

Bayesian Mechanics of Economic Choice: Computational Foundations of Economic Behavior

Samuel Montañez¹, Luis Alberto Quezada-Telléz²,
Ernesto Moya-Albor³

^{1,3} Universidad Panamericana,
Facultad de Ingeniería,
México

² Universidad Autónoma de Hidalgo,
México

{smontanez, emoya}@up.edu.mx, alquezada@ciencias.unam.mx

Abstract. This paper presents a theoretical unification of neuroeconomics with the Free Energy Principle (FEP) framework. We demonstrate that economic decision-making can be formulated as a variational inference problem where agents minimize expected free energy, balancing risk (aligning predictions with preferences) and ambiguity (reducing uncertainty). Our formal analysis establishes mathematical equivalence between divisive normalization in neuroeconomic models and precision-weighted prediction error minimization in active inference. We show how Expected Subjective Value Theory (ESVT) from neuroeconomics naturally emerges from the FEP under Gaussian assumptions, explaining context-dependent valuation, reference-dependence, and risk attitudes through a common computational mechanism and generative model. This unification has significant implications for artificial intelligence, providing computational principles for developing more human-like decision-making agents that balance exploration and exploitation in an information-theoretic way. By bridging Bayesian mechanics with divisive normalization, we provide a neurobiologically plausible foundation for economic behavior that encompasses both classical utility maximization and information-theoretic approaches to decision-making under uncertainty. By integrating thermodynamic principles of information processing, we demonstrate how economic decision-making operates under physical constraints, offering a theoretical foundation for AI systems that must optimize computational resources while managing uncertainty.

Keywords: Expected subjective value, variational free energy, active inference, neuroeconomics, Bayesian mechanics.

1 The Free Energy Principle as a Decision-Making Framework

Active inference has expanded beyond neuroscience into diverse domains [1, 2, 3, 4, 5, 6, 7], establishing itself as a unified framework for modeling decision-making under uncertainty. Its foundational principles—variational free energy minimization and belief updating through precision-weighted prediction errors [8]—provide mechanistic explanations for complex behaviors across multiple temporal scales.

This method which can be read as a physics of sentience is known as the new and growing field of Bayesian mechanics [8]. What distinguishes this approach is its ability to simultaneously address perception, learning, and action within a single coherent theoretical structure, accommodating both optimal and seemingly suboptimal behaviors that challenge traditional modeling approaches [9]. This integrative capacity makes active inference particularly valuable for artificial intelligence, where systems must similarly balance perception, learning, and action within unified computational architectures [1]. While reinforcement learning approaches separate perception from action and require external reward signals, active inference offers AI a more cohesive computational framework where perception, learning, and action emerge from the single imperative of free energy minimization [1].

The free energy principle (FEP) offers a neurobiologically plausible foundation for economic behavior by unifying Bayesian decision theory with statistical physics and information-theoretic approaches to uncertainty [10]. This framework shares mathematical equivalence with predictive coding [11] and relates conceptually to the Helmholtz free energy in statistical physics [12, 13]. Its computational architecture—minimizing the long-term average of sensory surprisal through an internal generative model—provides a principled basis for understanding both information processing and decision-making [14, 15, 16, 17].

Recent empirical evidence supports the role of dopamine in encoding precision of beliefs about optimal policies, demonstrating that humans employ hierarchical Bayesian inference to simultaneously determine both what they should do and how confident they should be a process that aligns more closely with active inference than with classical utility maximization [18]. This framework extends Barlow's efficient coding hypothesis [19], providing a formal basis for understanding economic decision-making as optimized information processing under biological constraints.

The application of FEP to economics has been limited despite its significant potential. While some research has applied these principles to temporal discounting [20] and bounded rationality [21], a formal connection between current neuroeconomic models and the FEP remains to be established. In this paper, we demonstrate that divisive normalization—a canonical neural computation central to Expected Subjective Value Theory (ESVT) [22] in neuroeconomics—emerges naturally from perceptual inference under the FEP. By proving this formal equivalence, we provide a theoretical foundation for understanding economic behaviors as manifestations of precision-weighted prediction error minimization in the brain, laying a bioinspired roadmap for developing AI agents capable of human-like economic decision-making and planning under the same principled information-theoretic framework [16].

2 Thermodynamic Foundations of Decision-Making

The minimization of complexity in free energy optimization has fundamental thermodynamic implications through Landauer's principle [23], which establishes a lower bound of $kT \ln(2)$ energy expenditure per bit erased. This physical constraint bridges information theory and thermodynamics, suggesting that cognitive efficiency has metabolic consequences. When economic agents employ parsimonious

representations of market dynamics, they reduce both complexity and associated energetic costs [24].

Under the FEP, neural information processing optimizes the trade-off between accuracy and complexity, manifesting biologically as the calibration of synaptic weights within hierarchical neural architectures [14]. This optimization process explains the prevalence of simplified heuristics and dimensionality reduction in successful economic strategies: such approaches preserve essential predictive power while minimizing metabolic expenditure [23] and stochasticity in choice [24].

The active inference framework formalizes this efficiency through precision-weighted prediction errors, where Bayesian belief updating is dynamically modulated by confidence estimates [8, 14]. This precision-weighting mechanism creates effective information compression, enabling adaptive decision-making even under severe computational constraints [21]. Economic agents implementing this mechanism naturally balance exploration (uncertainty reduction) with exploitation (preference satisfaction) without requiring exhaustive computation [16].

This thermodynamic perspective clarifies why seemingly "irrational" economic behaviors may represent optimal solutions under biological constraints [25]. Rather than implementing general-purpose rationality, economic cognition leverages domain-specific adaptations that exploit statistical regularities in environmental structure [26, 27]. These adaptations can be understood as instantiating a form of "Neural Darwinism," where neural architectures implement anti-entropic mechanisms that maintain functional organization against thermodynamic dissipation [14].

The resulting agent-environment system forms a Markov blanket structure that maintains integrity amid environmental fluctuations [8]. This statistical separation between internal and external states provides a formal basis for bounded rationality in economics, where access to market information is necessarily limited and costly. By minimizing expected free energy, economic agents effectively push "thermodynamically uphill" against disorder, implementing computationally efficient solutions [16]. This formulation provides a principled foundation for understanding economic decision-making as optimal inference under physical and informational constraints.

3 The Free Energy Principle and Active Inference

3.1 The Free Energy Principle Explained

The FEP posits that biological systems, including the human brain, strive to minimize free energy, which corresponds to reducing surprise or uncertainty by forming predictions about their environment [8, 10, 16]. This principle serves as a bridge between cognitive science and physics, suggesting that life, cognition, and evolution can be understood through a unified framework that aligns with variational principles such as Hamilton's principle of least action in classical mechanics [2].

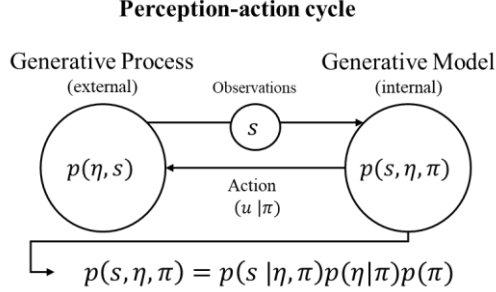


Fig. 1. Perception-action cycle in active inference. The figure illustrates the circular causal relationship between a generative process (external) and generative model (internal) connected through observations and actions. The generative process $p(\eta, s)$ represents the true causal structure governing the external states η and observations s , while the generative model $p(s, \eta, \pi)$ embodies the agent's beliefs about how observations and hidden states are generated, parameterized by policy π . Perception corresponds to inferring external states from observations, while action ($u | \pi$) influences the generative process according to policies inferred from the generative model. The joint distribution of the generative model factorizes as $p(s, \eta, \pi) = p(s | \eta, \pi) p(\eta | \pi) p(\pi)$, where $p(s | \eta, \pi)$ encodes the likelihood mapping, $p(\eta | \pi)$ represents conditional beliefs about external states given policies, and $p(\pi)$ encodes prior beliefs over policies. This formulation instantiates active inference, where agents select policies that minimize expected free energy, thereby reducing the divergence between the generative process and generative model through perception and action [9].

In active inference, the expected free energy (G) quantifies the probabilistic divergence between anticipated trajectories and preferred outcomes, while accounting for uncertainty reduction. Formally, $G(\alpha[\tau])$ represents the expected free energy of a policy or autonomous path $\alpha[\tau]$, equivalent to an action functional $\mathcal{A}(\alpha[\tau])$. The FEP states that agents select paths that minimize G formalized as [9]:

$$\alpha[\tau] = \arg \min_{\alpha[\tau]} G(\alpha[\tau]), \quad (1)$$

Perceptual inference minimizes free energy $F(s, \alpha)$ by updating internal states α :

$$\dot{\alpha}(\tau) = (Q_{\alpha\alpha} - \Gamma_{\alpha}) \nabla_{\alpha} F(s, \alpha). \quad (2)$$

This cyclical process creates a dynamic equilibrium between perception and action, formalized through the coupled equations [9]:

$$\dot{\alpha}(\tau) = (Q_{\alpha\alpha} - \Gamma_{\alpha}) \nabla_{\alpha} F(s, \alpha), \quad (3)$$

$$\alpha[\tau] = \arg \min_{\alpha[\tau]} G(\alpha[\tau]). \quad (4)$$

The closed perception-action loop implements an approximate Bayesian filtering scheme analogous to the Hamilton-Jacobi-Bellman (HJB) equation in continuous control, but uniquely optimizes both information gain and expected utility, explaining exploration-exploitation dynamics in economic decision-making [16, 18]. Active inference generalizes standard reinforcement learning by driving agents to minimize expected surprise rather than simply maximize rewards [28, 29]. This fundamental difference provides a more comprehensive framework for modeling decision-making

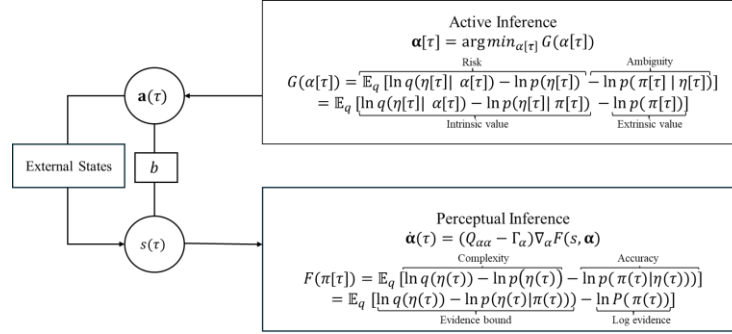


Fig. 2. Active and Perceptual Inference. This diagram illustrates the dual optimization processes in biological and sentient systems. The left panel depicts the relationship between external states $\eta \subset x$, sensory states ($s(\tau)$), and autonomous states $\alpha(\tau)$, linked by a Markov blanket b in current time τ . The autonomous states $\alpha = (a, \mu)$ comprise active states $a \subset x$, and internal states $\mu \subset x$, while the blanket states $b = (s, a)$ consist of sensory states $s(\tau)$ and active states $\alpha(\tau)$, collectively creating a partition of states that separates internal from external states $\eta(\tau)$. The top right panel shows Active Inference, where agents select policies $\pi(\tau)$ that minimize expected free energy $G(\alpha[\tau])$, balancing risk (aligning with preferences) against ambiguity (reducing uncertainty). The expected free energy represents the difference between posterior and prior beliefs about external states, which can be decomposed into intrinsic value (information gain) and extrinsic value (preference satisfaction). The bottom right panel represents Perceptual Inference, where agents update beliefs through gradient flows that minimize variational free energy $F(s, \alpha)$, optimizing the trade-off between accuracy and complexity. Here $q(\eta(\tau))$ denotes the recognition density or approximate posterior distribution over external states at time τ . The term $F(s, \alpha)$ is this free energy functional, and $\nabla_{\alpha} F(s, \alpha)$ is its gradient with respect to internal states, with precision weighting $Q_{\alpha\alpha}$ and friction Γ_{α} modulating the update speed. Together, these complementary processes enable adaptive self-organization through continual prediction and uncertainty minimization [8,9].

under uncertainty, where agents update beliefs and select actions to minimize surprise rather than merely accumulate rewards.

3.2 Variational Free Energy

When considering perceptual inference, we can express the variational free energy through the Kullback-Leibler (KL) divergence between an approximate posterior distribution $q(\eta)$ and the true posterior $p(\eta|s)$ [30]:

$$KL(q(\eta), p(\eta|s)) = \int q(\eta) \left(\ln \frac{q(\eta)}{p(\eta|s)} \right) d\eta, \quad (5)$$

By substituting the definition of conditional probability and taking logarithms [30]:

$$KL(q(\eta), p(\eta|s)) = -F + \ln p(s), \quad (6)$$

where F is the negative free energy. Assuming that $q(\eta)$ is a delta distribution, the negative free energy simplifies to [30]:

$$F = \int q(\eta) \ln \left(\frac{p(s|\eta)}{q(\eta)} \right) d\eta = \int q(\eta) \ln p(s|\eta) d\eta - \int q(\eta) \ln q(\eta) d\eta. \quad (7)$$

This assumption implies that $q(\eta) = \delta(\eta - \hat{\eta})$, where δ is the Dirac delta function and $\hat{\eta}$ represents a point estimate. This effectively transforms the variational problem into a maximum a posteriori (MAP) estimation, eliminating the entropy term $-\int q(\eta) \ln q(\eta) d\eta$ and simplifying the free energy formulation [8].

The expectation operator \mathbb{E}_q represents the weighted average with respect to the distribution q :

$$F = \mathbb{E}_{q(\eta)} \left[\ln \left(\frac{p(s|\eta)}{q(\eta)} \right) \right] = \int q(\eta) \ln \left(\frac{p(s|\eta)}{q(\eta)} \right) d\eta. \quad (8)$$

If we define a generative model as $m = p(s, \eta) = p(s|\eta) p(\eta)$, then [9]:

$$F = \mathbb{E}_{q(\eta)} \left[\ln \left(\frac{q(\eta)}{p(\eta|s)} \right) \right] - \ln p(s), \quad (9)$$

where $q(\eta) = p(s|\eta)$, free energy reduces to surprise $-\ln p(s)$.

Introducing policies the framework extends to active inference [9]:

$$F = \mathbb{E}_{q(\eta|\pi)} \left[\ln \left(\frac{q(\eta|\pi)}{p(\eta|s, \pi)} \right) \right] - \ln p(s|\pi). \quad (10)$$

With time dependencies (τ), the expectation operator becomes [9]:

$$\mathbb{E}_{q(\eta(\tau)|\pi(\tau))} [f(\eta(\tau), \pi(\tau))] = \int q(\eta(\tau)|\pi(\tau)) f(\eta(\tau), \pi(\tau)) d\eta(\tau). \quad (11)$$

The expectation notation encapsulates the integration over all possible states [9]:

$$F(\pi(\tau)) = \mathbb{E}_q [\ln q(\eta(\tau)) - \ln p(\eta(\tau)) - \ln p(\pi(\tau)|\eta(\tau))]. \quad (12)$$

This extension is often called "active inference" where the system not only infers hidden states but also selects policies that minimize expected free energy in the future. In active inference, the agent selects policies that are expected to minimize surprise (or maximize model evidence) in the future. This allows the framework to address not just perceptual inference but also decision-making and planning to fulfill their preferences or prior beliefs about desired states [16,31].

4 From Free Energy to Economics

4.1 Optimal Encoding Strategy

Steverson et al. [24] proposed that economic decision-making can be formalized as an optimization problem balancing expected utility maximization against information-processing costs. Making precise (non-stochastic) choices requires cognitive resources, which can be quantified using information theory. The optimization problem they propose takes the form:

$$\rho(x, A) \in \arg \max \sum_{x \in A} p(x) v(x) - C_A(\Delta H(p)), \quad (13)$$

where $\rho(x, A)$ is the probability of choosing option x from choice set A , $v(x)$ is the value of option x , and $C_A(\Delta H(p))$ represents the cognitive cost of reducing choice entropy. The entropy reduction term $\Delta H(p)$ measures deviation from random choice [24]:

$$\Delta H(p) := \ln|A| - H(p), \quad (14)$$

where $H(p)$ is the Shannon entropy of the choice distribution:

$$H(p) := - \sum_{x \in A} \left[\frac{p(x)}{p(A)} \right] \ln \left[\frac{p(x)}{p(A)} \right], \quad (15)$$

Solving this optimization problem, the resulting choice probabilities follow the form:

$$\rho(x, A) = \frac{\exp(\frac{\gamma v(x)}{\sigma + v(A)})}{\sum_{y \in A} \exp(\frac{\gamma v(y)}{\sigma + v(A)})}. \quad (16)$$

This expression is structurally identical to a softmax over divisively normalized values [32], providing a formal derivation of context-dependent valuation. This derivation shows that divisive normalization emerges naturally when agents optimize a trade-off between maximizing value and minimizing cognitive costs. The divisive term in the denominator $\sigma + v(A)$ effectively normalizes option values based on the overall value of the choice set, just as neurons in the brain normalize their firing rates based on surrounding activity [24].

This optimization approach aligns with principles from active inference, where agents must minimize entropy to maintain a stable identity within fluctuating environments [2]. This can be formalized using the entropy of the probability distribution $p(s^*)$ of finding the agent in a given state s^* of its state space S^* [1]:

$$H(S^*) = \int_{S^* \in S^*} (-\ln p(s^*)) p(s^*) ds^*, \quad (17)$$

where s^* can be replaced by observation space S . This entropy minimization corresponds precisely to the complexity term in the variational free energy functional, establishing a direct mathematical link between economic choice and the FEP framework. This connection has deep roots in statistical physics, where the Boltzmann distribution emerges as the probability distribution that minimizes Helmholtz free energy while maintaining a fixed average energy [1, 24]. The formal equivalence between thermodynamic systems and decision-making agents is not merely analogical—both involve systems that dynamically settle into probability distributions that optimize a free energy functional, subject to constraints [8, 13]. Just as physical systems minimize thermodynamic free energy to reach equilibrium, cognitive systems appear to minimize information-theoretic free energy to optimize behavior under constraints. Moreover, the entropy component parallels recent advances in reinforcement learning, where incorporating entropy regularization into reward objectives enhances algorithmic performance, stabilizes policy optimization, and improves generalization capabilities [2, 16].

4.2 From Free Energy to Expected Utility Theory

There is a clear path from the FEP to expected utility theory in economics. The FEP conceptualizes the world as a random dynamical system, with active inference explaining how self-organization emerges [8]. Within this framework, economic agents can be treated as adaptive systems whose behaviors are amenable to analysis through Bayesian mechanics [33]. This formulation offers a principled foundation for addressing unresolved challenges in economic decision-making, as it integrates effectively with recent advances in neuroeconomics through its accompanying neural process theory, enabling testable empirical predictions about neural responses [9].

When the epistemic value component is removed from the active inference framework, what remains is essentially the expected log probability of preferred outcomes. This is equivalent to maximizing expected utility ($\mathbb{E}[U_A]$) in economic theory [8,34]:

$$\mathbb{E} \ln p(\pi[\tau]) \cong \mathbb{E}[U_A] = \sum_{o \in O} P(o|A) U(o), \quad (18)$$

where $P(o|A)$ represents the conditional probability of obtaining outcome o given action A , and $U(o)$ denotes the utility or value associated with each possible outcome o in the set O . Expected utility maximization emerges as a special case of active inference when uncertainty reduction is not prioritized [8].

Rational choice theory posits that decision-makers employ rational calculations to optimize outcomes aligned with their preferences. Within this framework, preferences over independent outcomes remain consistent regardless of irrelevant factors—a fundamental principle known as independence [35]. Von Neumann and Morgenstern [34] interpreted conditional probabilities as objective chances within a perfectly rational framework rather than as beliefs about states. Furthermore, rational choice theory assumes that options possess absolute values independent of the value or existence of alternative options [36]. Active inference provides a more general formulation by reinterpreting utility functions as prior preference distributions, suggesting that observed behavior can be understood as Bayes optimal under some prior beliefs [18].

4.3 Expected Subjective Value Theory: A Neuroeconomic Model

Divisive normalization has emerged as a critical computation employed by the brain to facilitate decision-making. It functions as a canonical neural computation, contributing to efficient processing within neural circuits [32]. The work of Reynolds and Heeger [37] indicates that rectification can approximate a power law, resulting in contrast-response functions that align more closely with electrophysiological data than previous assumptions [22, 36]. Research has posited that the brain utilizes divisive normalization in a utility-like calculation during choice-making, which entails balancing the expected value of options against the entropic cost of reducing stochasticity [38-39].

Glimcher and Tymula's biophysical implementation in Expected Subjective Value Theory (ESVT) [22]—a neuroeconomics-based model of expected utility—demonstrate that normalization emerges from interacting excitatory and inhibitory neurons, described by coupled differential equations [22, 40]:

$$\tau \frac{dR}{dt} = -R + \frac{x}{1+G}, \quad (19)$$

$$\tau \frac{dG}{dt} = -G + R, \quad (20)$$

where R represents excitatory activity, G inhibitory activity, and x the objective value of the payoff (utility). This system converges to a unique equilibrium state:

$$\tau \frac{dG}{dt} = -G + R. \quad (21)$$

Two properties emerge: (1) this equilibrium state corresponds to standard divisive normalization, and (2) normalization arises from temporal integration of value inputs. In dynamic contexts, the action potential rate evolves as [22]:

$$R_t \propto \frac{x_t}{x_t + \sum_{k=0}^{t-1} D(k)x_k}, \quad (22)$$

where the denominator represents a weighting function $D(k)$ and a time-discounted average $(\sum_{k=0}^{t-1} D(k)x_k)$ of previously encountered payoffs (x_k) —effectively implementing reference-dependent valuation through neural computation [22]. This neurobiological implementation provides an understanding of how the brain might encode subjective values. As we will demonstrate in subsequent sections, this same normalization structure emerges from free energy minimization under specific assumptions.

The core of ESVT is a subjective value function mapping objective payoffs to neural representations [22]:

$$S_t(x) = \frac{x^\alpha}{x^\alpha + M_t^\alpha}, \quad (23)$$

where $S_t(x)$ represents the subjective value of payoff $x \in \mathbb{R}$ at time t , which corresponds to the neural firing rate encoding the value representation. The term M_t denotes the payoff expectation based on previously experienced outcomes, implementing a form of reference-dependence that emerges from neurobiological architecture. The predisposition parameter α controls the curvature of the value function, capturing individual differences in risk attitudes and value sensitivity, with lower values producing concave functions (risk aversion) and higher values yielding the characteristic sigmoid functions observed in prospect theory [22].

The payoff expectation is recursively computed as a time-weighted average of previous outcomes [22]:

$$S_t(x) = \frac{x^\alpha}{x^\alpha + M_t^\alpha}, \quad (24)$$

where $\gamma \in (0,1)$ represents the forgetting rate, capturing recency effects in expectation formation.

This formulation produces a subjective value function bounded between 0 and 1, consistent with neurobiological constraints on value encoding. Critically, unlike traditional utility formulations, ESVT provides a cardinal measure of subjective value that corresponds directly to neural firing rates observed in valuation regions of the

brain. ESVT captures not just ordering but also the intensity or magnitude of preferences, allowing for more precise predictions about behavior [22].

In ESVT neuronal firing rates represent excitatory input modulated by surrounding activity parallels the precision-weighting mechanisms in perceptual inference [22]. Research indicates that dynamic divisive normalization operates not only spatially but also temporally, influencing how perceptual evidence is weighted over time during decision-making tasks [41]. The context-dependence of divisive normalization has been linked to behavioral features previously unnoticed by economists, suggesting that it can predict how individuals adapt their preferences based on reward contexts during reinforcement learning [36,42].

4.4 From Free Energy to Divisive Normalization

In this section, we formally demonstrate that divisive normalization—the key computational mechanism in ESVT—emerges naturally from perceptual inference under the FEP framework, establishing a formal equivalence between these approaches.

The FEP and ESVT both characterize how neural systems optimize information processing under biological constraints, though they emerged from different disciplinary traditions. Both frameworks describe neural systems that maximize information transmission while minimizing metabolic costs. While divisive normalization describes how neurons encode information about the world [32], this precisely corresponds to perceptual inference within the FEP framework [8].

We demonstrate that divisive normalization emerges from perceptual inference under the FEP by considering a hierarchical generative model with Gaussian priors and likelihoods. Let's consider the free energy under Gaussian assumptions [9, 30]:

$$F = D_{KL}[q(\eta|\alpha)||p(\eta|s)] - \ln p(s). \quad (25)$$

At steady state ($\dot{\alpha} = 0$), internal states satisfy:

$$\nabla_{\alpha} F(s, \alpha) = 0. \quad (26)$$

Under Gaussian assumptions for recognition and generative densities is [8]:

$$q((\eta[\tau]|\alpha[\tau])) = \mathcal{N}(\mu_{\eta}, \Sigma_{\eta}), \quad (27)$$

$$p(\eta[\tau]|\alpha[\tau]) = \mathcal{N}(g(\pi[\tau]), \Pi_{\eta}^{-1}), \quad (28)$$

where Π_{η}^{-1} is the precision of prediction errors and Σ_{η} is the covariance matrix (inverse of the precision matrix Π_{η}). The gradient of free energy with respect to the conditional mean becomes:

$$\nabla_{\mu_{\eta}} F = \Pi_{\eta} (\mu_{\eta} - g(\pi[\tau])) - \nabla_{\mu_{\eta}} h(\eta)^T \Pi_s (s - h(g(\pi[\tau]))). \quad (29)$$

Lemma 1. (Divisive Normalization Equivalence): Under a hierarchical generative model with Gaussian priors and likelihoods, perceptual inference through free energy minimization converges to a representation that is equivalent to divisive normalization as formulated in Expected Subjective Value Theory.

Let a generative model be defined with Gaussian recognition density $q((\eta[\tau]|\alpha[\tau])) = \mathcal{N}(\mu_{\eta}, \Sigma_{\eta})$ and generative density $p(\eta[\tau]|\alpha[\tau]) = \mathcal{N}(g(\pi[\tau]), \Pi_{\eta}^{-1})$.

At the steady state where $\nabla_{\alpha} F(s, \alpha) = 0$, the optimal internal representation μ_{η} takes the form:

$$\mu_{\eta} = \frac{g(\pi[\tau]) + \sum_{\eta} \nabla_{\mu_{\eta}} h(\eta)^T \Pi_s (s - h(g(\pi[\tau])))}{1 + \sum_{\eta} \nabla_{\mu_{\eta}} h(\eta)^T \Pi_s \nabla_{\mu_{\eta}} h(\eta)}, \quad (30)$$

where the numerator represents direct input (predicted value), and the denominator represents a baseline term plus contextual modulation that implements precision-weighted prediction error minimization. This representation is structurally equivalent to the divisive normalization formulation in ESVT [22].

Proof: Consider the variational free energy under Gaussian assumptions [9]:

$$F = D_{KL}[q(\eta|\alpha)||p(\eta|s)] - \ln p(s), \quad (31)$$

At steady state, internal states satisfy $\nabla_{\alpha} F(s, \alpha) = 0$. The gradient of free energy with respect to the conditional mean is:

$$\nabla_{\mu_{\eta}} F = \Pi_{\eta} (\mu_{\eta} - g(\pi[\tau])) - \nabla_{\mu_{\eta}} h(\eta)^T \Pi_s (s - h(g(\pi[\tau]))), \quad (32)$$

where Π_{η} is the precision of the prior (inverse covariance), μ_{η} is the conditional mean (posterior expectation), $h(\eta)^T$ is the mapping from hidden states to observations, Π_s is the precision of sensory prediction errors, and $(s - h(g(\pi[\tau])))$ represents the sensory prediction error. This can be directly mapped to the ESVT's divisive normalization equation [22]:

$$S_t(x) = \frac{x^{\alpha}}{x^{\alpha} + M_t^{\alpha}}, \quad (33)$$

with the following correspondences:

$S_t(x) \leftrightarrow \mu_{\eta}$ (neural activity encodes expected causes),

$x^{\alpha} \leftrightarrow g(\pi[\tau])$ (direct input maps to prior prediction),

$x^{\alpha} \leftrightarrow 1$ (baseline term maps to constant normalization factor),

$M_t^{\alpha} \leftrightarrow \sum_{\eta} \nabla_{\mu_{\eta}} h(\eta)^T \Pi_s \nabla_{\mu_{\eta}} h(\eta)$ (contextual modulation maps to precision-weighted prediction error terms).

Therefore, divisive normalization in neural systems can be understood as implementing precision-weighted prediction error minimization as prescribed by the FEP.

This equivalence explains both why divisive normalization appears throughout the nervous system and provides a theoretical foundation for understanding seemingly irrational economic behaviors as optimal inference under constraints. The metabolic efficiency principle aligns with our economic agent model, suggesting that context-dependent valuation and reference-dependence emerge naturally from free energy minimization. By demonstrating this mathematical correspondence, we provide a mechanistic explanation for how Bayesian mechanics can describe economic choice.

5 Conclusion

Our formal demonstration of mathematical equivalence between divisive normalization and free energy minimization under Gaussian assumptions establishes a fundamental principle of neural computation. Through Lemma 1, we proved that the steady-state solution to variational inference under the FEP yields precisely the divisive normalization model central to modern neuroeconomic models. This equivalence explains the pervasiveness of normalization across neural systems and reveals it not merely as an implementation detail, but as a necessary property of self-organizing systems maintaining integrity against environmental entropy. By establishing this bridge between neural implementation and economic theory, our analysis unifies ESVT [22] with active inference [9]. This unification has profound implications for artificial intelligence development, particularly for systems that require adaptive decision-making and planning under uncertainty.

The computational mechanisms we identify could lead to more efficient AI architectures that naturally balance exploration and exploitation, mirroring the specialized resource-constrained optimization that biological systems have evolved to implement. This integration aligns with Barlow's efficient coding hypothesis [19], demonstrating that the neural substrates of economic decision-making reflect computational imperatives constrained by biological reality. The precision-weighting mechanism explains how context-dependent valuation and reference-dependence emerge naturally from free energy minimization. Our framework explains the exploration-exploitation trade-off fundamental to economic decisions as the natural balance between ambiguity minimization and risk management within the active inference formalism. For AI, these principles provide a computational neuroscience foundation for designing systems that can adapt to context-dependent valuations and make decisions that appear irrational under classical utility theory yet are optimal when accounting for informational and computational constraints. AI agents implementing these principles could better model human economic behavior while also achieving greater computational efficiency through principled dimensionality reduction similar to what Barlow called “economy of thought” [19].

This synthesis offers behavioral scientists a comprehensive framework for understanding decision-making that spans from computational principles to neural mechanisms. By adopting this perspective, researchers can develop more nuanced models of economic behavior that account for the cognitive processes underlying decision-making, particularly how agents navigate uncertainty through predictive modeling.

Future AI research could leverage these insights to develop agents that not only maximize reward but actively maintain their internal model integrity through Bayesian mechanics, potentially leading to more robust and adaptive systems in complex, changing environments. Furthermore, future research should examine whether the correspondence we identified extends beyond Gaussian assumptions. The application of the FEP to economics enables a deeper understanding of economic behavior through the solid foundations of information theory and statistical mechanics, offering a roadmap for physics-informed and bioinspired AI that is simultaneously more human-like in its decision processes and more principled in its computational implementation.

Acknowledgments. Samuel Montañez and Ernesto Moya-Albor extend their gratitude to the Faculty of Engineering at Universidad Panamericana for their support throughout this research endeavor. Some of the graphical representations and explanations were taken or inspired by papers created by K. Friston and colleagues.

References

1. Ueltzhöffer, K.: Deep active inference. *Biological Cybernetics* 112(6), pp. 547–573 (2018) doi: 10.1007/s00422-018-0785-7.
2. Ramstead, M.J., Constant, A., Badcock, P.B., Friston, K.J.: Variational ecology and the physics of sentient systems. *Phys Life Rev* 31, pp. 188–205 (2019) doi: 10.1016/j.plrev.2018.12.002.
3. Kawahara, D., Ozeki, S., Mizuuchi, I.: A curiosity algorithm for robots based on the free energy principle. In: *IEEE/SICE International Symposium on System Integration (SII)*, pp. 53–59 (2022) doi: 10.1109/SII52469.2022.9708814
4. Van de Maele, T., Verbelen, T., Çatal, O., De Boom, C., Dhoedt, B.: Active vision for robot manipulators using the free energy principle. *Frontiers in Neurorobotics*, 15, pp. 642780 (2021) doi: 10.3389/fnbot.2021.642780.
5. Wirkuttis, N., Ohata, W., Tani, J.: Turn-taking mechanisms in imitative interaction: Robotic social interaction based on the free energy principle. *Entropy* 25(2), pp. 263 (2023) doi: 10.3390/e25020263.
6. Smith, R., Ramstead, M.J.D., Kiefer, A.: Active inference models do not contradict folk psychology. *Synthese*, 200(2), pp. 81 (2022) doi: 10.1007/s11229-022-03624-y.
7. Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K.J., Ramstead, M.J.: Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation*, 33(2), pp. 398–446 (2021) doi: 10.1162/neco_a_01341.
8. Friston, K., Da Costa, L., Sajid, N., Heins, C., Ueltzhöffer, K., Pavliotis, G.A., Parr, T.: The free energy principle made simpler but not too simple. *Physics Reports* 1024, pp. 1–29 (2023)
9. Smith, R., Friston, K.J., Whyte, C.J.: A step-by-step tutorial on active inference and its application to empirical data. *J Math Psychol* 107: pp. 102632 (2021) doi: 10.1016/j.jmp.2021.102632.
10. Parr, T., Friston, K.J.: Active inference, Bayesian optimal design, and expected utility. *Int J Forecast* 35(2), pp. 219–233 (2019)
11. Rao, R.P.N., Ballard, D.H.: Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), pp. 79–87 (1999) doi: 10.1038/4580.
12. Dayan, P., Hinton, G.E., Neal, R.M., Zemel, R.S.: The helmholtz machine. *Neural Computation*, 7(5), pp. 889–904 (1995)
13. Friston, K. J., Daunizeau, J., Kilner, J., Kiebel, S. J.: Action and behavior: a free-energy formulation. *Biological Cybernetics* (2010) doi: 10.1007/s00422-010-0364-z.
14. Friston, K.: The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 11, pp. 127–138 (2010) doi: 10.1038/nrn2787.
15. Friston, K.: I am ther before I think. In: *The Unconscious*. Routledge, London, pp 127–151 (2016)
16. Da Costa, L., Sajid, N., Parr, T., Friston, K., Smith, R.: Reward Maximization through Discrete Active Inference. *Neural Computation* 35(5), pp. 807–852 (2023)
17. Friston, K.J., Salvatori, T., Isomura, T., Tschantz, A., Kiefer, A., Verbelen, T., Sajid, N., Adams, R.A., Parr, T., Ramstead, M.J.: Active Inference and Intentional Behavior. *Neural Computation* 37(4), pp. 666–700 (2025)

18. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G.: Active inference and learning. *Neuroscience & Biobehavioral Reviews* 68: pp. 862–879 (2016).
19. Barlow, H.B.: Possible principles underlying the transformations of sensory messages. In: Rosenblith WA (ed) *Sensory communication*. MIT Press, Cambridge, pp 217–234 (1961)
20. Henriksen, M.: Variational free energy and economics optimizing with biases and bounded rationality. *Front Psychol* 11, pp. 549187 (2020) doi: 10.3389/fpsyg.2020.549187.
21. Genewein, T., Leibfried, F., Grau-Moya, J., Braun, D.: Bounded Rationality, Abstraction, and Hierarchical Decision-Making: An Information-Theoretic Optimality Principle. *Frontiers in Robotics and AI*, 2, pp. 27 (2015) doi: 10.3389/frobt.2015.00027.
22. Tymula, A., Glimcher, P.: Expected subjective value theory (ESVT): A representation of decision under risk and certainty. Available at SSRN 2783638 (2022)
23. Landauer, R.: Irreversibility and heat generation in the computing process. *IBM J Res Dev* 5, pp. 183–191 (1961) doi: 10.1147/rd.53.0183.
24. Steverson, K., Brandenburger, A., Glimcher, P.: Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior & Organization*, 164, pp. 148–165 (2019) doi: 10.1016/j.jebo.2019.05.026.
25. Tooby, J., Cosmides, L., Barrett, H.C.: The second law of thermodynamics is the first law of psychology: evolutionary developmental psychology and the theory of tandem, coordinated inheritances: comment on Lickliter and Honeycutt. *Psychol Bull* 129, pp. 858–865 (2003) doi: 10.1037/0033-2909.129.6.858.
26. Cosmides, L., Tooby, J.: Better than rational: Evolutionary psychology and the invisible hand. *Am Econ Rev* 84, pp. 327–332 (1994)
27. Cosmides, L., Barrett, H.C., Tooby, J.: Adaptive specializations, social exchange, and the evolution of human intelligence. *Proc Natl Acad Sci USA* 107, pp. 9007–9014 (2010) doi: 10.1073/pnas.0914623107.
28. Friston, K.J., Lin, M., Frith, C.D., Pezzulo, G., Hobson, J.A., Ondobaka, S.: Active inference, curiosity and insight. *Neural Computation* 29, pp. 2633–2683 (2017) doi: 10.1162/neco_a_00999.
29. Millidge, B.: Deep active inference as variational policy gradients. *Journal of Mathematical Psychology* 96, pp. 102348 (2020) doi: 10.1016/j.jmp.2020.102348.
30. Bogacz, R.: A tutorial on the free-energy framework for modelling perception and learning. *J Math Psychol* 76, pp. 198–211 (2017) doi: 10.1016/j.jmp.2015.11.003.
31. Kaplan, R., Friston, K.J.: Planning and navigation as active inference. *Biol Cybernet* 112, pp. 323–343 (2018) doi: 10.1007/s00422-018-0753-2.
32. Carandini, M., Heeger, D.J.: Normalization as a canonical neural computation. *Nat Rev Neurosci* 13, pp. 51–62 (2012) doi: 10.1038/nrn3136.
33. Da Costa, L., Friston, K., Heins, C., Pavliotis, G.A.: Bayesian mechanics for stationary processes. *Proceedings of the Royal Society A* 477(2256), pp. 20210518 (2021)
34. Von Neumann, J., Morgenstern, O.: *Theory of games and economic behavior*. Princeton University Press, Princeton (1944)
35. Hemmatian, B., Varshney, L.R., Pi, F., Barbey, A.K.: The utilitarian brain: Moving beyond the Free Energy Principle. *Cortex*, 170, pp. 69–79 (2024) doi: 10.1016/j.cortex.2023.11.013.
36. Louie, K., Grattan, L.E., Glimcher, P.W.: Reward value-based gain control: divisive normalization in parietal cortex. *The Journal of Neuroscience*, 31(29), pp. 10627–10639 (2011) doi: 10.1523/JNEUROSCI.1237-11.2011.
37. Reynolds, J.H., Heeger, D.J.: The normalization model of attention. *Neuron* 61(2), pp. 168–185 (2009)
38. Webb, R., Glimcher, P.W., Louie, K.: The normalization of consumer valuations: Context-dependent preferences from neurobiological constraints. *Management Science* 67(1), 93–125 (2021) doi: 10.1287/mnsc.2019.3557.

39. Kurtz-David, V., Sinha, S., Alladi, V., Bucher, S., Brandenburger, A., Louie, K., Glimcher, P., Tymula, A.: A Tale of Two Environments: Divisive Normalization and the Flexibility of Choice. *bioRxiv* (2024) doi: 10.1101/2024.08.25.609561.
40. LoFaro, T., Louie, K., Webb, R., Glimcher, P.W.: The Temporal Dynamics of Cortical Normalization Models of Decision-making. *Letters in Biomathematics*, 1(2), pp. 209–220 (2014). <http://www.lettersinbiomath.org>
41. Keung, W., Hagen, T. A., & Wilson, R. C.: A divisive model of evidence accumulation explains uneven weighting of evidence over time. *Nature Communications*, 11(1), pp. 2160 (2020) doi: 10.1038/s41467-020-15630-0.
42. Bossaerts, P., Murawski, C.: From behavioural economics to neuroeconomics to decision neuroscience: the ascent of biology in research on human decision making. *Current Opinion in Behavioral Sciences*, 5, pp. 37–42 (2015) doi: 10.1016/j.cobeha.2015.07.001.