > 0$ whenever there is genuine uncertainty about the underlying probability q .

It is therefore useful to distinguish between two objects. The first is the ordinary first-order entropy $H(P_A) = \ln 2$. The second is the *ambiguity-adjusted uncertainty measure*

$$\widetilde{H}(P_A) := H(P_A) + \Delta = \ln 2 + \Delta, \quad (71)$$

which augments first-order entropy by the second-order ambiguity premium. In contrast, the risky urn involves no second-order uncertainty, so

$$\widetilde{H}(P_R) = H(P_R) = \ln 2.$$

This distinction resolves the apparent contradiction. The predictive distribution of the ambiguous urn is symmetric, but the agent still perceives it as more uncertain because the symmetry is generated by averaging over unknown probability laws rather than by knowledge of a fixed one. The ambiguous urn is therefore informationally richer, even though its first-order entropy is unchanged.

The ambiguity premium Δ captures exactly this informational gap. It will be derived and calibrated formally in the next subsections. For now, the crucial point

is conceptual: the extra uncertainty of the ambiguous urn is not a first-order feature of outcomes, but a second-order feature of beliefs about the process generating them.

This distinction explains why risk-averse agents—who dislike entropy and, more generally, unresolved uncertainty⁶—systematically prefer to bet on urn U_R , despite both urns yielding the same first-order expected value. Their preference is driven not by differences in expected utility under known probabilities, but by a utility penalty attached to unresolved model uncertainty. In this sense, the Ellsberg paradox is not generated by different payoffs or even by different first-order beliefs, but by the informational premium attached to ambiguity itself.

The next subsection investigates how this ambiguity premium interacts with preferences and enters utility in a tractable way.

12.3 Utility Comparison and Resolution

The ambiguity premium $\Delta > 0$ is an objective *epistemic* property of the ambiguous urn. It measures the additional uncertainty generated by unresolved second-order beliefs and, as such, does not itself belong to preferences. Preferences remain encoded by the primitive utility function $u(x)$, with the parameter γ governing the agent’s attitude toward risk. What ambiguity changes is the informational environment in which that primitive utility is evaluated.

To keep the notation clear, it is useful to distinguish three objects from the outset. First, $u(x)$ is the Bernoulli utility over outcomes. Second, P denotes the first-order predictive belief under which expected utility is computed. Third, $\widetilde{H}(P)$ denotes the *effective* uncertainty associated with that belief structure. In the present Ellsberg setting, $\widetilde{H}(P)$ coincides with ordinary Shannon entropy for the risky urn, but includes the additional ambiguity premium Δ for the ambiguous urn. Since $\widetilde{H}(P)$ is completely determined by P , it is more transparent in this section to write the information-adjusted utility as a function of P , rather than

⁶The relationship between risk preferences and entropy is explained in (Munoz-Alonso, 2025).

listing the uncertainty term separately as an argument.

Recall from the previous subsection that the risky and ambiguous urns share the same first-order predictive entropy, but differ in their effective uncertainty:

$$\widetilde{H}(P_R) = \ln 2, \quad \widetilde{H}(P_A) = \ln 2 + \Delta, \quad \Delta > 0.$$

The term Δ therefore captures the extra uncertainty generated by second-order ignorance, even when the first-order predictive distributions of the risky and ambiguous urns coincide.

With this notation, we write the information-adjusted utility as

$$\mathbb{U}_\gamma(x; P) := \mathbb{E}_P[u(x)] - \lambda(\gamma) \widetilde{H}(P), \quad (72)$$

where $\mathbb{E}_P[u(x)]$ is baseline expected utility, $\widetilde{H}(P)$ is the ambiguity-adjusted uncertainty measure, and $\lambda(\gamma)$ is a coefficient governing the value or cost of one additional unit of epistemic uncertainty for an agent with risk preference parameter γ .⁷

In what follows, we use the notation $\mathbb{U}_\gamma(x; P)$ throughout, with the subscript γ indicating risk preferences and the argument P denoting the epistemic environment. Equation (72) should therefore be read as follows: preferences are still represented by the primitive utility $u(x)$, while the epistemic structure enters through the belief P and its associated uncertainty measure $\widetilde{H}(P)$. In this sense, entropy remains an informational object rather than a primitive taste argument.

This notation is also preferable here because it avoids confusion with earlier expressions in the thesis where the informational term appeared explicitly as an argument of the utility object. In the present section, that informational term is absorbed into the epistemic argument P , since $\widetilde{H}(P)$ is uniquely determined by the belief structure under consideration. The result is a more compact expression, but the underlying economic interpretation remains the same: utility depends

⁷The notation $\lambda(\gamma)$ here refers to the reduced-form ambiguity weight, not to the KL multipliers used elsewhere in the thesis. Its role is to translate epistemic uncertainty into utility space.

on outcomes through $u(x)$, and on the informational environment through the effective uncertainty attached to beliefs.

The effect of this uncertainty surplus on utility is governed jointly by the magnitude of $\widetilde{H}(P)$ and the sign and size of the coefficient $\lambda(\gamma)$. Appendix A.4 provides the underlying derivation of entropy-augmented utility, whereas Appendix B.4 restates that result in the notation of Part II, thereby linking the general framework to the applied developments of the present paper. The coefficient $\lambda(\gamma)$ bridges first-order attitudes toward risk and second-order attitudes toward beliefs. In the construction developed in Appendix A.4, $\lambda(\gamma)$ emerges as the shadow price of a Kullback–Leibler minimization between a pure-risk density and an uncertainty density based on a best guess about the former. In a binary framework, the former may be represented by a Bernoulli distribution over outcomes, while the latter may be represented by a Beta distribution over the Bernoulli parameter. In that sense, $\lambda(\gamma)$ functions as an *exchange rate*: one extra nat of epistemic uncertainty requires $\lambda(\gamma)$ units of compensating utility.

A higher γ steepens the utility function, amplifies the effect of adverse states, and therefore increases the utility relevance of unresolved uncertainty. In the case of a CRRA utility function, this mechanism yields an effective ambiguity weight of the form

$$\lambda(\gamma) = \frac{\gamma(1-\gamma)}{12\lambda} \bar{x}^{-\gamma-1},$$

which scales the uncertainty term in utility. Because the first- and second-order problems are solved jointly, these objects cannot be wholly independent: the more concave $u(\cdot)$, the larger the marginal utility loss from adverse states, and hence the greater the compensation the agent requires for bearing one extra unit of ambiguity.

For $0 < \gamma < 1$, the coefficient $\lambda(\gamma)$ is positive, so ambiguity enters utility as a penalty. As $\gamma \rightarrow 0$ (risk neutrality), or as the informational constraint becomes negligible, $\lambda(\gamma) \rightarrow 0$, and ambiguity becomes unpriced. For sufficiently negative values, $\lambda(\gamma) < 0$, so ambiguity becomes a reward rather than a penalty. This

simple structure allows the model to accommodate ambiguity aversion, neutrality, and ambiguity seeking within a single expected-utility representation.

Combining equation (72) with

$$\widetilde{H}(P_R) = \ln 2, \quad \widetilde{H}(P_A) = \ln 2 + \Delta,$$

we obtain the utility comparison between the risky and ambiguous urns. Because both urns have the same first-order predictive mean, the expected-utility term $\mathbb{E}_P[u(x)]$ is the same in both cases; only the uncertainty adjustment differs. For a risk-averse individual, ambiguity lowers total utility, so that $\mathbb{U}_\gamma(P_R) > \mathbb{U}_\gamma(P_A)$. For a risk-neutral individual, the penalty vanishes and $\mathbb{U}_\gamma(P_R) = \mathbb{U}_\gamma(P_A)$. For an ambiguity-seeking individual, ambiguity raises utility, leading to $\mathbb{U}_\gamma(P_A) > \mathbb{U}_\gamma(P_R)$.

Table 4: Utility comparison under different ambiguity attitudes. Preference order reverses with the sign of the ambiguity weight $\lambda(\gamma)$.

| Urn | Baseline Expected Utility | Ex- Uncertainty Term | Total Utility \mathbb{U}_γ | Preference |
|--|------------------------------|--------------------------------------|--|-----------------|
| <i>Ambiguity-averse</i> ($\lambda(\gamma) > 0$) | | | | |
| P_R | $\mathbb{E}_P[u(x)]$ | $-\lambda(\gamma) \ln 2$ | $\mathbb{E}_P[u(x)] - \lambda(\gamma) \ln 2$ | $P_R \succ P_A$ |
| P_A | $\mathbb{E}_P[u(x)]$ | $-\lambda(\gamma)(\ln 2 + \Delta)$ | $\mathbb{E}_P[u(x)] - \lambda(\gamma)(\ln 2 + \Delta)$ | |
| <i>Ambiguity-neutral</i> ($\lambda(\gamma) = 0$) | | | | |
| P_R | $\mathbb{E}_P[u(x)]$ | 0 | $\mathbb{E}_P[u(x)]$ | $P_R \sim P_A$ |
| P_A | $\mathbb{E}_P[u(x)]$ | 0 | $\mathbb{E}_P[u(x)]$ | |
| <i>Ambiguity-seeking</i> ($\lambda(\gamma) < 0$) | | | | |
| P_R | $\mathbb{E}_P[u(x)]$ | $+ \lambda(\gamma) \ln 2$ | $\mathbb{E}_P[u(x)] + \lambda(\gamma) \ln 2$ | $P_A \succ P_R$ |
| P_A | $\mathbb{E}_P[u(x)]$ | $+ \lambda(\gamma) (\ln 2 + \Delta)$ | $\mathbb{E}_P[u(x)] + \lambda(\gamma) (\ln 2 + \Delta)$ | |

This framework provides a minimal yet sufficient resolution to the Ellsberg paradox. By introducing a single, interpretable adjustment to classical expected utility—namely, an ambiguity premium scaled by the function $\lambda(\gamma)$ —it captures the empirically observed preference for known risks over ambiguous alternatives while remaining fully within the expected utility paradigm.

It requires no appeal to non-additive probabilities, multiple priors, second-order

utility, or Choquet integrals—just one extra term in the information-adjusted utility index. The resulting structure is parsimonious, microfounded, and consistent with Bayesian reasoning. Most importantly, it makes ambiguity attitudes endogenous to preferences rather than imposing them axiomatically. Ambiguity aversion becomes a special case followed by the majority of agents, but not a universal one. Higher values of the risk aversion parameter γ lead to stronger penalization of ambiguous distributions. In the limit $\gamma \rightarrow 0$, the ambiguity penalty vanishes, restoring standard expected utility.

The framework also sheds light on the failure of the Sure-Thing Principle⁸⁹. In the present model, the violation arises because utility depends on the structure of beliefs, not merely on realized payoffs. The entropy adjustment introduces a penalty that is sensitive to unresolved model uncertainty. As a result, preferences may differ across acts even when their first-order expected utilities coincide. The paradox is therefore captured within a refined expected utility framework that tolerates belief-based deviations without abandoning internal coherence.

Proposition 2 (Entropy-Penalised Expected Utility and Ellsberg). *Let u be strictly increasing and concave, and suppose choices are represented by the information-adjusted utility in equation (72),*

$$\mathbb{U}_\gamma(x; P) = \mathbb{E}_P[u(x)] - \lambda(\gamma) \widetilde{H}(P), \quad \lambda(\gamma) \in \mathbb{R}.$$

Assume further that the ambiguous urn exhibits a strictly larger effective uncertainty level than the risky urn, i.e.

$$\widetilde{H}(P_A) - \widetilde{H}(P_R) = \Delta > 0.$$

⁸⁹Originally formulated in Savage’s *Foundations of Statistics* (1954), the Sure-Thing Principle states that if an agent prefers act f to act g in every state of the world, then the agent should prefer f to g unconditionally.

⁹The Ellsberg paradox reveals a fundamental limitation of the Sure-Thing Principle. In ambiguous settings, preferences depend not only on outcomes, but also on the perceived reliability of the probabilities themselves. The informational structure therefore becomes relevant for choice, violating the independence logic implicit in Savage’s axiom.

Then the agent strictly prefers the risky urn over the ambiguous urn, $P_R \succ P_A$, if and only if $\lambda(\gamma) > 0$.

Conversely, if $\lambda(\gamma) < 0$, then $P_A \succ P_R$. If $\lambda(\gamma) = 0$, the agent is indifferent: $P_A \sim P_R$.

12.4 Calibration of the Ambiguity Premium Δ

The goal of this subsection is twofold. First, we calibrate the ambiguity premium Δ using a concrete probabilistic model, so that the informational cost of ambiguity can be assigned a precise numerical value. Second, we endogenize Δ by deriving it from the structure of second-order uncertainty, rather than postulating it exogenously.

The central idea is that Δ arises naturally when the agent holds a second-order prior over unknown probabilities, as in the Ellsberg urn. In that case, ambiguity is not a primitive taste parameter, but the Jensen gap between the entropy of the predictive distribution and the expected entropy of the underlying first-order distributions. This gap is strictly positive whenever beliefs are genuinely imprecise.

To model second-order uncertainty about binary probabilities, we use the Beta distribution. This choice is natural in the present context. The Beta family accommodates a wide range of priors, from complete ignorance, $\text{Beta}(1, 1)$, to highly concentrated beliefs such as $\text{Beta}(100, 100)$. It also allows the relevant entropy objects to be expressed in closed form using standard digamma identities; see, for example, Cover and Thomas (2006), MacKay (2003), and Appendix B.5 .

In a frequentist Bernoulli experiment, the agent treats the success probability as a known constant $p \in [0, 1]$. The corresponding first-order belief vector is

$$P_R = (p, 1 - p),$$

and uncertainty is purely risk.

To model ambiguity, we instead allow the agent to be uncertain about the true

success probability. Let the unknown parameter Θ be distributed as

$$\Theta \sim \text{Beta}(\alpha, \beta),$$

so that the agent entertains a distribution over Bernoulli laws. The induced predictive first-order distribution is then

$$P_A = \int_0^1 (\theta, 1 - \theta) f_{\text{Beta}}(\theta; \alpha, \beta) d\theta = \left(\frac{\alpha}{\alpha + \beta}, \frac{\beta}{\alpha + \beta} \right).$$

Ambiguity is therefore captured by treating θ as a random variable rather than a known constant. The mean of the Beta prior determines the predictive belief, while the variance determines the degree of second-order uncertainty: the more dispersed the prior, the larger the ambiguity.

The ambiguity premium is defined as the Jensen gap

$$\Delta(\alpha, \beta) = H(P_A) - \mathbb{E}_{\Theta} [H(P(\Theta))] > 0, \quad (73)$$

where $P(\Theta) = (\Theta, 1 - \Theta)$. Equation (73) measures the extra uncertainty generated by unresolved second-order beliefs.

In the special symmetric Ellsberg benchmark, where $\alpha = \beta$, the predictive distribution is

$$P_A = \left(\frac{1}{2}, \frac{1}{2} \right),$$

so that

$$H(P_A) = \ln 2.$$

In this case,

$$\Delta = \ln 2 - \mathbb{E}_{\Theta} [H(P(\Theta))] > 0. \quad (74)$$

Thus, even though the predictive distribution is the same as in the risky urn, the agent still faces additional effective uncertainty because the underlying probability law is unknown.

More generally, for arbitrary (α, β) , the predictive entropy is

$$H(P_A) = -\frac{\alpha}{\alpha + \beta} \log\left(\frac{\alpha}{\alpha + \beta}\right) - \frac{\beta}{\alpha + \beta} \log\left(\frac{\beta}{\alpha + \beta}\right).$$

The expected entropy of the conditional Bernoulli law is

$$\mathbb{E}_\Theta[H(P(\Theta))] = \psi(\alpha + \beta + 1) - \frac{\alpha}{\alpha + \beta} \psi(\alpha + 1) - \frac{\beta}{\alpha + \beta} \psi(\beta + 1), \quad (75)$$

where $\psi(\cdot)$ denotes the digamma function. Substituting into equation (73) yields the closed-form expression

$$\begin{aligned} \Delta(\alpha, \beta) = & -\frac{\alpha}{\alpha + \beta} \log\left(\frac{\alpha}{\alpha + \beta}\right) - \frac{\beta}{\alpha + \beta} \log\left(\frac{\beta}{\alpha + \beta}\right) \\ & - \psi(\alpha + \beta + 1) + \frac{\alpha}{\alpha + \beta} \psi(\alpha + 1) + \frac{\beta}{\alpha + \beta} \psi(\beta + 1). \end{aligned} \quad (76)$$

Equation (76) makes the structure of the ambiguity premium explicit. It depends both on the predictive mean and on the dispersion of the second-order prior. When the prior is diffuse, the Jensen gap is large; when the prior is highly concentrated, the gap becomes small.

The table below reports two benchmark cases. The first corresponds to a uniform prior, Beta(1, 1), which represents maximal ambiguity. The second corresponds to a highly concentrated prior, Beta(100, 100), which approximates near certainty.

Table 5: Ambiguity premium $\Delta(\alpha, \beta)$ under different Beta priors.

| Prior Type | (α, β) | Δ (nats) |
|------------------------------|-------------------|-----------------|
| Uniform (maximum ambiguity) | (1, 1) | 0.193 |
| Informative (near-certainty) | (100, 100) | 0.002 |

These values are economically intuitive. Under Beta(1, 1), the agent faces maximal second-order uncertainty, so the Jensen gap is sizeable. Under Beta(100, 100), the prior is sharply concentrated around 1/2, so ambiguity is nearly absent and Δ becomes negligible.

Figure 9 broadens this comparison by plotting three quantities across several Beta priors: the entropy of the predictive distribution $H(P_A)$, the expected entropy of the underlying Bernoulli laws $\mathbb{E}_\Theta[H(P(\Theta))]$, and the resulting ambiguity premium Δ . The figure makes clear that Δ is largest when the second-order prior is diffuse or U-shaped, and smallest when beliefs are highly concentrated.

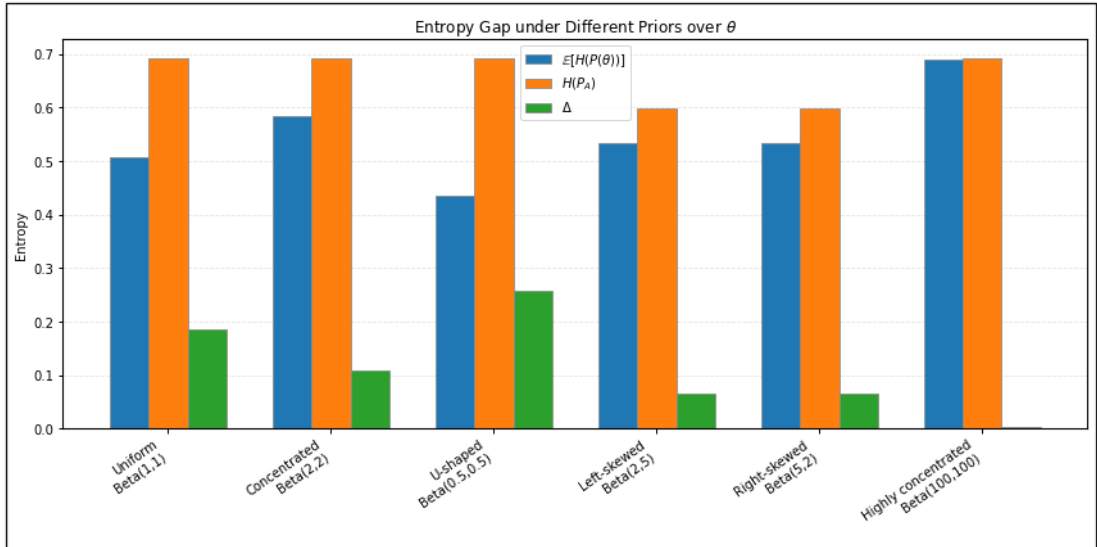


Figure 9: Entropy decomposition under alternative Beta priors. The bars compare the entropy of the predictive distribution $H(P_A)$, the expected entropy of the conditional Bernoulli laws $\mathbb{E}_\Theta[H(P(\Theta))]$, and the ambiguity premium $\Delta = H(P_A) - \mathbb{E}_\Theta[H(P(\Theta))]$. Diffuse or U-shaped priors generate larger ambiguity premia, while highly concentrated priors make Δ nearly vanish.

In a Bayesian updating framework, posterior learning reduces the variance of the Beta prior. As information accumulates and the posterior tightens, one expects

$$\alpha_t + \beta_t \uparrow \quad \implies \quad \Delta_t \downarrow 0.$$

This formalizes the intuition that ambiguity is not permanent: as agents learn more about the underlying probability law, second-order uncertainty diminishes.

In summary, the ambiguity premium Δ emerges endogenously from second-order beliefs modeled as a Beta distribution and can be computed in closed form. This result confirms that ambiguity aversion need not be imposed axiomatically; it can be derived from the informational structure of beliefs itself. The Ellsberg

paradox is thus linked to a measurable epistemic premium rather than to an ad hoc departure from expected utility.

12.5 Ambiguity Premiums in the Information Stochastic Discount Factor

This subsection shows how ambiguity enters the equilibrium pricing kernel through differences in ambiguity premia across agents. The basic idea is simple: if two agents face different levels of second-order uncertainty, the resulting ambiguity premia, Δ_i and Δ_j , generate a wedge in valuation. That wedge can be embedded directly into the stochastic discount factor and compared with the information-based kernel developed in Munoz-Alonso (2025). The result is a unified entropy-based pricing kernel in which first-order disagreement and second-order ambiguity appear as closely related sources of valuation distortion.

To see the connection, recall that the information-based equilibrium kernel in Munoz-Alonso (2025) can be written as

$$m^{ij}(s) = \exp\left(\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)}\right), \quad (77)$$

where τ_i and τ_j denote the informational multipliers of agents i and j , and $u_i^*(s)$, $u_j^*(s)$ are the corresponding state-contingent marginal utility objects. This kernel captures disagreement through a relative tilt in marginal valuation across states.

In the present Ellsberg setting, we introduce disagreement through differences in ambiguity premia. The counterpart ambiguity wedge is

$$\kappa_{ij}^\Delta := \lambda(\gamma)(\Delta_i - \Delta_j), \quad (78)$$

where $\lambda(\gamma)$ is the ambiguity weight introduced in the previous subsection. This term is state independent if Δ_i and Δ_j are scalar agent-level ambiguity premia, but it plays exactly the same formal role as a disagreement wedge in valuation.

Under this identification, the pricing restriction for any asset with gross return R_{t+1}^A may be written as

$$\mathbb{E}_t \left[\frac{\delta_i}{\delta_j} \exp(\kappa_{ij} \Delta) \exp(-\tau_\rho \rho_{t+1}) R_{t+1}^A \right] = 1, \quad (79)$$

where ρ_{t+1} denotes the benchmark payoff or return channel entering the pricing kernel, and τ_ρ is its associated sensitivity parameter. Equation (79) is the ambiguity-augmented counterpart of the information-based Euler equation: the benchmark pricing channel remains in place, while ambiguity adds an exponential valuation wedge around it.

It is useful to define the baseline kernel, in the absence of ambiguity, as

$$m_{ji,t+1}^{(0)} = \frac{\delta_i}{\delta_j} \exp(-\tau_\rho \rho_{t+1}). \quad (80)$$

This is the part of the pricing kernel driven by time discounting and the benchmark return channel alone. Ambiguity then enters multiplicatively through the surcharge $\exp(\kappa_{ij} \Delta)$, so that the full ambiguity-adjusted stochastic discount factor becomes (see Appendix B.6)

$$m_{ji,t+1}^{\text{amb}} = m_{ji,t+1}^{(0)} \exp[\lambda(\gamma)(\Delta_i - \Delta_j)]. \quad (81)$$

Equation (81) makes the role of ambiguity transparent. The kernel contains three components: relative time discounting, the benchmark payoff channel, and an ambiguity premium reflecting the difference in second-order uncertainty across agents. When $\Delta_i = \Delta_j$, the ambiguity wedge vanishes and the model collapses back to the benchmark kernel. When $\tau_\rho \rho_{t+1} = 0$, pricing is driven purely by ambiguity.

To connect this expression with standard utility-based pricing, assume CRRA utility,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}.$$

Then the marginal utility ratio along the benchmark channel can be written as

$$\frac{u'_i(c_{t+1})}{u'_i(c_t)} = \left(\frac{c_{t+1}}{c_t}\right)^{-\gamma} = \exp(-\gamma\rho_{t+1}), \quad (82)$$

where $\rho_{t+1} := \log(c_{t+1}/c_t)$ is the benchmark log return or log consumption-growth channel. Under standard isoelastic preferences, this identifies $\tau_\rho = \gamma$. Substituting into equation (80), the ambiguity-augmented kernel becomes

$$m_{ji,t+1}^{\text{amb}} = \frac{\delta_i}{\delta_j} \exp[\lambda(\gamma)(\Delta_i - \Delta_j)] \exp(-\gamma\rho_{t+1}). \quad (83)$$

The corresponding Euler equation is therefore

$$\mathbb{E}_t \left[\frac{\delta_i}{\delta_j} \exp[\lambda(\gamma)(\Delta_i - \Delta_j)] \exp(-\gamma\rho_{t+1}) R_{t+1}^A \right] = 1. \quad (84)$$

This is the equilibrium pricing restriction implied by ambiguity-adjusted preferences. Ambiguity is priced exactly once, through the exponential wedge $\exp[\lambda(\gamma)(\Delta_i - \Delta_j)]$, while the benchmark payoff channel continues to perform the same role as in the baseline information-sensitive SDF.

This formulation also clarifies the sense in which the ambiguity-based and information-based kernels are equivalent. Comparing equations (77) and (81), the two approaches can be aligned under the identification

$$\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)} \longleftrightarrow \lambda(\gamma)(\Delta_i - \Delta_j). \quad (85)$$

In other words, the ambiguity premium plays the role of a second-order informational wedge. Both kernels preserve the same benchmark pricing structure; they differ only in the informational source of the valuation distortion.

The ambiguity-augmented SDF therefore generalizes the original information-based kernel by incorporating second-order uncertainty explicitly. In the baseline version, valuation differences arise from disagreement over first-order probabilities and their effect on marginal utilities. In the present version, ambiguity premia

Δ_i and Δ_j generate an additional wedge even when first-order expected returns coincide. This makes it possible to price ambiguity directly while remaining within the same equilibrium logic.

Finally, equation (85) suggests a useful interpretation of $\lambda(\gamma)$. It measures the marginal change in pricing tilt induced by a unit change in ambiguity. The larger the ambiguity gap $\Delta_i - \Delta_j$, the larger the wedge in valuation, and hence the larger the compensation required in equilibrium. In this way, second-order uncertainty enters asset pricing through the same structural channel as disagreement: a distortion in relative marginal valuation around a common benchmark.

12.6 Concluding Remarks

The Ellsberg paradox has long stood as one of the clearest challenges to standard expected utility theory, because it reveals that agents care not only about outcomes and first-order probabilities, but also about the informational structure underlying those probabilities. The framework developed in this section provides a tractable resolution by introducing ambiguity as an epistemic premium rather than as a breakdown of rational choice.

The central object is the ambiguity premium Δ , which captures the additional effective uncertainty generated by unresolved second-order beliefs. This premium does not alter primitive preferences directly; instead, it modifies the informational environment in which expected utility is evaluated. Once weighted by the risk-sensitive coefficient $\lambda(\gamma)$, it yields a simple and economically interpretable criterion for ambiguity aversion, neutrality, or seeking. In this way, the Ellsberg paradox appears not as an anomaly requiring a rejection of expected utility, but as a special case of information-adjusted utility under incomplete knowledge.

The section also shows that Δ can be derived endogenously from a probabilistic model of second-order beliefs. Using a Beta prior over unknown Bernoulli parameters, the ambiguity premium emerges as a Jensen gap and can be computed in closed form. This gives the framework empirical discipline: ambiguity is no longer

an abstract preference primitive, but a measurable feature of the informational structure faced by agents.

Finally, embedding ambiguity premia into the information-based stochastic discount factor connects the Ellsberg problem directly to asset pricing. Differences in ambiguity premia across agents generate valuation wedges that are formally analogous to disagreement wedges in the broader thesis. The result is a unified entropy-based pricing kernel in which first-order disagreement and second-order ambiguity are treated within the same structural logic.

Taken together, these results provide a minimal yet sufficient resolution of the Ellsberg paradox. They preserve coherent preferences, additive probabilities, and Bayesian structure, while allowing unresolved uncertainty itself to matter for valuation and choice. In this sense, the contribution of the section is not to abandon standard decision theory, but to refine it so that it can accommodate ambiguity in a disciplined and economically meaningful way.

13 Belief Convergence as a Microfoundation for Informational Efficiency

This section develops a dynamic theory of belief convergence in environments where no new information arrives. Its central contribution is to provide a microfoundation for informationally efficient markets. Rather than assuming efficiency from the outset, as is common in much of asset-pricing theory, we show how efficiency can emerge endogenously as the limiting outcome of rational updating, common pricing discipline, and the progressive erosion of disagreement.

The key idea is simple. When agents stop receiving fresh signals, they can no longer sustain arbitrarily different belief wedges without economic cost. Probability mass is not created; it is only redistributed across states. Under a common benchmark discipline, this process compresses relative distortions, gradually aligns subjective valuations, and drives prices toward their no-arbitrage benchmark. Informational efficiency is therefore not a primitive assumption, but an asymptotic property of a system in which disagreement cannot rationally persist forever in the absence of new information.

The section proceeds in four steps. First, we define disagreement as the benchmark-invariant wedge between agents' beliefs and formalize price-consistent updating. Second, we model the no-signal regime through a common Fokker–Planck discipline that conserves informational mass while contracting relative entropy. Third, we show that this contraction has a direct pricing implication: the informational wedge in the stochastic discount factor disappears asymptotically, so market prices converge toward the benchmark valuation rule. Finally, we study the role of risk aversion and the special case in which the limiting benchmark is also the maximum-entropy distribution.

13.1 Framework and Price-Consistent Belief Dynamics

Disagreement as the dynamic object. Let $\eta_{i,t}$ and $\eta_{j,t}$ denote two agents' subjective beliefs over a common state space \mathcal{S} at date t . Define the benchmark-invariant disagreement wedge

$$\kappa_{ij,t}(s) := \log\left(\frac{\eta_{i,t}(s)}{\eta_{j,t}(s)}\right), \quad m_t^{\eta_i|\eta_j}(s) := \frac{\eta_{i,t}(s)}{\eta_{j,t}(s)} = \exp(\kappa_{ij,t}(s)).$$

Convergence of disagreement means $\kappa_{ij,t}(s) \rightarrow 0$ for all s , equivalently

$$m_t^{\eta_i|\eta_j}(s) \rightarrow 1,$$

so that the relative state-tilt between agents disappears over time. At the pricing level, this is the relevant dynamic object: convergence is not merely the movement of each belief in isolation, but the gradual elimination of the wedge that separates one agent's valuation from another's.

Learning from prices. A convenient microfoundation for belief dynamics is an iterated learning rule in which agent i updates a prior $\eta_{i,t}$ into a posterior $\eta_{i,t+1}$ after observing the market price P_t :

$$\underbrace{\eta_{i,t}}_{\text{prior}} \xrightarrow{\text{observe } P_t} \underbrace{\eta_{i,t+1}}_{\text{posterior}}.$$

A tractable reduced-form representation is the KL-regularized projection

$$\eta_{i,t+1} = \arg \min_{\eta \in \Delta(\mathcal{S})} \{D_{\text{KL}}(\eta \| \eta_{i,t}) + \tau \mathcal{C}(P_t, \eta)\}, \quad \tau > 0, \quad (86)$$

where $\mathcal{C}(P_t, \eta)$ penalizes price inconsistency. A convenient specification is the squared pricing-error cost

$$\mathcal{C}(P_t, \eta) = \frac{1}{2\sigma_P^2} \left(P_t - \hat{P}(\eta) \right)^2, \quad \hat{P}(\eta) := \mathbb{E}_\eta[g(s)], \quad P_t = \hat{P}(\eta) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_P^2). \quad (87)$$

Repeated price-based updating therefore generates a mean-reverting pull toward a price-consistent benchmark. In this sense, prices discipline beliefs: they do not determine them mechanically, but they compress belief configurations that are persistently inconsistent with the common no-arbitrage structure.

Bayesian learning as a finite-signal foundation. The reduced-form step above can be interpreted as a disciplined approximation to Bayesian learning. Let s_t denote a sequence of signals with likelihood $\ell(s_t | s)$, and suppose agent i updates beliefs by Bayes' rule:

$$\eta_{i,t}(s) = \frac{\eta_{i,0}(s) \ell(s_t | s)}{\sum_{s'} \eta_{i,0}(s') \ell(s_t | s')}. \quad (88)$$

If the signal process is asymptotically identifying and all agents observe the same signal sequence, then posterior beliefs converge toward a common benchmark $\bar{\nu}$. If the signal flow stops after some finite date $T < \infty$, Bayesian learning alone no longer moves beliefs after T ; but the no-signal discipline introduced below provides a continuation dynamics that regularizes residual disagreement and pushes beliefs toward a common invariant benchmark.

13.2 No-Signal Dynamics, Conservation, and Disagreement Contraction

A common no-signal discipline. Assume that, after the last relevant signal has arrived, both beliefs evolve under the same Fokker–Planck operator

$$\partial_t \eta_{k,t}(s) = -\nabla \cdot \left(A(s, t) \eta_{k,t}(s) \right) + \frac{1}{2} \nabla^2 \left(B^2(s, t) \eta_{k,t}(s) \right), \quad k \in \{i, j\}, \quad (89)$$

with no-flux boundaries.

Meaning of the notation. Equation (89) describes the evolution of a belief density over the state space. Here $\eta_{k,t}(s)$ is the probability weight assigned by agent k to state s at time t ; $\partial_t \eta_{k,t}(s)$ is its time derivative; $A(s, t)$ is the *drift* term, capturing the systematic directional component of belief revision; and $B^2(s, t)$ is the *diffusion* coefficient, capturing local variance or adjustment noise in the updating process. The operator $\nabla \cdot (\cdot)$ measures net probability outflow from a neighborhood of the state space, while $\nabla^2(\cdot)$ is the Laplacian and captures local spreading of probability mass. In one dimension these reduce to ordinary derivatives with respect to the state coordinate.

The Fokker–Planck equation is therefore a law of motion for beliefs: drift transports probability mass in a systematic direction, while diffusion spreads it locally across nearby states. Because the equation is written in conservation form, it does not create beliefs ex nihilo; it only reallocates existing probability mass across states.

Informational mass conservation. Under no-flux boundaries, each belief conserves total mass:

$$\frac{d}{dt} \int_{\mathcal{S}} \eta_{k,t}(s) ds = 0, \quad k \in \{i, j\}.$$

Economically, this means that updating does not create new probability. In the no-signal regime, learning is therefore a process of redistribution rather than invention: disagreement can be reshaped, transported, and smoothed across states, but it cannot be sustained indefinitely by a fresh source of informational asymmetry. In this precise sense, what is conserved is informational mass.

Continuity form and informational flow. In one dimension, (89) can be written in continuity form

$$\partial_t \eta(s, t) + \partial_s J(s, t) = 0, \quad J(s, t) := A(s, t)\eta(s, t) - \frac{1}{2} \partial_s (B^2(s, t)\eta(s, t)). \quad (90)$$

The current $J(s, t)$ measures the net flow of probability mass across states. The first component, $A(s, t)\eta(s, t)$, is directed transport induced by systematic learning, while the second, $-\frac{1}{2} \partial_s (B^2\eta)$, is diffusive spreading induced by local uncertainty in adjustment. This representation makes the economic mechanism transparent: beliefs evolve by moving probability mass through the state space, not by changing the total amount of probability assigned.

Drift discipline and benchmark convergence. Conservation alone does not guarantee convergence. The crucial economic content lies in the *discipline* embodied in the drift term $A(s, t)$. In the no-signal regime, drift pulls beliefs toward a benchmark-compatible region of the state space, while diffusion governs how noisy or dispersed that adjustment path is. Conservation prevents arbitrary jumps in aggregate probability; drift determines the direction of motion; diffusion determines the roughness of the trajectory.

Proposition 3 (Disagreement contraction under a common no-signal discipline). *Assume the common dynamics (89) are ergodic and admit a unique invariant distribution $\bar{\nu}$. Then*

$$\eta_{i,t} \rightarrow \bar{\nu}, \quad \eta_{j,t} \rightarrow \bar{\nu} \quad \text{as } t \rightarrow \infty,$$

and therefore

$$\kappa_{ij,t}(s) = \log \left(\frac{\eta_{i,t}(s)}{\eta_{j,t}(s)} \right) \rightarrow 0 \quad \text{for all } s \in \mathcal{S}.$$

Equivalently,

$$m_t^{\eta_i | \eta_j}(s) = \exp(\kappa_{ij,t}(s)) \rightarrow 1.$$

Moreover, relative entropy

$$D_t := D_{\text{KL}}(\eta_{i,t} \parallel \eta_{j,t}) = \int_{\mathcal{S}} \eta_{i,t}(s) \log\left(\frac{\eta_{i,t}(s)}{\eta_{j,t}(s)}\right) ds$$

is non-increasing along the common no-signal dynamics.

Reduced-form local decay. Near the invariant benchmark, a convenient reduced-form implication is exponential decay of the disagreement wedge:

$$\kappa_{ij,t}(s) \approx \kappa_{ij,0}(s) e^{-\phi t}, \quad \phi > 0. \quad (91)$$

Equation (91) should be read as a local approximation or benchmark law of motion, not as a primitive theorem. Economically, it captures the idea that state-by-state belief ratios mean revert toward 1 once no new information sustains divergence.

13.3 From Disagreement Convergence to Informationally Efficient Markets

Pricing implication: convergence of relative valuations. Since the endogenous InSDF factorizes as

$$m_t(s) = m_t^C(s) \exp(\kappa_{ij,t}(s)),$$

Proposition 3 implies that, absent new information, the disagreement wedge vanishes asymptotically and pricing collapses to the benchmark kernel $m_t^C(s)$. In that limit, subjective valuations converge toward common no-arbitrage valuations:

$$V_i(A) \rightarrow P(A), \quad V_j(A) \rightarrow P(A).$$

This is the key efficiency result of the section. The efficient-market benchmark is not assumed; it is recovered as the limiting situation in which belief wedges

have been compressed to zero by common updating discipline. In that sense, informationally efficient markets arise from the dynamics of convergence themselves. Prices become efficient not because agents were assumed to agree from the beginning, but because, in the absence of fresh informational asymmetries, they cannot rationally disagree forever.

Why efficiency emerges. The mechanism has three layers. First, common prices discipline beliefs through repeated consistency checks. Second, the no-signal conservation law prevents the indefinite creation of new informational wedges. Third, ergodic drift discipline pulls all surviving beliefs toward the same invariant benchmark. Together, these forces imply that informational distortions are self-compressing. The market therefore approaches an efficiency limit in which prices reflect the common informational benchmark and no persistent relative pricing wedge remains.

Irreversibility and the arrow of time. This gives the process an intrinsic direction. In the absence of fresh signals, persistent disagreement cannot be maintained indefinitely without economic cost. To move away from the benchmark once learning has compressed the wedge, an agent must either absorb expected portfolio losses or sustain positions vulnerable to arbitrage by agents whose beliefs are closer to the common pricing discipline. In that precise sense, convergence is effectively irreversible under no new information. Arbitrage can therefore be understood as a *temporary rebellion against the arrow of time*.

13.4 The Role of Risk Aversion in the Convergence Process

Risk aversion plays a dual and nontrivial role in the dynamics of belief convergence. It influences both the shape of the agent's target belief distribution and the speed at which the agent updates toward that target.

Higher γ flattens optimal beliefs. In Paper 1, higher risk aversion leads to flatter optimal belief distributions and therefore to higher entropy. Risk-averse agents assign greater value to informational neutrality and penalize highly concentrated belief mass. In that sense, they are entropy-seeking: they prefer smoother belief profiles because those profiles reduce exposure to surprise.

Higher γ can also slow adjustment. Within the Fokker–Planck representation, dispersion of the updating path is governed by diffusion B^2 . A natural reduced-form mapping is

$$B^2(s, t) \propto \frac{1}{\gamma}, \quad (92)$$

so that higher γ corresponds to lower diffusion and more cautious updating. Under a common drift discipline, more risk-averse agents may therefore target flatter beliefs while adjusting toward them more slowly. This creates endogenous informational inertia.

The two effects operate in opposite directions: greater risk aversion pushes agents toward higher-entropy beliefs, but may also make their path toward the common benchmark more conservative. This helps explain why agents may agree on the long-run desirability of informational symmetry while reaching it at different speeds.

The simulation below illustrates how disagreement, entropy, and relative pricing evolve over time for agents with different levels of risk aversion. Consistent with the theoretical discussion, higher γ tends to flatten target beliefs but slows the speed of convergence toward the common informational benchmark.

Convergence Dynamics under Informational Stasis

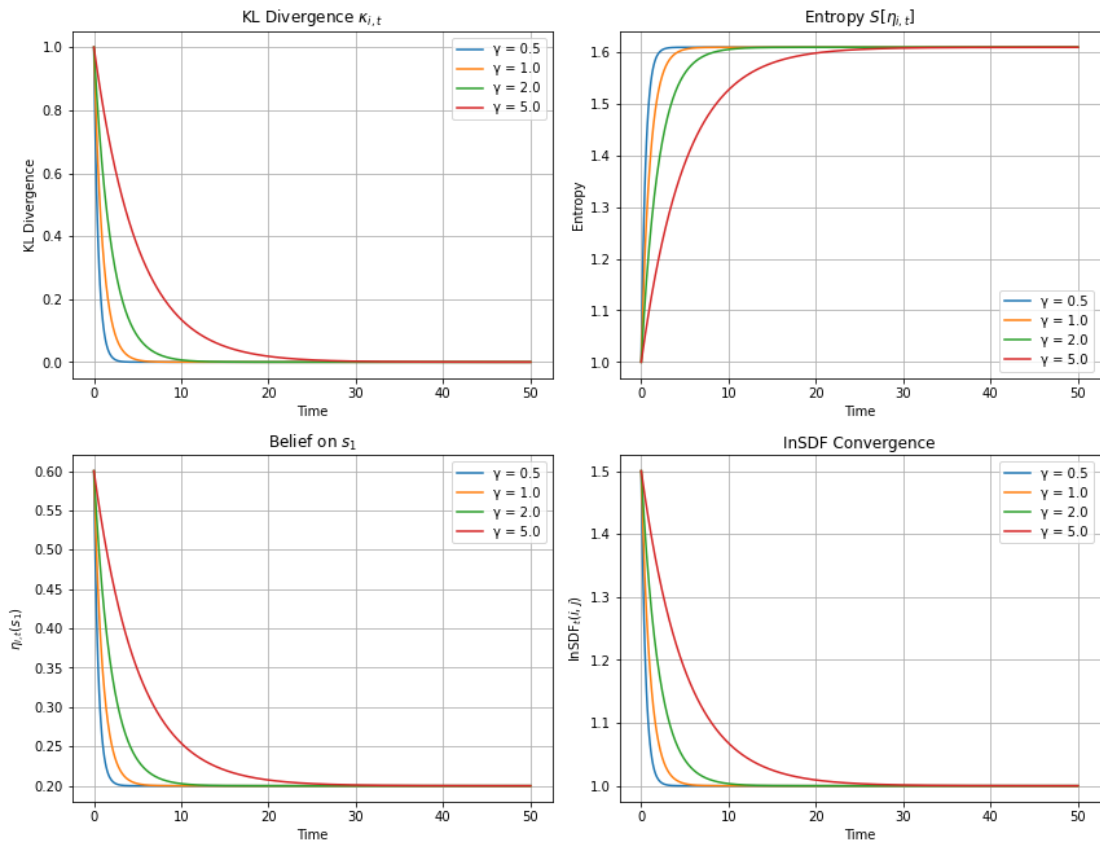


Figure 10: Belief convergence, entropy, and relative pricing under different levels of risk aversion. The figure illustrates how disagreement wedges compress over time, how entropy evolves along the convergence path, and how the relative pricing kernel tends toward unity. Higher risk aversion generates flatter target beliefs but slower adjustment, creating endogenous informational inertia even though convergence eventually occurs.

As a portfolio-level implication of the same convergence mechanism, the next figure shows that entropy-adjusted allocations move progressively toward the benchmark CAPM portfolio as ambiguity premia shrink through learning.

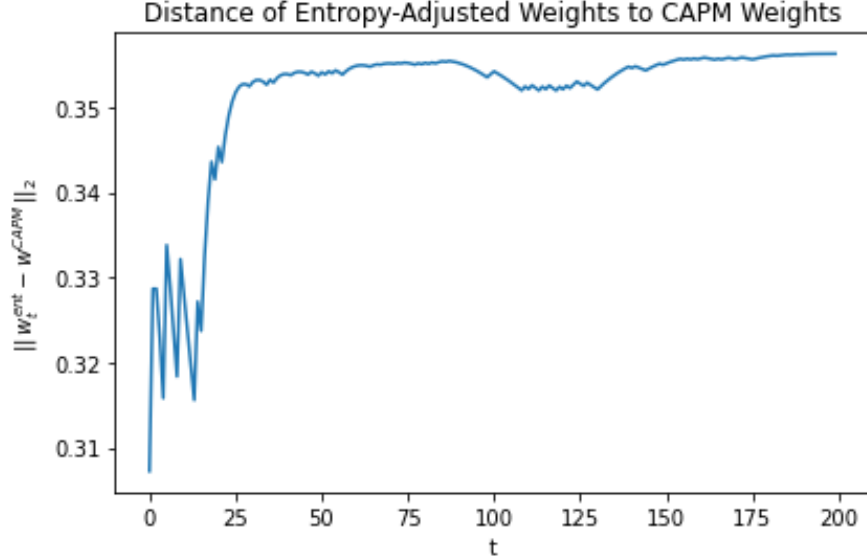


Figure 11: Distance between entropy-adjusted portfolio weights and the CAPM benchmark over time. The horizontal axis t represents simulated periods, while the vertical axis reports the Euclidean distance $\|w_t^{\text{ent}} - w^{\text{CAPM}}\|_2$. Here $w^{\text{CAPM}} = \Sigma^{-1}(\mu/\gamma)$ is the benchmark CAPM allocation and $w_t^{\text{ent}} = \Sigma^{-1}(\mu/\gamma - \frac{1}{2}h_t)$ is the entropy-adjusted allocation implied by the time-varying ambiguity premium Δ_t . The data are simulated: Δ_t is generated through Bayesian updating of a Beta prior over an unknown Bernoulli parameter, while the distortion term h_t is scaled from a Monte Carlo covariance between simulated returns and an informational shock. The figure illustrates how, as learning reduces ambiguity, entropy-adjusted allocations converge toward the benchmark CAPM portfolio.

This result complements the preceding convergence dynamics. The disappearance of informational wedges is reflected not only in beliefs and relative pricing kernels, but also in portfolio choice. As second-order uncertainty declines, the entropy-adjusted allocation converges toward the standard CAPM benchmark, reinforcing the interpretation of market efficiency as an endogenous limiting outcome of the model.

13.5 Symmetry, Maximum Entropy, and the Special Uniform Benchmark

The general convergence result above delivers convergence to the invariant benchmark $\bar{\nu}$. A special and economically important case arises when no signal breaks

symmetry across states and the benchmark itself is uniform:

$$\bar{\nu}(s) = \frac{1}{|\mathcal{S}|}, \quad H(\bar{\nu}) = \log |\mathcal{S}|.$$

In that case, convergence to the benchmark also implies convergence to maximum entropy:

$$H(\eta_{k,t}) \uparrow H(\bar{\nu}).$$

As shown formally in Appendix B.7, in a finite state space the uniform distribution is the unique symmetric belief and the unique maximum-entropy benchmark. Hence, in the no-signal symmetric case, convergence toward the benchmark coincides with convergence toward maximum entropy.

Thus, the uniform benchmark is the symmetric efficiency limit of the no-signal regime. Disagreement disappears, informational premia vanish, and the market converges toward the highest-entropy configuration compatible with the absence of exploitable structure.

More generally, beliefs converge to the unique invariant benchmark $\bar{\nu}$; only when the benchmark itself is uniform does convergence coincide with maximum entropy in the literal sense. This distinction is important. Efficiency does not require uniformity in general. What it requires is convergence toward a common benchmark under which relative informational wedges vanish.

13.6 Summary and Conclusion

In the absence of new information, belief convergence is not accidental. It is the product of deep internal forces: Bayesian learning, price-consistent updating, conservation of informational mass, and common drift discipline. As time passes, informational asymmetries dissolve not because they are assumed away, but because they become increasingly difficult to sustain within a rational pricing system.

The central implication is that informationally efficient markets can be understood as a dynamic limit rather than as a primitive axiom. When fresh informational

asymmetries stop arriving, disagreement wedges contract, subjective valuations align, and prices converge toward the benchmark no-arbitrage valuation rule. In this sense, the section provides a formal route through which efficiency can emerge endogenously from within the belief system itself.

This convergence process reveals that the key conserved quantity in the system is not capital, not probability mass, but *information*. With no arrival of new signals, agents are left to process and refine their existing priors. In doing so, they gravitate toward informational symmetry — the most agnostic, least surprising belief distribution consistent with the constraints they face.

The role of entropy is also clarified. Entropy is not the benchmark in every case; the general benchmark is the invariant distribution $\bar{\nu}$. But in the special symmetric case, the two coincide, and convergence to the benchmark is simultaneously convergence to maximum entropy (see Appendix B.7 for the formal finite-state argument). This gives a precise meaning to the idea that, absent new information, markets move toward informational neutrality and pricing symmetry.

Agents are therefore not forced to agree by assumption. Rather, given enough time and no new signals, they cannot rationally disagree forever. That is where efficiency comes from: not from an arbitrary postulate, but from the endogenous compression of disagreement under a common informational discipline.

14 Final Remarks

This paper has developed a unified asset-pricing framework in which disagreement, ambiguity, and informational structure are treated not as peripheral frictions, but as central determinants of valuation. By extending classical models with entropy-weighted and information-adjusted preferences, the analysis shows that heterogeneity in beliefs is neither mere noise nor a temporary anomaly. It is a systematic and quantifiable force that shapes prices, trading activity, and equilibrium risk premia. The paper's first contribution is to show how disagreement can be embedded tractably into canonical pricing environments. In the binomial setting, the κ -tilted extension of the Cox–Ross–Rubinstein model makes clear that disagreement does not price assets in isolation. Rather, it distorts valuation around a common benchmark measure. This preserves the no-arbitrage logic of the classical model while allowing subjective beliefs to affect state weights, option values, hedging positions, and ultimately market activity. In this way, the framework provides a disciplined explanation for why heterogeneous beliefs may generate trading volume and liquidity even when prices remain internally consistent. The second contribution is the development of the Information CAPM (InCAPM), which extends the benchmark market-risk framework by introducing an informational channel of systematic risk. The informational beta captures an asset's exposure to disagreement-intensive states, while the associated informational premium measures how strongly such exposure is priced in equilibrium. This yields a richer interpretation of the Security Market Line: expected returns depend not only on covariance with the market benchmark, but also on exposure to belief dispersion and informational asymmetry. The model therefore offers a natural way to interpret empirical deviations from the sharpest CAPM predictions, including flatter average SMLs and regime-dependent variation in its slope. The third contribution is to show that ambiguity can be incorporated into the same structural framework. By introducing the ambiguity premium Δ , derived from

second-order uncertainty over first-order probabilities, the paper provides a minimal and coherent resolution of the Ellsberg paradox. Ambiguity aversion emerges not from abandoning expected utility, additive beliefs, or Bayesian structure, but from recognizing that unresolved uncertainty about the probability model itself carries an epistemic cost. When embedded into the stochastic discount factor, differences in ambiguity premia across agents generate valuation wedges that are formally analogous to the disagreement wedges studied elsewhere in the paper. This unifies first-order disagreement and second-order ambiguity within a single entropy-based pricing logic. A further implication of the framework is dynamic. Once convergence is brought into the analysis, informational efficiency no longer needs to be assumed ex ante as a primitive property of markets. Instead, it can be understood as an endogenous limit of the model. In the absence of new information, belief wedges compress over time under common pricing discipline, subjective valuations align, and relative distortions in the pricing kernel disappear. Informationally efficient markets thus emerge as the long-run outcome of disagreement convergence rather than as an imposed axiom. This provides a microfoundation for efficiency within a setting that begins explicitly from heterogeneous beliefs. Taken together, these results offer a broader interpretation of informational risk in financial markets. Prices are shaped not only by fundamentals and intertemporal substitution, but also by the way agents form, distort, and eventually reconcile beliefs. The framework developed here therefore builds a bridge between subjective probability formation and objective market outcomes. It opens several directions for future work, including empirical estimation of informational betas, the study of ambiguity premia in derivative markets, and the role of belief convergence in macro-financial stability. More generally, it suggests that many classical asset-pricing relations can be reinterpreted as limiting cases of a richer informational structure in which disagreement and ambiguity are disciplined, priced, and ultimately resolved through market dynamics.

15 Extensions and Future Work

The framework developed in these two papers opens several natural avenues for future research. The model is intentionally tractable: disagreement, ambiguity, and informational premia are introduced in a way that preserves analytical transparency and close contact with classical asset-pricing benchmarks. Precisely for that reason, the next step is not only to refine the theory further, but also to expand its empirical and computational scope.

A first extension concerns *empirical validation*. The central challenge is to identify observable and economically meaningful proxies for disagreement and ambiguity that can be mapped into the model's informational wedges. A natural starting point is the existing disagreement literature, which has used analyst forecast dispersion, short interest, trading volume around information events, and option-market activity as proxies for belief heterogeneity. Particularly promising for the present framework are option-based measures, since option prices are inherently forward-looking and often reveal asymmetries in beliefs more directly than spot-market variables. Put–call imbalances, option-implied skewness, synthetic long–short disagreement measures, and more recent model-free measures of option-implied belief dispersion appear especially relevant candidates for estimating the informational beta and testing whether disagreement premia are priced in equilibrium.

A second extension concerns *higher-dimensional state spaces and non-closed-form solutions*. The current analysis deliberately focuses on low-dimensional environments in order to preserve closed-form intuition. Yet many of the mechanisms developed here—especially those involving heterogeneous beliefs, ambiguity premia, and convergence dynamics—are likely to be even more relevant in richer settings with many states, many assets, and many interacting agents. In such environments, closed-form derivations may no longer be feasible, and the natural next step is to rely on simulation-based and machine-learning methods. Deep

neural networks, sequence-space methods, and other numerical tools from modern computational economics provide a promising route for solving high-dimensional equilibrium problems in which informational wedges evolve jointly with prices, allocations, and beliefs.

A third extension concerns *market microstructure, liquidity, and segmentation*. One of the appealing implications of the present framework is that disagreement generates trading activity without requiring arbitrage opportunities. This suggests a richer interaction between informational wedges and market structure than has been explored here. Future work could study how disagreement and ambiguity interact with short-sale constraints, dealer balance-sheet capacity, option market segmentation, or liquidity frictions. Such an extension would connect the model more directly to the empirical evidence showing that disagreement often manifests itself through volume, spreads, and differential pricing across trading venues and asset classes.

A fourth extension concerns *dynamic convergence and informational efficiency*. Once convergence is interpreted as an endogenous implication of the model rather than as an external assumption, one can study more systematically how informational efficiency emerges, breaks down, and is restored. This opens the door to dynamic empirical tests: for example, whether disagreement premia compress after major public announcements, whether ambiguity premia decline with learning, and whether the speed of convergence varies systematically with market stress, volatility, or the composition of investors. In this way, the framework could provide a bridge between static asset pricing, belief dynamics, and the broader question of how markets become informationally efficient over time.

Finally, a broader extension concerns *cross-market applications*. While the present paper focuses mainly on equities, options, and stylized portfolio settings, the same informational logic may be relevant in other markets where beliefs are heterogeneous and model uncertainty is central. Credit markets, macro-financial forecasting, and term-structure applications all appear promising in this respect.

More generally, the framework suggests that many observed pricing anomalies may be reinterpreted not as isolated empirical irregularities, but as manifestations of disagreement and ambiguity operating through a common informational channel.

Taken together, these extensions point in the same direction. The contribution of the present paper is to show that disagreement and ambiguity can be incorporated into asset pricing in a disciplined and tractable way. The next stage is to test, scale, and generalize that logic: to measure informational wedges empirically, to solve the model in higher-dimensional settings, and to study how disagreement, ambiguity, and efficiency interact in real markets.

Appendices

A Appendices to Part I: The Model

A.1 Derivation of the Entropy-based SDF

To lighten notation we drop the agent subscript and write $\eta = \eta_i$. Let $S = \{s_1, \dots, s_n\}$ be a finite state space. Fix a *reference benchmark claim* with state log return

$$r_{t+1}^B(s) := \log R_{t+1}^B(s),$$

and let \hat{r}^B denote the external pricing anchor (benchmark requirement).

The problem is

$$\begin{aligned} \arg \min_{\eta(s) \geq 0} \quad & \sum_{s \in S} \eta(s) \ln \left(\frac{\eta(s)}{\bar{\nu}(s)} \right) \\ \text{s.t.} \quad & \sum_{s \in S} \eta(s) = 1, \\ & \sum_{s \in S} \eta(s) r_{t+1}^B(s) = \hat{r}^B, \end{aligned}$$

where $\bar{\nu}$ is the benchmark reference measure.

Introduce Lagrange multipliers $\lambda_0, \lambda_1 \in \mathbb{R}$. The Lagrangian is

$$\mathcal{L} = \sum_{s \in S} \eta(s) \ln \left(\frac{\eta(s)}{\bar{\nu}(s)} \right) + \lambda_0 \left(\sum_{s \in S} \eta(s) - 1 \right) + \lambda_1 \left(\sum_{s \in S} \eta(s) r_{t+1}^B(s) - \hat{r}^B \right).$$

The first-order condition for each s is

$$\frac{\partial \mathcal{L}}{\partial \eta(s)} = \ln \left(\frac{\eta(s)}{\bar{\nu}(s)} \right) + 1 + \lambda_0 + \lambda_1 r_{t+1}^B(s) = 0,$$

hence

$$\eta(s) = \bar{\nu}(s) \exp\left(-1 - \lambda_0 - \lambda_1 r_{t+1}^B(s)\right).$$

Imposing normalization yields

$$1 = \sum_{s \in S} \eta(s) = e^{-1-\lambda_0} \sum_{s \in S} \bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}.$$

Define the partition function

$$Z(\lambda_1) := \sum_{s \in S} \bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}.$$

Then

$$e^{-1-\lambda_0} = \frac{1}{Z(\lambda_1)}.$$

Substituting back, the optimal distorted distribution is

$$\eta(s) = \frac{\bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}.$$

The implied likelihood ratio (Radon–Nikodym derivative) is

$$m^{\eta|\bar{\nu}}(s) = \frac{\eta(s)}{\bar{\nu}(s)} = \frac{e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}.$$

Finally, λ_1 is determined by the anchor constraint $\sum_s \eta(s) r_{t+1}^B(s) = \hat{r}^B$, i.e.

$$\hat{r}^B = \frac{\sum_{s \in S} \bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)} r_{t+1}^B(s)}{Z(\lambda_1)} = \sum_{s \in S} \bar{\nu}_{\lambda_1}(s) r_{t+1}^B(s), \quad \bar{\nu}_{\lambda_1}(s) = \frac{\bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}.$$

Remark (absorbing the normalizer). The optimal belief can be written as

$$\eta(s) = \bar{\nu}(s) \exp\left(-1 - \lambda_0 - \lambda_1 r_{t+1}^B(s)\right),$$

where

$$Z(\lambda_1) = \sum_{u \in S} \bar{\nu}(u) e^{-\lambda_1 r_{t+1}^B(u)} \quad \text{and} \quad e^{-1-\lambda_0} = \frac{1}{Z(\lambda_1)}.$$

Equivalently, defining $\tilde{\lambda}_0 := 1 + \lambda_0 + \log Z(\lambda_1)$, we may write

$$\eta(s) = \bar{\nu}(s) \exp\left(-\tilde{\lambda}_0 - \lambda_1 r_{t+1}^B(s)\right).$$

This is a purely algebraic reparametrization: the probability vector η is unchanged, and identification of λ_1 is unaffected.

A.2 Verification of kernel properties

Let

$$m^{\eta|\bar{\nu}}(s) = \frac{e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}, \quad Z(\lambda_1) = \sum_{u \in S} \bar{\nu}(u) e^{-\lambda_1 r_{t+1}^B(u)}, \quad \eta(s) = \bar{\nu}(s) m^{\eta|\bar{\nu}}(s).$$

1) **Positivity.** Since $e^{-\lambda_1 r_{t+1}^B(s)} > 0$ and $Z(\lambda_1) > 0$, we have $m^{\eta|\bar{\nu}}(s) > 0$ for all s .

2) **Normalization.**

$$\sum_{s \in S} \bar{\nu}(s) m^{\eta|\bar{\nu}}(s) = \frac{1}{Z(\lambda_1)} \sum_{s \in S} \bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)} = \frac{Z(\lambda_1)}{Z(\lambda_1)} = 1.$$

3) **Anchor consistency.**

$$\sum_{s \in S} \bar{\nu}(s) m^{\eta|\bar{\nu}}(s) r_{t+1}^B(s) = \sum_{s \in S} \eta(s) r_{t+1}^B(s) = \hat{r}^B.$$

4) **Monotonicity and uniqueness of λ_1 .** Define

$$f(\lambda_1) := \sum_{s \in S} \bar{\nu}_{\lambda_1}(s) r_{t+1}^B(s), \quad \bar{\nu}_{\lambda_1}(s) := \frac{\bar{\nu}(s) e^{-\lambda_1 r_{t+1}^B(s)}}{Z(\lambda_1)}.$$

Then

$$f'(\lambda_1) = -\text{Var}_{\bar{\nu}_{\lambda_1}}(r_{t+1}^B) \leq 0,$$

with strict inequality whenever r_{t+1}^B is non-degenerate under $\bar{\nu}_{\lambda_1}$. Moreover,

$$\lim_{\lambda_1 \rightarrow +\infty} f(\lambda_1) = \min_{s \in S} r_{t+1}^B(s), \quad \lim_{\lambda_1 \rightarrow -\infty} f(\lambda_1) = \max_{s \in S} r_{t+1}^B(s).$$

Hence for any \hat{r}^B in the convex hull of $\{r_{t+1}^B(s) : s \in S\}$ there exists a unique λ_1 such that $f(\lambda_1) = \hat{r}^B$.

A.3 Decomposition of a disagreement price premium

Let ν be a benchmark measure and η an alternative belief, with RN derivative

$$m^{\eta\nu}(s) := \frac{d\eta}{d\nu}(s) = \frac{\eta(s)}{\nu(s)}.$$

Let $D(s)$ denote the benchmark discounting term used in pricing (e.g. $D(s) = R^{-1}(s)$ or $D(s) = \hat{m}_{t+1}(s)$). Define the benchmark price functional

$$P^\nu(A) := \sum_{s \in S} \nu(s) A(s) D(s), \quad P^\eta(A) := \sum_{s \in S} \eta(s) A(s) D(s).$$

Then the disagreement premium is

$$\begin{aligned} P^\eta(A) - P^\nu(A) &= \sum_{s \in S} (\eta(s) - \nu(s)) A(s) D(s) \\ &= \sum_{s \in S} \nu(s) (m^{\eta\nu}(s) - 1) A(s) D(s). \end{aligned}$$

Since $\sum_{s \in S} \nu(s) m^{\eta\nu}(s) = 1$, we may write the premium as a covariance under ν :

$$P^\eta(A) - P^\nu(A) = \text{Cov}_\nu \left(m^{\eta\nu}(s), A(s) D(s) \right).$$

In particular, assets whose discounted payoffs $A(s)D(s)$ covary positively with the likelihood ratio $m^{\eta\nu}(s)$ command a positive disagreement premium under η relative to ν .

A.4 Derivation of Entropy-Augmented Utility

We derive the entropy-augmented expected-utility representation used in the main text. Let

$$H_\gamma^u(p) := - \sum_s p_s U(w_s, \gamma) \log p_s.$$

The target reduced-form representation is

$$\mathbb{E}_p[\mathbb{U}(w, \gamma, p)] = \mathbb{E}_p[U(w, \gamma)] - \lambda(\gamma) H_\gamma^u(p). \quad (93)$$

Consider the KL-disciplined objective

$$\max_{p \in \Delta} \left\{ \sum_i p(i) u(x_i) - \alpha D_{\text{KL}}(p||q) \right\}, \quad D_{\text{KL}}(p||q) = \sum_i p(i) \log \frac{p(i)}{q(i)}. \quad (94)$$

If $q(i) = 1/n$, then

$$D_{\text{KL}}(p||q) = \log n - H(p), \quad H(p) := - \sum_i p(i) \log p(i),$$

so the objective is equivalent, up to a constant, to

$$\sum_i p(i) u(x_i) + \alpha H(p).$$

The first-order condition gives the exponential tilt

$$p^*(i) \propto q(i) \exp\left(\frac{u(x_i)}{\alpha}\right).$$

For small dispersion around $\bar{u} := \sum_i q(i)u(x_i)$,

$$p^*(i) \approx q(i) \left(1 + \frac{u(x_i) - \bar{u}}{\alpha}\right),$$

which implies the standard second-order approximation

$$\sum_i p^*(i)u(x_i) \approx \bar{u} - \frac{1}{2\alpha} \text{Var}_q(u(x)).$$

Under CRRA utility,

$$u(x) = \frac{x^{1-\gamma} - 1}{1-\gamma}, \quad u''(x) = -\gamma x^{-\gamma-1},$$

and for $\bar{x} := \mathbb{E}_q[x]$,

$$\text{Var}_q(u(x)) \approx (u'(\bar{x}))^2 \text{Var}_q(x), \quad u'(\bar{x}) = \bar{x}^{-\gamma}.$$

Hence

$$\sum_i p^*(i)u(x_i) \approx u(\bar{x}) - \frac{1}{2\alpha} \bar{x}^{-2\gamma} \text{Var}_q(x).$$

Using utility-weighted entropy, a parsimonious sign-consistent mapping is

$$\lambda(\gamma) = \lambda_0 \gamma, \quad \lambda_0 > 0, \tag{95}$$

so entropy is penalized for $\gamma > 0$, neutral at $\gamma = 0$, and rewarded for $\gamma < 0$.

A.5 Derivation of Utility-Maximizing Subjective Probabilities

We derive the optimal belief distribution η_i^* when beliefs are choice variables and preferences enter through an information-weighted objective.

Fix agent i and write $U_i(s) := U_i(w_i(s), \gamma_i)$. Consider

$$\Delta := \left\{ \eta_i : \eta_i(s) \geq 0, \sum_{s \in S} \eta_i(s) = 1 \right\},$$

and the program

$$\max_{\eta_i \in \Delta} \sum_{s \in S} \eta_i(s) U_i(s) \left(1 + \alpha_i \log \eta_i(s) \right), \quad \alpha_i > 0. \tag{96}$$

The Lagrangian is

$$\mathcal{L}(\eta_i, \theta_i) = \sum_{s \in S} \eta_i(s) U_i(s) \left(1 + \alpha_i \log \eta_i(s)\right) - \theta_i \left(\sum_{s \in S} \eta_i(s) - 1\right).$$

The first-order condition is

$$U_i(s) \left(1 + \alpha_i + \alpha_i \log \eta_i(s)\right) = \theta_i. \quad (97)$$

Hence

$$\log \eta_i(s) = \frac{\theta_i}{\alpha_i U_i(s)} - \frac{1 + \alpha_i}{\alpha_i}, \quad (98)$$

so

$$\eta_i(s) = \exp\left(-\frac{1 + \alpha_i}{\alpha_i}\right) \exp\left(\frac{\theta_i/\alpha_i}{U_i(s)}\right). \quad (99)$$

Define

$$\tau_i := \frac{\theta_i}{\alpha_i}.$$

Then

$$\eta_i(s) = \exp\left(-\frac{1 + \alpha_i}{\alpha_i}\right) \exp\left(\frac{\tau_i}{U_i(s)}\right). \quad (100)$$

Imposing normalization and defining

$$Z_i := \sum_{u \in S} \exp\left(\frac{\tau_i}{U_i(u)}\right), \quad (101)$$

we obtain

$$\eta_i^*(s) = \frac{1}{Z_i} \exp\left(\frac{\tau_i}{U_i(w_i(s), \gamma_i)}\right). \quad (102)$$

A.6 Full Derivation of the Endogenous InSDF

This appendix derives the endogenous entropy-based stochastic discount factor (InSDF).

Let

$$\eta_i^*(s) = \frac{\bar{\nu}(s) \exp\left(\frac{\lambda_i}{u_i^*(s)}\right)}{Z_i},$$

$$\eta_j^*(s) = \frac{\bar{\nu}(s) \exp\left(\frac{\lambda_j}{u_j^*(s)}\right)}{Z_j},$$

with

$$Z_k = \sum_{s \in S} \bar{\nu}(s) \exp\left(\frac{\lambda_k}{u_k^*(s)}\right).$$

Define the disagreement wedge

$$m^{\eta_i|\eta_j}(s) := \frac{\eta_i^*(s)}{\eta_j^*(s)} = \frac{Z_j}{Z_i} \exp\left(\frac{\lambda_i}{u_i^*(s)} - \frac{\lambda_j}{u_j^*(s)}\right). \quad (103)$$

Equivalently,

$$\kappa(s) := \log\left(\frac{Z_j}{Z_i}\right) + \frac{\lambda_i}{u_i^*(s)} - \frac{\lambda_j}{u_j^*(s)}, \quad m^{\eta_i|\eta_j}(s) = \exp(\kappa(s)). \quad (104)$$

Introduce consumption growth

$$\rho_{t+1}(s) := \log\left(\frac{C_{t+1}(s)}{C_t}\right), \quad (105)$$

and the benchmark consumption kernel

$$m_{t+1}^C(s) := \delta \exp\left(-\gamma \rho_{t+1}(s)\right). \quad (106)$$

Then the endogenous InSDF is

$$m_{t+1}(s) := m_{t+1}^C(s) m^{\eta_i|\eta_j}(s) = \delta \exp\left(-\gamma \rho_{t+1}(s)\right) \exp(\kappa(s)). \quad (107)$$

For any traded asset A with gross return $R_{t+1}^A(s)$, the Euler restriction is

$$\mathbb{E}_t^{\bar{\nu}}[m_{t+1}(s) R_{t+1}^A(s)] = 1. \quad (108)$$

Equivalently,

$$\mathbb{E}_t^{\bar{\nu}}[\delta \exp(-\gamma \rho_{t+1}(s)) \exp(\kappa(s)) R_{t+1}^A(s)] = 1. \quad (109)$$

If $\eta_i^* = \eta_j^*$, then $\kappa(s) \equiv 0$ and the restriction collapses to the standard benchmark consumption Euler equation.

B Appendices to Part II: Applications

B.1 From the Common Reference Measure $\bar{\nu}$ to the κ -Tilted Binomial Measure and SDF

This appendix derives the κ -tilted binomial pricing measure used in the main text and shows its associated stochastic discount factor representation. The logic follows the same sequence as in Paper 1 (see Munoz-Alonso (2025)). We begin with the common reference measure $\bar{\nu}$, project it onto the no-arbitrage set to obtain the benchmark binomial pricing measure ν^* , and then introduce disagreement as a relative state-dependent tilt around that benchmark.

1. Optimal beliefs relative to the common reference measure $\bar{\nu}$. In Paper 1, agent i 's optimal subjective belief over a finite state space S is derived relative to the common reference measure $\bar{\nu}$ and takes the form

$$\eta_i^*(s) = \frac{\bar{\nu}(s) \exp(\tau_i/U_i(s))}{\sum_{s' \in S} \bar{\nu}(s') \exp(\tau_i/U_i(s'))}, \quad s \in S, \quad (110)$$

while for agent j ,

$$\eta_j^*(s) = \frac{\bar{\nu}(s) \exp(\tau_j/U_j(s))}{\sum_{s' \in S} \bar{\nu}(s') \exp(\tau_j/U_j(s'))}. \quad (111)$$

Here $\bar{\nu}$ is the common benchmark belief, $U_i(s)$ and $U_j(s)$ are the agents' state-contingent utility objects, and τ_i, τ_j are the corresponding information-cost multipliers.

Taking the ratio of (110) and (111) yields the relative likelihood ratio

$$\frac{\eta_i^*(s)}{\eta_j^*(s)} = \exp(\kappa_{ij}(s)), \quad (112)$$

where

$$\kappa_{ij}(s) := \log\left(\frac{Z_j}{Z_i}\right) + \frac{\tau_i}{U_i(s)} - \frac{\tau_j}{U_j(s)}, \quad Z_\ell := \sum_{s' \in S} \bar{\nu}(s') \exp(\tau_\ell/U_\ell(s')). \quad (113)$$

2. From $\bar{\nu}$ to the no-arbitrage benchmark ν^* . In the one-period binomial model, the state space is $S = \{u, d\}$, with gross risky returns

$$\rho(u) = u, \quad \rho(d) = d, \quad u > d > 0,$$

and gross risk-free return ρ_f . The benchmark pricing measure ν^* is obtained by projecting $\bar{\nu}$ onto the set of measures satisfying the CRR martingale restriction:

$$\nu^* = \arg \min_{\nu} \sum_{s \in \{u, d\}} \nu(s) \log\left(\frac{\nu(s)}{\bar{\nu}(s)}\right) \quad \text{s.t.} \quad \sum_s \nu(s) = 1, \quad \sum_s \nu(s) \rho(s) = \rho_f. \quad (114)$$

The solution is the exponential tilt

$$\nu^*(s) = \frac{\bar{\nu}(s) e^{-\lambda \rho(s)}}{\sum_{s' \in \{u, d\}} \bar{\nu}(s') e^{-\lambda \rho(s')}}. \quad (115)$$

In the two-state CRR setting, this projection reproduces the standard benchmark probability

$$q^* = \frac{\rho_f - d}{u - d}, \quad \nu^*(u) = q^*, \quad \nu^*(d) = 1 - q^*. \quad (116)$$

3. Disagreement as a tilt around ν^* . Once the benchmark pricing measure ν^* is fixed, disagreement between agents i and j induces the tilted measure

$$\nu_{ij}^\kappa(s) = \frac{\nu^*(s)e^{\kappa_{ij}(s)}}{\sum_{s' \in \{u,d\}} \nu^*(s')e^{\kappa_{ij}(s')}}. \quad (117)$$

4. The disagreement-adjusted up-state probability. Applying (117) to the up state gives

$$q_{\kappa,ij} = \nu_{ij}^\kappa(u) = \frac{q^* e^{\kappa_{ij}(u)}}{q^* e^{\kappa_{ij}(u)} + (1 - q^*) e^{\kappa_{ij}(d)}}. \quad (118)$$

Substituting (113) into (118), the constant term $\log(Z_j/Z_i)$ cancels under normalization. Therefore

$$q_{\kappa,ij} = \frac{q^* \exp\left(\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)}\right)}{q^* \exp\left(\frac{\tau_i}{U_i(u)} - \frac{\tau_j}{U_j(u)}\right) + (1 - q^*) \exp\left(\frac{\tau_i}{U_i(d)} - \frac{\tau_j}{U_j(d)}\right)}. \quad (119)$$

5. Reduced-form scalar tilt. For numerical work and figures, it is convenient to summarize the state dependence in a scalar parameter κ . A simple specialization is

$$\kappa_{ij}(u) = \kappa \ln u, \quad \kappa_{ij}(d) = \kappa \ln d. \quad (120)$$

Substituting (120) into (118) gives

$$q_\kappa = \frac{q^* e^{\kappa \ln u}}{q^* e^{\kappa \ln u} + (1 - q^*) e^{\kappa \ln d}} = \frac{q^* u^\kappa}{q^* u^\kappa + (1 - q^*) d^\kappa}. \quad (121)$$

6. The associated stochastic discount factor. The disagreement-adjusted stochastic discount factor is the benchmark CRR kernel multiplied by the likelihood-ratio wedge:

$$m_{ij}^\kappa(s) = \frac{1}{\rho_f} \frac{d\nu_{ij}^\kappa}{d\nu^*}(s) = \frac{1}{\rho_f} \frac{e^{\kappa_{ij}(s)}}{\sum_{s' \in \{u,d\}} \nu^*(s')e^{\kappa_{ij}(s')}}. \quad (122)$$

Therefore any one-period payoff X_1 is priced according to

$$P_0(X) = \mathbb{E}_{\nu^*} [m_{ij}^\kappa(s) X_1(s)] = \frac{1}{\rho_f} \mathbb{E}_{\nu_{ij}^\kappa} [X_1]. \quad (123)$$

7. Nesting of the CRR benchmark. If disagreement vanishes, or more generally if it is identical across states, then

$$\kappa_{ij}(u) = \kappa_{ij}(d),$$

and equations (118)–(119) collapse to

$$q_{\kappa,ij} = q^*.$$

Hence $\nu_{ij}^{\kappa} = \nu^*$ and $m_{ij}^{\kappa}(s) = 1/\rho_f$.

B.2 Derivation of the Information CAPM (InCAPM)

This appendix derives the Information CAPM (InCAPM) reported in Section 11. The logic is the same as in the CRR extension of the previous section, but the benchmark object is now different. In the CRR model, disagreement distorted valuation around the no-arbitrage binomial benchmark. Here, disagreement distorts valuation around the benchmark market return channel.

Step 1: The full stochastic discount factor has two sources of risk. The InCAPM stochastic discount factor is

$$m_{t,T}(s) = \left(\frac{\delta_j}{\delta_i} \right) \exp(\kappa_{t,T}(s)) \exp(-\tau_M \rho_M(s)), \quad (124)$$

Step 2: Euler equation for the asset being priced. For any asset A , with state-contingent return $R_{t,T}^A(s)$, the Euler equation is

$$1 = \mathbb{E}_t^{\bar{\nu}} [m_{t,T}(s) R_{t,T}^A(s)]. \quad (125)$$

Step 3: Isolating the informational component. Define

$$m_t^{\kappa}(s) := \left(\frac{\delta_j}{\delta_i} \right) \exp(\kappa_{t,T}(s)). \quad (126)$$

and

$$Y_A(s) := e^{-\tau_A R_{t,T}^A(s)}, \quad Y_f := e^{-\tau_A \rho_f}. \quad (127)$$

The informational pricing condition is

$$\mathbb{E}_t^{\bar{\nu}}[m_t^\kappa(s) Y_A(s)] = 1. \quad (128)$$

Step 4: Covariance decomposition. Applying covariance decomposition yields

$$\mathbb{E}[m_t^\kappa] \mathbb{E}[e^{-\tau_A R_{t,T}^A(s)}] + \text{Cov}(m_t^\kappa, e^{-\tau_A R_{t,T}^A(s)}) = 1. \quad (129)$$

Since the risk-free benchmark is constant across states,

$$\mathbb{E}[m_t^\kappa] e^{-\tau_A \rho_f} = 1, \quad \mathbb{E}[m_t^\kappa] = e^{\tau_A \rho_f}. \quad (130)$$

Hence

$$\mathbb{E}[e^{-\tau_A R_{t,T}^A(s)}] - e^{-\tau_A \rho_f} = -e^{-\tau_A \rho_f} \text{Cov}(m_t^\kappa, e^{-\tau_A R_{t,T}^A(s)}). \quad (131)$$

Step 5: Informational beta. Define

$$\beta_{\kappa,A} := \frac{\text{Cov}(m_t^\kappa, e^{-\tau_A R_{t,T}^A(s)})}{\text{Var}(m_t^\kappa)}. \quad (132)$$

Step 6: Defining the disagreement-linked payoff. Define

$$e^{-\tau_A \rho_\kappa(s)} = m_t^\kappa(s), \quad (133)$$

so that

$$\rho_\kappa(s) = -\frac{1}{\tau_A} \left[\log\left(\frac{\delta_j}{\delta_i}\right) + \kappa_{t,T}(s) \right]. \quad (134)$$

Then

$$\mathbb{E}[e^{-\tau_A \rho_\kappa(s)}] - e^{-\tau_A \rho_f} = -e^{-\tau_A \rho_f} \text{Var}(m_t^\kappa). \quad (135)$$

Step 7: The Information CAPM. Dividing gives

$$\mathbb{E}\left[e^{-\tau_A R_{i,T}^A(s)}\right] - e^{-\tau_A \rho_f} = \beta_{\kappa,A} \left(\mathbb{E}\left[e^{-\tau_A \rho_\kappa(s)}\right] - e^{-\tau_A \rho_f}\right). \quad (136)$$

B.3 KL-Minimizing Aggregate Benchmark under Heterogeneous Beliefs

This appendix derives the aggregate benchmark distribution that minimizes the weighted average Kullback–Leibler distance from individual beliefs. The result is used in the main text to motivate the endogenous benchmark interpretation of prices under disagreement.

Suppose agent i holds subjective beliefs $p_i(n)$ over a finite state space $n \in \mathcal{N}$, and let w_i denote the relative market weight of that agent, with $\sum_i w_i = 1$. Consider

$$\min_q \quad \mathcal{L}(q) = \sum_i w_i \sum_{n \in \mathcal{N}} p_i(n) \log\left(\frac{p_i(n)}{q(n)}\right) \quad \text{s.t.} \quad \sum_{n \in \mathcal{N}} q(n) = 1. \quad (137)$$

The first-order condition yields

$$q^*(n) = \sum_i w_i p_i(n). \quad (138)$$

B.4 Link to the Entropy-Augmented Utility of Paper 1

The full derivation of the entropy-augmented utility representation is developed in Paper 1 and is not repeated here. For the purposes of the present section, we use only the reduced-form result that information-adjusted utility can be written as

$$\mathbb{U}_\gamma(x; P) = \mathbb{E}_P[u(x)] - \lambda(\gamma) \widetilde{H}(P),$$

where $u(x)$ is primitive Bernoulli utility, P is the first-order predictive belief, $\widetilde{H}(P)$ is the effective uncertainty associated with that belief structure, and $\lambda(\gamma)$ maps epistemic uncertainty into utility units. In the Ellsberg application, the only

additional object is the ambiguity premium Δ , which enters through

$$\widetilde{H}(P_A) = \ln 2 + \Delta, \quad \widetilde{H}(P_R) = \ln 2.$$

B.5 Closed-Form Derivation of the Ambiguity Premium $\Delta(\alpha, \beta)$

Let

$$\Theta \sim \text{Beta}(\alpha, \beta), \quad P(\Theta) = (\Theta, 1 - \Theta),$$

so that the agent entertains a second-order distribution over Bernoulli laws. The predictive distribution is

$$P_A = \mathbb{E}[P(\Theta)] = \left(\frac{\alpha}{\alpha + \beta}, \frac{\beta}{\alpha + \beta} \right).$$

The ambiguity premium is

$$\Delta(\alpha, \beta) := H(P_A) - \mathbb{E}_\Theta [H(P(\Theta))]. \quad (139)$$

The predictive entropy is

$$H(P_A) = -\frac{\alpha}{\alpha + \beta} \log\left(\frac{\alpha}{\alpha + \beta}\right) - \frac{\beta}{\alpha + \beta} \log\left(\frac{\beta}{\alpha + \beta}\right). \quad (140)$$

For $\Theta \sim \text{Beta}(\alpha, \beta)$,

$$\mathbb{E}[\Theta \log \Theta] = \frac{\alpha}{\alpha + \beta} \left(\psi(\alpha + 1) - \psi(\alpha + \beta + 1) \right), \quad (141)$$

$$\mathbb{E}[(1 - \Theta) \log(1 - \Theta)] = \frac{\beta}{\alpha + \beta} \left(\psi(\beta + 1) - \psi(\alpha + \beta + 1) \right). \quad (142)$$

Hence

$$\begin{aligned} \Delta(\alpha, \beta) &= -\frac{\alpha}{\alpha + \beta} \log\left(\frac{\alpha}{\alpha + \beta}\right) - \frac{\beta}{\alpha + \beta} \log\left(\frac{\beta}{\alpha + \beta}\right) \\ &\quad - \psi(\alpha + \beta + 1) + \frac{\alpha}{\alpha + \beta} \psi(\alpha + 1) + \frac{\beta}{\alpha + \beta} \psi(\beta + 1). \end{aligned} \quad (143)$$

B.6 Equivalence Between the Ambiguity-Based and Information-Based SDFs

This appendix clarifies the structural equivalence between the ambiguity-based pricing kernel developed in the Ellsberg section and the information-based stochastic discount factor of Munoz-Alonso (2025).

Recall the information-based disagreement kernel:

$$m^{ij}(s) = \exp\left(\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)}\right). \quad (144)$$

The ambiguity wedge is

$$\kappa_{ij}^\Delta := \lambda(\gamma)(\Delta_i - \Delta_j). \quad (145)$$

The baseline kernel is

$$m_{ji,t+1}^{(0)} = \frac{\delta_i}{\delta_j} \exp(-\tau_\rho \rho_{t+1}), \quad (146)$$

and the ambiguity-adjusted kernel is

$$m_{ji,t+1}^{\text{amb}} = m_{ji,t+1}^{(0)} \exp(\kappa_{ij}^\Delta). \quad (147)$$

The structural identification is

$$\frac{\tau_i}{u_i^*(s)} - \frac{\tau_j}{u_j^*(s)} \longleftrightarrow \lambda(\gamma)(\Delta_i - \Delta_j). \quad (148)$$

B.7 Maximum Entropy, Symmetry of Beliefs, and Informational Efficiency

This appendix formalizes the connection between three ideas used in the main text: maximum entropy, symmetry of beliefs, and the limiting notion of informational

efficiency. The purpose is to show that, in the special case where no state is privileged by incoming information, the uniform distribution is both the unique symmetric belief and the unique maximum-entropy benchmark. In that environment, convergence toward the benchmark is therefore equivalent to convergence toward maximum entropy.

Finite state space and entropy. Let $\mathcal{S} = \{s_1, \dots, s_n\}$ be a finite state space, with $n := |\mathcal{S}| \geq 2$. Let $\eta \in \Delta(\mathcal{S})$ denote a probability distribution over states, where

$$\Delta(\mathcal{S}) := \left\{ \eta \in \mathbb{R}_+^n : \sum_{k=1}^n \eta(s_k) = 1 \right\}.$$

The Shannon entropy of η is

$$H(\eta) := - \sum_{k=1}^n \eta(s_k) \log \eta(s_k). \quad (149)$$

Symmetry of beliefs. By symmetry of beliefs we mean invariance under a relabeling of states. Let Π denote the set of permutations of $\{1, \dots, n\}$. For any permutation $\pi \in \Pi$, define the permuted belief $\pi \cdot \eta$ by

$$(\pi \cdot \eta)(s_k) := \eta(s_{\pi^{-1}(k)}).$$

A belief distribution is symmetric if

$$\eta = \pi \cdot \eta \quad \text{for every } \pi \in \Pi.$$

This requires equal probability mass on every state, so the only symmetric distribution is the uniform one:

$$\eta^{\text{unif}}(s_k) = \frac{1}{n} \quad \forall k = 1, \dots, n. \quad (150)$$

Uniformity as the maximum-entropy benchmark. Consider the problem

$$\max_{\eta \in \Delta(\mathcal{S})} H(\eta).$$

The first-order conditions imply that the maximizing distribution must assign the same probability to every state. Hence the unique maximizer is η^{unif} , and its entropy is

$$H(\eta^{\text{unif}}) = \log n. \quad (151)$$

This establishes the equivalence between symmetry and maximum entropy in the finite-state case: the most symmetric belief is also the least informative one, in the sense of having maximal entropy.

Entropy as distance from symmetry. The same result can be written in relative-entropy form. The KL divergence from an arbitrary belief η to the uniform distribution is

$$D_{\text{KL}}(\eta \| \eta^{\text{unif}}) = \sum_{k=1}^n \eta(s_k) \log \frac{\eta(s_k)}{1/n}.$$

Expanding this expression yields

$$D_{\text{KL}}(\eta \| \eta^{\text{unif}}) = \log n - H(\eta). \quad (152)$$

Since $D_{\text{KL}}(\eta \| \eta^{\text{unif}}) \geq 0$, with equality if and only if $\eta = \eta^{\text{unif}}$, maximizing entropy is equivalent to minimizing KL distance to the symmetric benchmark.

Connection with informational efficiency. This identity clarifies the limiting argument used in the convergence section. In the no-signal case, no state receives new privileged informational content. If the benchmark distribution is therefore symmetric, $\bar{\nu} = \eta^{\text{unif}}$, then convergence of beliefs toward the benchmark implies

$$\eta_{i,t} \rightarrow \bar{\nu} = \eta^{\text{unif}},$$

and hence

$$H(\eta_{i,t}) \rightarrow H(\bar{\nu}) = \log n.$$

Thus, in the symmetric no-signal case, convergence toward the common benchmark coincides exactly with convergence toward maximum entropy.

Interpretation. This does not mean that every efficient market must literally correspond to a uniform distribution. In general, beliefs converge toward the invariant benchmark $\bar{\nu}$, which need not be uniform. The stronger maximum-entropy conclusion applies only in the special case where no informational asymmetry breaks symmetry across states. In that case, informational efficiency corresponds to the disappearance of relative wedges and the restoration of the most symmetric feasible belief configuration.

References

- Abbas, A. E. (2006). Maximum entropy utility. *Operations Research* 54(2), 277–290.
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109(3), 659–684.
- Akrenius, M. (2020). Information theory meets expected utility: The entropic roots of probability weighting functions. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 42.
- Alfarano, S., T. Lux, and F. Wagner (2005). Estimation of agent-based models: the case of an asymmetric herding model. *Computational Economics* 26, 19–49.
- Alvarez, F. and U. J. Jermann (2005). Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica* 73(6), 1977–2016.
- An, Y., Y. Su, and C. Wang (2024). Quantity, risk, and return: A beta-times-quantity (btq) asset pricing model. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4098609.
- Ardakani, O. M. (2022). Option pricing with maximum entropy densities: The inclusion of higher-order moments. *Journal of Futures Markets* 42(10), 1821–1836.
- Axtell, R. L. and J. D. Farmer (2022). Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature*, 1–101.
- Backus, D., C. Mikhail, and S. Zin (2014). Sources of entropy in representative agent models. *The Journal of Finance* 69(1), 51–99.
- Banerjee, S., R. Kaniel, and I. Kremer (2009). Price drift as an outcome of differences in higher-order beliefs. *The Review of Financial Studies* 22(9), 3707–3734.

- Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of financial economics* 49(3), 307–343.
- Bera, A. K. and S. Y. Park (2008). Optimal portfolio diversification using the maximum entropy principle. *Econometric Reviews* 27(4-6), 484–512.
- Bhamra, H. S., R. Uppal, and J. Walden (2019). The geography of beliefs. Technical report, Working paper.
- Bhamra, H. S., R. Uppal, and J. Walden (2022). Psychological distance and subjective beliefs. *Available at SSRN 4243916*.
- Bhattacharyya, R., M. B. Kar, S. Kar, and D. D. Majumder (2009). Mean-entropy-skewness fuzzy portfolio selection by credibility theory approach. In *International Conference on Pattern Recognition and Machine Intelligence*, pp. 603–608. Springer.
- Boehme, R. D., B. R. Danielsen, and S. M. Sorescu (2006). Short-sale constraints, differences of opinion, and overvaluation. *Journal of Financial and Quantitative Analysis* 41(2), 455–487.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives* 36(3), 223–244.
- Borwein, J. M. and Q. J. Zhu (2017). Entropy maximization in finance. In *Jonathan M. Borwein Commemorative Conference*, pp. 275–295. Springer, Cham.
- Branch, W. A. and G. W. Evans (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic theory* 127(1), 264–295.
- Branger, N. (2004). Pricing derivative securities using cross-entropy: An economic analysis. *International journal of theoretical and applied finance* 7(01), 63–81.

- Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7(3), 265–296.
- Brock, W. A. and C. H. Hommes (2001). A rational route to randomness. In *Growth Theory, Nonlinear Dynamics and Economic Modelling*, pp. 402–438. Edward Elgar Publishing.
- Brock, W. A., C. H. Hommes, et al. (1997). Models of complexity in economics and finance. Technical report.
- Cabrales, A., O. Gossner, and R. Serrano (2013). Entropy and the value of information for investors. *American Economic Review* 103(1), 360–377.
- Campbell, J. Y. (2014). Empirical asset pricing: Eugene fama, lars peter hansen, and robert shiller. *The Scandinavian Journal of Economics* 116(3), 593–634.
- Campbell, J. Y. (2017). *Financial decisions and markets: a course in asset pricing*. Princeton University Press.
- Cavagnaro, D. R., M. A. Pitt, R. Gonzalez, and J. I. Myung (2013). Discriminating among probability weighting functions using adaptive design optimization. *Journal of risk and uncertainty* 47, 255–289.
- Chiarella, C. (1992). The dynamics of speculative behaviour. *Annals of operations research* 37(1), 101–123.
- Cover, T. M. and J. A. Thomas (2006). *Elements of information theory*. Wiley Interscience.
- Das, S., C. M. Kuhnen, and S. Nagel (2020). Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies* 33(1), 395–432.
- Day, R. H. and W. Huang (1990). Bulls, bears and market sheep. *Journal of Economic Behavior & Organization* 14(3), 299–329.

- De Grauwe, P. and H. Dewachter (1993). A chaotic model of the exchange rate: the role of fundamentalists and chartists. *Open economies review* 4, 351–379.
- De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990). Noise trader risk in financial markets. *Journal of political Economy* 98(4), 703–738.
- Diether, K. B., C. J. Malloy, and A. Scherbina (2002). Differences of opinion and the cross section of stock returns. *The journal of finance* 57(5), 2113–2141.
- Dionisio, A., A. H. Reis, and L. Coelho (2008). Utility function estimation: The entropy approach. *Physica A: Statistical Mechanics and its Applications* 387(15), 3862–3867.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics* 75(4), 643–669.
- Eyster, E., M. Rabin, and D. Vayanos (2019). Financial markets where traders neglect the informational content of prices. *The Journal of Finance* 74(1), 371–399.
- Farmer, J. D. (2002, November). Market force, ecology and evolution. *Industrial and Corporate Change* 11(5), 895–953.
- Farmer, J. D. and D. Foley (2009, August). The economy needs agent-based modelling. *Nature* 460(7256), 685–686.
- Friedman, C. and S. Sandow (2016). *Utility-based learning from data*. Chapman and Hall/CRC.
- Ghosh, A., C. Julliard, and A. P. Taylor (2019). An information-theoretic asset pricing model. Technical report, Working Paper, London School of Economics.
- Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review* 111(5), 1481–1522.

- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Golan, A., G. Judge, and D. Miller (1997). *Maximum entropy econometrics: Robust estimation with limited data*. John Wiley and Sons.
- Gray, R. M. (2011). *Entropy and information theory*. Springer Science & Business Media.
- Grossman, S. J. and J. E. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American economic review* 70(3), 393–408.
- Guiasu, S. (1971). Weighted entropy. *Information and Control* 18(4), 323–332. Expands on weighted entropy in English.
- Gulko, L. (1999). The entropy theory of stock option pricing. *International Journal of Theoretical and Applied Finance* 2(03), 331–355.
- Gulko, L. (2002). The entropy theory of bond option pricing. *International Journal of Theoretical and Applied Finance* 5(04), 355–383.
- Gârleanu, N. and L. H. Pedersen (2021). Active and passive investing: Understanding samuelson’s dictum. *Review of Asset Pricing Studies* 12, 389–446.
- Hansen, L. P. and P. J. Maenhout (2006). Robust control and model misspecification. *Journal of Economic Theory* 128(1), 45–90.
- Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton university press.
- Hansen, L. P. and T. J. Sargent (2010). Wanting robustness in macroeconomics. In *Handbook of monetary economics*, Volume 3, pp. 1097–1157. Elsevier.
- Hansen, L. P. and T. J. Sargent (2011). Robustness and ambiguity in continuous time. *Journal of Economic Theory* 146(3), 1195–1223.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2001). Robustness and uncertainty aversion. *Manuscript, University of Chicago*.

- Harańczyk, G., W. Słomczyński, and T. Zastawniak (2008). Relative and discrete utility maximising entropy. *Open Systems & Information Dynamics* 15(04), 303–327.
- Harrison, J. M. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics* 92(2), 323–336.
- Hommes, C. (2013). *Behavioral rationality and heterogeneous expectations in complex economic systems*. Cambridge University Press.
- Hommes, C. H. (2001). Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance* 1(1), 149.
- Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of computational economics* 2, 1109–1186.
- Hong, H. and D. A. Sraer (2016). Speculative betas. *The Journal of Finance* 71(5), 2095–2144.
- Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of finance* 54(6), 2143–2184.
- Hong, H. and J. C. Stein (2007). Disagreement and the stock market. *Journal of Economic perspectives* 21(2), 109–128.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control* 17(5–6), 953–969.
- Imran, F. N. (2023). Entropic value-at-risk analysis of downside risks in major equity indices.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review* 106(4), 620.

- Kirkby, R. (2019). Bewley–huggett–aiyagari models: computation, simulation, and uniqueness of general equilibrium. *Macroeconomic Dynamics* 23(6), 2469–2508.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica* 73(6), 1849–1892.
- Kurz, M. (1994). On the structure and diversity of rational beliefs. *Economic theory* 4(6), 877–900.
- Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics* 47(1), 13–37.
- Liu, A., J. Chen, S. Y. Yang, and A. G. Hawkes (2020a). The flow of information in trading: An entropy approach to market regimes. *Entropy* 22, 1064.
- Liu, A., J. Chen, S. Y. Yang, and A. G. Hawkes (2020b). Information transition in trading and its effect on market efficiency: An entropy approach. *Quantitative Finance* 20, 143–157.
- Lucas Jr, R. E. (1978). Asset prices in an exchange economy. *Econometrica: journal of the Econometric Society*, 1429–1445.
- Lux, T. (2009). Rational forecasts or social opinion dynamics? identification of interaction effects in a business climate survey. *Journal of Economic Behavior & Organization* 72(2), 638–655.
- Lux, T. (2012). Estimation of an agent-based model of investor sentiment formation in financial markets. *Journal of Economic Dynamics and Control* 36(8), 1284–1302.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6), 1447–1498.

- MacKay, D. J. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge: Cambridge University Press.
- Maćkowiak, B., F. Matějka, and M. Wiederholt (2023). Rational inattention: A review. *Journal of Economic Literature* 61(1), 226–273.
- MacLean, L., L. Yu, and Y. Zhao (2022). A generalized entropy approach to portfolio selection under a hidden markov model. *Journal of Risk and Financial Management* 15(8), 337.
- Marschak, J. (1960). Remarks on the economics of information. Cowles Foundation Discussion Paper 146, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.
- Marschak, J. (1971). Economics of information systems. *Journal of the American Statistical Association* 66(333), 192–219.
- Marschak, J. (1974). Towards an economic theory of organization and information. In *Economic Information, Decision and Prediction*, Theory and Decision Library, pp. 29–62. D. Reidel, Dordrecht. Originally written in 1954.
- Mercurio, P. J., Y. Wu, and H. Xie (2020). An entropy-based approach to portfolio optimization. *Entropy* 22(3), 332.
- Milgrom, P. and N. Stokey (1982). Information, trade and common knowledge. *Journal of economic theory* 26(1), 17–27.
- Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of Finance* 32(4), 1151–1168.
- Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica* 34(4), 768–783.
- Munoz-Alonso, I. (2025). Information-based asset pricing with endogenous beliefs : Equilibrium. PhD Thesis.

- Olbryś, J. and E. Majewska (2022). Regularity in stock market indices within turbulence periods: The sample entropy approach. *Entropy* 24(7), 921.
- Ormos, M. and D. Zibriczky (2014). Entropy-based financial asset pricing. *Quantitative Finance* 14, 587–603.
- Patra, S. and G. S. Hiremath (2022). An entropy approach to measure the dynamic stock market efficiency. *The Journal of Financial Markets* 49, 1–20.
- Peng, L. (2005). Learning with information capacity constraints. *Journal of Financial and Quantitative Analysis* 40(2), 307–329.
- Pola, G. (2016). On entropy and portfolio diversification. *Journal of Portfolio Management* 42, 19–34.
- Risso, W. A. (2008). The informational efficiency and the financial crashes. *Research in international business and finance* 22(3), 396–408.
- Rouge, R. and N. El Karoui (2000). Pricing via utility maximization and entropy. *Mathematical Finance* 10(2), 259–276.
- Savage, L. J. (1954). The foundations of. *Statistics*, 11–34.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal* 27(3), 379–423.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19(3), 425–442.
- Shleifer, A. and L. H. Summers (1990). The noise trader approach to finance. *Journal of Economic perspectives* 4(2), 19–33.
- Sims, C. A. (1998). Stickiness. In *Carnegie-rochester conference series on public policy*, Volume 49, pp. 317–356. Elsevier.
- 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.
- Treynor, J. (1999). Toward a theory of market value of risky assets. In R. A. Korajczyk (Ed.), *Asset Pricing and Portfolio Performance*, pp. 15–22. University of Chicago Press.
- Veldkamp, L. L. (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.
- Xu, J. (2017). No trade under heterogeneous beliefs and risk aversion. *Journal of Economic Theory* 169, 428–452.
- Yang, J. and W. Qiu (2014). Normalized expected utility-entropy measure of risk. *Entropy* 16(7), 3590–3604.
- Yu, J. (2011). Disagreement and return predictability of stock portfolios. *Journal of Financial Economics* 99(1), 162–183.
- Yu, J.-R. and W.-Y. Lee (2011). Portfolio rebalancing model using multiple criteria. *European Journal of Operational Research* 209(2), 166–175.
- Yu, X. (2020). Risk-neutrality of rnd and option pricing within an entropy framework. *Entropy* 22(8), 836.
- Zanardo, E. (2017). *Essays in Microeconomics*. Phd thesis, Columbia University.
- Zunino, L., M. Zanin, B. M. Tabak, D. G. Pérez, and O. A. Rosso (2009). Forbidden patterns, permutation entropy and stock market inefficiency. *Physica A: Statistical Mechanics and its Applications* 388(14), 2854–2864.
- Zunino, L., M. Zanin, B. M. Tabak, D. G. Pérez, and O. A. Rosso (2010). Complexity-entropy causality plane: A useful approach to quantify the stock market inefficiency. *Physica A: Statistical Mechanics and its Applications* 389(9), 1891–1901.