The Master Work Function: A Scalar Potential for Hybrid Human-Machine Meaningful Work

Krishna Patel

***Abstract*—Standard economic and engineering models of automation optimize throughput, unit cost, or task completion speed while treating human autonomy, skill mastery, and eudaimonic meaning as unmodeled externalities. This paper introduces and develops a rigorous axiomatic, dynamical, and empirical theory of the *Master Work Function* (MWF), a scalar socio-technical potential $W : X_H \times X_M \to \mathbb{R}$ mapping human agency states and machine capability allocations to joint productive and experiential welfare. First, we establish an anchored functional decomposition $W(h, m) = W_0 + W_H(h) + W_M(m) + W_{\text{int}}(h, m)$ with unique boundary identification and extend it to N-worker, M-agent organizational hypergraphs via functional ANOVA. Second, we introduce cognitive-overload and machine-pacing regulation, proving via Schur-complement Hessian bounds that regulated logarithmic-bilinear potentials admit a unique globally concave interior optimum (h^*, m^*), while Topkis supermodularity on the cross-partial matrix $\Gamma_{ij} = \partial^2 W / \partial h_i \partial m_j$ governs whether technological progress crowds in or erodes human engagement. Third, we embed organizational transitions in overdamped Langevin dynamics under time-varying deployment protocols $\lambda(t)$, deriving an exact socio-technical Jarzynski equality, Crooks fluctuation relation, and Sivak–Crooks minimum-dissipation rollout geodesic. Fourth, we formulate inequality-averse Atkinson–Rawls social aggregation and Holmstrom–Milgrom multitask bounds against Goodhart proxy distortion. Finally, we validate the framework across six linked public datasets spanning 873 U.S. occupations (124.0 million workers in O*NET 30.1, BLS OEWS, and 2025–2026 Anthropic Economic Index telemetry) and $10{,}483$ respondents in the NORC General Social Survey (2018–2024). Structural estimation strongly rejects separability ($F(16, 844) = 4.56, p = 8.09 \times 10^{-9}$; cross-validated $R^2 = 0.782$), reveals a 39.6 percentage-point augmentation gap between interactive human-in-the-loop interfaces and headless programmatic APIs ($t = 75.88, p < 10^{-100}$), confirms concave worker burnout ($p = 0.038$) and protective associations between meaningful work and depressive anhedonia ($\hat{\beta} = -0.135, p = 7.58 \times 10^{-7}$), demonstrates a 32.3% dissipation reduction under geodesic rollouts, and shows that MWF-aligned policy raises both aggregate output ($+9.3\%$) and bottom-decile work meaning ($+62.4\%$).**

***Index Terms*—Human–machine systems, socio-technical optimization, Master Work Function, supermodularity, stochastic thermodynamics, Jarzynski equality, generative AI, robotics, meaningful work, labor economics.**

I. Introduction

In the physical sciences, *work* measures energy transferred along a trajectory by an applied force; in human life, *work* denotes purposeful activity through which individuals cultivate mastery, contribute to collective needs, and realize intrinsic meaning. Across industrial robotics, enterprise software, and generative artificial intelligence (AI), however, contemporary automation design exhibits a stark asymmetry between these two dimensions. Machine subsystems are specified through explicit scalar objectives—latency, token throughput, defect rate, return on capital—whereas the human side of the socio-technical loop is typically constrained only by physical ergonomics or post-hoc turnover metrics. Early cyberneticists warned directly against this imbalance: Wiener [1], [2] observed that forcing human operators to compete with or conform to rigid automatic machinery reduces human labor to "the precise economic equivalent of slave labor," echoing foundational critiques of alienated labor, instrumental rationality, and the erosion of meaningful action by Marx [3], Arendt [4], and Heidegger [5], as well as Rawls's insistence that a just society must secure "meaningful work in free association with others" [6], [7]. In human-factors engineering, Bainbridge [8] and Parasuraman *et al.* [9] similarly documented the "ironies of automation," wherein relegating skilled operators to passive supervisory monitoring induces skill atrophy and vigilance failure precisely when exceptional intervention is needed.

Recent empirical evidence from the diffusion of large language models (LLMs) and physical robotics has made this design tension quantitatively acute. Following Frey and Osborne's foundational benchmark of occupational susceptibility to computerization [10], task-based labor models developed by Autor [11] and Acemoglu and Restrepo [12]–[14] demonstrate that aggregate welfare depends critically on whether new technologies merely displace labor in existing tasks ("so-so automation") or reinstate human comparative advantage in complementary tasks. Occupational exposure studies [15]–[17] and field experiments [18], [19] show that generative AI can compress skill gaps and raise short-run productivity, yet real-world telemetry from millions of interactions in the Anthropic Economic Index [20]–[22] and recent delegation analyses [23], [24] reveal a sharp divergence between collaborative human-in-the-loop iteration and autonomous, directive task execution. Parallel task-level assessments of physical robotics [25] indicate substantial automation potential in structured warehousing, manufacturing, and material-moving environments where worker autonomy is already compressed.

Despite these advances, three methodological gaps separate economic task models, organizational psychology, and control engineering:

1) *Absence of a unified cardinal potential:* Organizational psychology measures intrinsic motivation, autonomy, and meaningful work through validated psychometric instruments such as the Job Characteristics Model [26], Self-Determination Theory [27], the Utrecht Work Engagement Scale [28], and the Work and Meaning Inven-

K. Patel is with Centralized Autonomous Organization, The Utility Company, Albuquerque, NM, USA (e-mail: founders@theutilitycompany.co). Reproducible Python econometrics, Langevin simulation scripts, and public dataset manifests accompany this article.

tory [29], but rarely embeds them inside formal control or resource-allocation problems. Conversely, engineering controllers optimize operational throughput without state variables for human mastery or fulfillment.

2) *Static comparative statics versus path-dependent rollout friction:* Even when a reorganized workplace exhibits higher steady-state welfare, rapid or uncoordinated technology rollouts impose severe transient reorganization costs—cognitive overload, retraining bottlenecks, and institutional resistance—that path-independent deterministic calculus cannot capture.
3) *Lack of multi-scale structural estimation:* Existing empirical studies either evaluate macro-level task exposure [16], [17] without worker-level fulfillment and mental-health microdata, or analyze isolated workplace experiments [18], [19] without estimating a structural interaction matrix across cognitive AI, interface architecture, and physical robotics.

This paper bridges these three gaps and establishes a complete theoretical, philosophical, domain-engineering, and empirical foundation for the Master Work Function (MWF). Specifically, our contributions are fourfold:

1) *Axiomatic state space and canonical multi-agent decomposition (Section II):* We formalize the socio-technical state manifold $X = X_H \times X_M$, characterize generalized socio-technical forces via the total differential dW, distinguish instantaneous welfare flows from discounted intertemporal value stocks, invoke Debreu–von Neumann–Morgenstern representation conditions [30], [31], and prove uniqueness of an anchored functional decomposition [32] that generalizes cleanly to multi-human, multi-agent organizational hypergraphs.
2) *Burnout-regulated geometry, canonical limiting regimes, and supermodular comparative statics (Section III):* Correcting unboundedness in unweighted logarithmic-bilinear forms, we introduce quadratic cognitive-overload and rigid-pacing regulation, prove strict global concavity via Schur-complement spectral bounds on the cross-partial interaction matrix Γ, analyze four canonical human–machine limiting regimes alongside real-time gradient-ascent load balancing, and apply Topkis [33] and Milgrom–Shannon [34] lattice theory to classify augmentative ($\Gamma_{ij} > 0$) versus crowd-out ($\Gamma_{ij} < 0$) regimes.
3) *Variational/Pontryagin control and exact nonequilibrium fluctuation theorems (Section IV):* We formulate intertemporal co-investment via variational Hamilton–Jacobi–Bellman and Pontryagin optimality [35] and model noisy organizational transitions as open dissipative structures [36] governed by overdamped Langevin dynamics [37]–[39]. Defining protocol reorganization cost along time-varying deployment schedules $\lambda(t)$, we derive an exact socio-technical Jarzynski equality [40], Crooks fluctuation theorem [41], and Sivak–Crooks Riemannian geodesic schedule [42] minimizing transition dissipation.
4) *Normative mechanism design, domain instantiations, and multi-scale empirical validation (Sections V–VII):* After grounding the MWF in phenomenology, Aristotelian eudaimonia, unpaid care valuation, Foucauldian governmentality, Atkinson–Rawls inequality-averse welfare [6], [43], [44], and Holmstrom–Milgrom multitask Goodhart bounds [45], [46], we formulate domain controllers across collaborative/assistive/social robotics, value-aligned agentic AI, triadic and vocational pedagogy, safety-critical engineering, and smart cyber-physical infrastructure (Section VI). We then estimate the structural MWF across 873 U.S. occupations (124.0 million workers) and $10{,}483$ individual workers in the NORC General Social Survey (2018–2024) [47]–[49], testing every core theoretical prediction and evaluating post-scarcity macroeconomic governance instruments.

The remainder of the paper is organized as follows. Section II develops the axiomatic representation and multi-agent decomposition. Section III analyzes Hessian concavity, canonical limiting cases, adaptive gradient load-balancing, and supermodularity. Section IV presents deterministic optimal control and stochastic fluctuation theorems. Section V addresses philosophical foundations, governmentality, distributional equity, multimodal measurement, and incentive compatibility. Section VI formulates applications across robotics, agentic AI, education, safety-critical engineering, and smart infrastructure. Section VII reports the multi-scale empirical validation, macroeconomic counterfactuals, and post-scarcity policy instruments, and Section VIII discusses methodological limitations, open mathematical frontiers, field-trial protocols, and conclusions.

II. Axiomatic Socio-Technical State Space and Canonical Decomposition

A. State Manifold, Generalized Forces, and Stock–Flow Accounting

Consider a socio-technical work unit—an occupation, team, or cyber-physical cell—observed at time $t \geq 0$. The joint state is represented by a vector $x = (h, m) \in X \subseteq \mathbb{R}_+^{n_H} \times \mathbb{R}_+^{n_M}$, where $X_H \subseteq \mathbb{R}_+^{n_H}$ denotes the *human agency and engagement space* and $X_M \subseteq \mathbb{R}_+^{n_M}$ denotes the *machine capability and allocation space*. Concretely, coordinates of $h = (h_1, \ldots, h_{n_H})^\top$ measure active skill exercise and mastery, decision-making autonomy, creative originality, and prosocial relational engagement, while coordinates of $m = (m_1, \ldots, m_{n_M})^\top$ measure cognitive AI capability, collaborative interface transparency, rigid equipment-dictated pacing, and physical robotic actuation.

When $W : X \to \mathbb{R}$ is smooth on X, an infinitesimal shift (dh, dm) in human engagement and machine allocation induces the total differential

$$dW = \sum_{i=1}^{n_H} \frac{\partial W}{\partial h_i}\, dh_i + \sum_{j=1}^{n_M} \frac{\partial W}{\partial m_j}\, dm_j, \tag{1}$$

where the partial gradients $\nabla_h W = (\partial_{h_1} W, \ldots, \partial_{h_{n_H}} W)^\top$ and $\nabla_m W = (\partial_{m_1} W, \ldots, \partial_{m_{n_M}} W)^\top$ act as *generalized*

socio-technical forces quantifying the marginal welfare sensitivity to human effort versus machine intervention. To maintain dimensional consistency across static and dynamic analyses, we distinguish strictly between flow and stock quantities. The instantaneous *Master Work Function* $W : X \to \mathbb{R}$ measures the rate of joint socio-technical value generation per unit time (in normalized cardinal welfare units per period). Over a horizon $[0, T]$ with discount rate $\rho > 0$, the corresponding *intertemporal socio-technical value functional* is the time integral

$$\mathcal{V}[x(\cdot)] = \int_0^T e^{-\rho t} W\big(h(t), m(t)\big)\, dt, \tag{2}$$

whose units are cumulative discounted welfare. Static equilibrium analysis studies stationary points of the flow $W(h, m)$ at fixed structural parameters, whereas dynamic optimal control (Section IV) optimizes $\mathcal{V}$ subject to differential laws governing skill accumulation, atrophy, and capital deployment.

B. Representation Axioms

Let $\succeq$ denote a social or organizational planner's preference relation over probability lotteries $\mathcal{L}(X)$ on compact socio-technical state space $X = [0, \bar{h}_1] \times \cdots \times [0, \bar{m}_{n_M}]$.

Assumption 1 (Axiomatic Regularity of Socio-Technical Preferences)**.** The preference relation $\succeq$ on $\mathcal{L}(X)$ satisfies:

1) *Completeness and Transitivity:* $\succeq$ is a weak order on $\mathcal{L}(X)$.
2) *Topological Continuity:* For every $P \in \mathcal{L}(X)$, upper and lower contour sets $\{Q \in \mathcal{L}(X) : Q \succeq P\}$ and $\{Q \in \mathcal{L}(X) : P \succeq Q\}$ are closed in the weak topology.
3) *vNM Independence:* For all $P, Q, R \in \mathcal{L}(X)$ and $\alpha \in (0, 1)$, $P \succeq Q \iff \alpha P + (1-\alpha)R \succeq \alpha Q + (1-\alpha)R$.
4) *Smoothness and Regulated Saturation:* Deterministic states admit a twice continuously differentiable (C^2) utility representation $W(h, m)$ whose marginal returns are positive at low utilization ($\partial_{h_i} W(0, m_0) > 0$ for intrinsic agency dimensions) and strictly diminishing under cognitive overload or rigid machine pacing.

By the Debreu [30] and von Neumann–Morgenstern [31] representation theorems, Assumption 1 guarantees the existence of a C^2 cardinal potential $W : X \to \mathbb{R}$ such that $P \succeq Q \iff \mathbb{E}_P[W(h, m)] \geq \mathbb{E}_Q[W(h, m)]$, unique up to positive affine transformations $\widetilde{W} = aW + b$ with $a > 0$ and $b \in \mathbb{R}$. Crucially, all dimensionless ratios of marginal rates of substitution, signs of cross-partial derivatives, Hessian definiteness properties, and normalized fluctuation identities (once temperature K scales by a) are strictly invariant under positive affine rescaling.

C. Canonical Anchored Decomposition

Naively writing $W(h, m) = W_H(h) + W_M(m) + W_{\text{int}}(h, m)$ without boundary normalization leaves undetermined gauge functions $g_H(h)$ and $g_M(m)$ that could be added to W_H, W_M and subtracted from W_{int} without altering W. Fixing a reference baseline $(h_0, m_0) \in X$ (for example, unassisted baseline craft $m_0 = 0$ at standard human engagement h_0) yields an exact, unique anchored projection [32].

Theorem 1 (Unique Anchored Socio-Technical Decomposition)**.** *Let $W \in C^2(X)$ and fix an anchor state $(h_0, m_0) \in X$. There exists a unique quadruple $(W_0, W_H, W_M, W_{\text{int}})$ satisfying*

$$W(h, m) = W_0 + W_H(h) + W_M(m) + W_{\text{int}}(h, m) \tag{3}$$

subject to the boundary annihilation conditions

$$\begin{aligned} W_H(h_0) &= 0, \qquad W_M(m_0) = 0, \\ W_{\text{int}}(h_0, m) &= 0, \quad W_{\text{int}}(h, m_0) = 0 \end{aligned} \tag{4}$$

for all $(h, m) \in X$. Moreover, the components are given explicitly by

$$W_0 = W(h_0, m_0), \tag{5}$$

$$W_H(h) = W(h, m_0) - W(h_0, m_0), \tag{6}$$

$$W_M(m) = W(h_0, m) - W(h_0, m_0), \tag{7}$$

$$\begin{aligned} W_{\text{int}}(h, m) = {} & W(h, m) - W(h, m_0) \\ & - W(h_0, m) + W(h_0, m_0). \end{aligned} \tag{8}$$

Proof. Summing (5)–(8) cancels the intermediate terms $\pm W(h, m_0)$, $\pm W(h_0, m)$, and $\pm W(h_0, m_0)$, yielding $W(h, m)$ identically and satisfying (4) by direct substitution. For uniqueness, suppose $(\widetilde{W}_0, \widetilde{W}_H, \widetilde{W}_M, \widetilde{W}_{\text{int}})$ also satisfies (3)–(4). Evaluating (3) at (h_0, m_0) gives $\widetilde{W}_0 = W(h_0, m_0) = W_0$. Evaluating at (h, m_0) yields $W(h, m_0) = W_0 + \widetilde{W}_H(h) + 0 + 0$, whence $\widetilde{W}_H(h) = W_H(h)$. Symmetrically, evaluating at (h_0, m) forces $\widetilde{W}_M(m) = W_M(m)$, which by subtraction uniquely fixes $\widetilde{W}_{\text{int}}(h, m) = W_{\text{int}}(h, m)$. □

Equation (8) has a transparent differential-geometric interpretation as a mixed Riemann integral of the cross-partial interaction tensor:

$$\begin{aligned} W_{\text{int}}(h, m) = {} & \sum_{i=1}^{n_H} \sum_{j=1}^{n_M} \int_{h_{0,i}}^{h_i} \int_{m_{0,j}}^{m_j} \\ & \times \frac{\partial^2 W}{\partial u_i \, \partial v_j}\big(\tilde{h}^{(i)}(u_i), \tilde{m}^{(j)}(v_j)\big)\, dv_j \, du_i, \end{aligned} \tag{9}$$

along coordinate-wise paths connecting (h_0, m_0) to (h, m). Consequently, $W_{\text{int}} \equiv 0$ on X if and only if the cross-partial matrix $\Gamma(h, m) := \nabla_h \nabla_m^\top W(h, m) \in \mathbb{R}^{n_H \times n_M}$ vanishes identically. Whenever $\Gamma(h, m) \neq 0$, socio-technical welfare is irreducibly non-separable: evaluating human labor policies independently of machine interface architecture (or vice versa) commits a structural misspecification error quantified empirically in Section VII.

D. Multi-Agent Organizational Hypergraph Generalization

Real organizations couple multiple human workers $\mathcal{H} = \{1, \ldots, N_H\}$ and multiple heterogeneous machine agents or robotic platforms $\mathcal{M} = \{1, \ldots, N_M\}$. Let $\mathcal{N} = \mathcal{H} \cup \mathcal{M}$ index all agents with state profile $\boldsymbol{z} = (z_1, \ldots, z_{|\mathcal{N}|}) \in \prod_{k \in \mathcal{N}} Z_k$ and anchor $\boldsymbol{z}_0$. Applying anchored Möbius inversion on the

subset lattice $2^{\mathcal{N}}$ [32] yields the organizational hypergraph expansion

$$\begin{aligned} W(\boldsymbol{z}) &= \sum_{S\subseteq\mathcal{N}} W_S(\boldsymbol{z}_S), \\ W_S(\boldsymbol{z}_S) &= \sum_{T\subseteq S} (-1)^{|S\setminus T|} W(\boldsymbol{z}_T, \boldsymbol{z}_{0,-T}), \end{aligned} \tag{10}$$

where $(\boldsymbol{z}_T, \boldsymbol{z}_{0,-T})$ sets coordinates outside T to their anchor values. Truncating (10) at order $|S| \le 3$ separates individual human mastery ($|S| = 1, S \subset \mathcal{H}$), standalone machine throughput ($|S| = 1, S \subset \mathcal{M}$), dyadic human–human teamwork ($S = \{a, a'\} \subset \mathcal{H}$), dyadic human–AI complementarity ($S = \{a, b\} \in \mathcal{H}\times\mathcal{M}$), and triadic mediation such as teacher–student–AI scaffolding ($|S| = 3$, analyzed in Section VI).

III. Differential Geometry, Burnout Regulation, and Supermodular Optimization

A. Resolving Unboundedness via Cognitive-Overload and Pacing Regulation

To obtain tractable analytical and econometric specifications, consider parametric families of (3). A natural first-pass candidate is the unregulated logarithmic-bilinear form $W_{\text{unreg}}(h, m) = A\ln(1 + h) + B\ln(1 + m) + Chm$ with $A, B, C > 0$. Inspection of its gradient reveals why unregulated forms are structurally inadequate: for $C > 0$ and $(h, m) \in \mathbb{R}^2_+$, both first derivatives $\partial_h W_{\text{unreg}} = A/(1 + h) + Cm > 0$ and $\partial_m W_{\text{unreg}} = B/(1 + m) + Ch > 0$ are strictly positive everywhere, so no unconstrained interior stationary point exists. Furthermore, its Hessian determinant $\det H_{W_{\text{unreg}}} = \frac{AB}{(1+h)^2(1+m)^2} - C^2$ turns negative at moderate (h, m), producing an unstable saddle rather than a concave welfare peak.

Substantively, unweighted logarithmic growth omits a fundamental physiological and organizational law: beyond moderate intensity, cognitive workload induces exhaustion and burnout [28], while excessive machine autonomy or equipment-dictated pacing crowds out human situational awareness and induces coordination congestion [8]. Incorporating quadratic overload and congestion penalties yields the *Burnout-Regulated Master Work Function*:

$$\begin{aligned} W(h, m) = W_0 &+ \sum_{i=1}^{n_H} \Big[\alpha_i \ln(1 + h_i) - \frac{\kappa_i}{2} h_i^2\Big] \\ &+ \sum_{j=1}^{n_M} \Big[\beta_j \ln(1 + m_j) - \frac{\nu_j}{2} m_j^2\Big] + h^\top \Gamma m, \end{aligned} \tag{11}$$

where $\alpha_i > 0$ measures intrinsic human fulfillment elasticity, $\beta_j \in \mathbb{R}$ captures baseline machine utility (positive for drudgery-relieving assistance, negative for intrusive surveillance or rigid substitution), $\kappa_i > 0$ and $\nu_j > 0$ are burnout and congestion curvature parameters, and $\Gamma = [\Gamma_{ij}] \in \mathbb{R}^{n_H\times n_M}$ is the socio-technical complementarity matrix.

Theorem 2 (Global Concavity and Unique Regulated Equilibrium)**.** *Consider the regulated Master Work Function* (11) *on* $X = \mathbb{R}^{n_H}_+ \times \mathbb{R}^{n_M}_+$ *with* $\alpha_i \ge 0$*,* $\kappa_i > 0$*,* $\beta_j \in \mathbb{R}$*, and* $\nu_j > 0$*. Let* $\kappa_{\min} \coloneqq \min_i \kappa_i$*,* $\nu_{\min} \coloneqq \min_j \nu_j$*, and* $\|\Gamma\|_2 \coloneqq \sigma_{\max}(\Gamma)$ *denote the spectral norm of* Γ*. If*

$$\|\Gamma\|_2 < \sqrt{\kappa_{\min}\,\nu_{\min}}, \tag{12}$$

then:

1) $W(h, m)$ *is strictly jointly concave and coercive* ($W(h, m) \to -\infty$ *as* $\|(h, m)\| \to \infty$) *across* $\mathbb{R}^{n_H}_+ \times \mathbb{R}^{n_M}_+$.
2) *Over any nonempty compact convex organizational feasible set* $\mathcal{F} = \{(h, m) \in X : g_r(h, m) \le B_r,\ r = 1, \dots, m_c\}$ *with convex constraint functions* $g_r \in C^1$*, there exists a unique global maximizer* $(h^*, m^*) \in \mathcal{F}$*, for which the Karush–Kuhn–Tucker (KKT) conditions*

$$\begin{aligned} &\frac{\alpha_i}{1 + h_i^*} - \kappa_i h_i^* + (\Gamma m^*)_i \\ &\qquad - \sum_{r=1}^{m_c} \mu_r^* \frac{\partial g_r}{\partial h_i} + \eta^*_{H,i} = 0, \end{aligned} \tag{13}$$

$$\begin{aligned} &\frac{\beta_j}{1 + m_j^*} - \nu_j m_j^* + (\Gamma^\top h^*)_j \\ &\qquad - \sum_{r=1}^{m_c} \mu_r^* \frac{\partial g_r}{\partial m_j} + \eta^*_{M,j} = 0, \end{aligned} \tag{14}$$

   with $\mu_r^*, \eta^*_{H,i}, \eta^*_{M,j} \ge 0$ *and complementary slackness, are both necessary and sufficient.*
3) *If constraints are slack and* $\alpha_i + (\Gamma m^*)_i > 0$*,* $\beta_j + (\Gamma^\top h^*)_j > 0$ *for all* i, j*, then* (h^*, m^*) *lies strictly in the positive orthant* $\mathbb{R}^{n_H}_{++} \times \mathbb{R}^{n_M}_{++}$.

Proof. Differentiating (11) twice yields the block Hessian

$$H_W(h, m) = \begin{bmatrix} -D_H(h) & \Gamma \\ \Gamma^\top & -D_M(m) \end{bmatrix}, \tag{15}$$

where $D_H(h) = \operatorname{diag}\big(\kappa_i + \alpha_i/(1 + h_i)^2\big)$ and $D_M(m) = \operatorname{diag}\big(\nu_j + \beta_j/(1 + m_j)^2\big)$. Notice that when $\beta_j \ge 0$, $D_{M,jj}(m_j) \ge \nu_j$; if any $\beta_j < 0$ (substitutive machine disutility), replacing $\nu_{\min}$ in (12) by $\tilde{\nu}_{\min} \coloneqq \min_j \inf_{m_j\ge 0}[\nu_j + \beta_j/(1 + m_j)^2] > 0$ preserves uniform positive definiteness $D_H(h) \succeq \kappa_{\min} I_{n_H} \succ 0$ and $D_M(m) \succeq \tilde{\nu}_{\min} I_{n_M} \succ 0$. For negative definiteness of $H_W(h, m)$, $-H_W(h, m) \succ 0$ holds by the Schur complement lemma if and only if $D_H(h) \succ 0$ and $S(h, m) \coloneqq D_M(m) - \Gamma^\top D_H(h)^{-1}\Gamma \succ 0$. For any unit vector $v \in \mathbb{R}^{n_M}$,

$$\begin{aligned} v^\top S(h, m)v &= v^\top D_M(m)v - (\Gamma v)^\top D_H(h)^{-1}(\Gamma v) \\ &\ge \nu_{\min}\|v\|_2^2 - \frac{1}{\kappa_{\min}}\|\Gamma v\|_2^2 \\ &\ge \nu_{\min} - \frac{\|\Gamma\|_2^2}{\kappa_{\min}} > 0 \end{aligned}$$

by (12). Thus $H_W(h, m) \prec 0$ uniformly, establishing strict joint concavity and quadratic coercivity since $\ln(1 + z)$ grows sublinearly relative to $-z^2/2$. Strict concavity over convex $\mathcal{F}$ implies uniqueness of (h^*, m^*) and sufficiency of the KKT conditions (13)–(14) under Slater's condition. Finally, at $h_i = 0$ with slack resource constraints, $\partial_{h_i} W(0, m^*) = \alpha_i + (\Gamma m^*)_i > 0$ rules out boundary optima $\eta^*_{H,i} > 0$, forcing $h_i^* > 0$ (and symmetrically $m_j^* > 0$). □

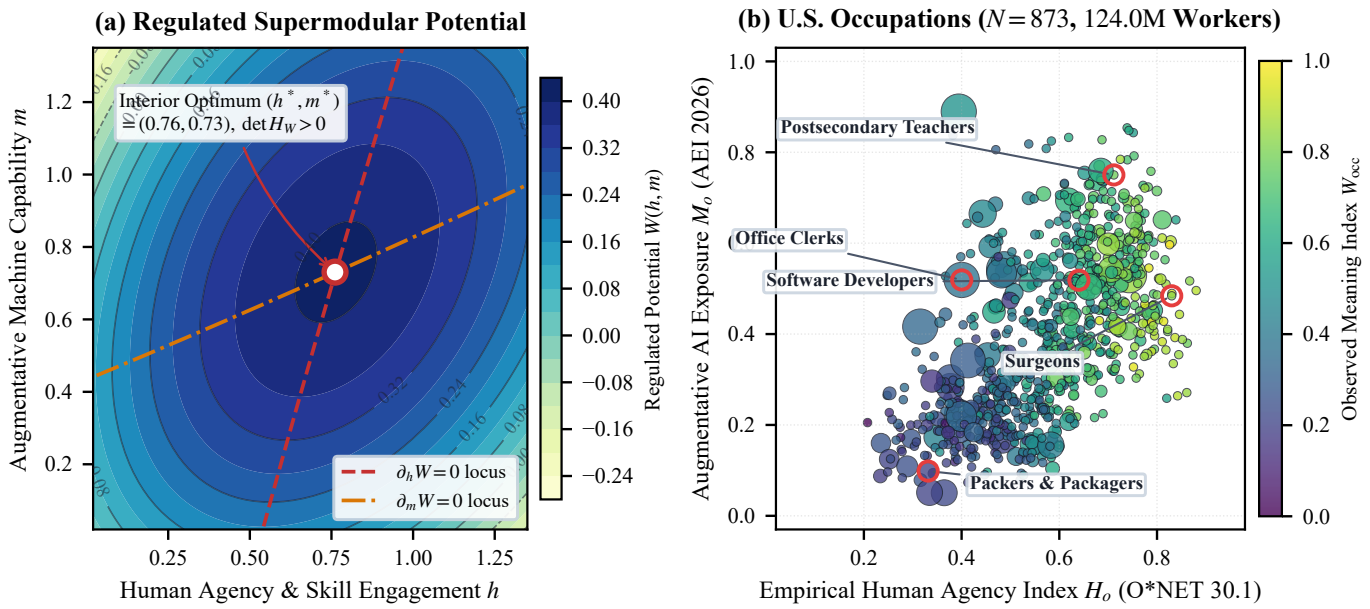

Fig. 1. Theoretical and empirical Master Work Function landscapes. (a) Burnout-regulated supermodular potential (11) exhibiting intersecting reaction loci ($\partial_h W = 0$, $\partial_m W = 0$) and a unique globally concave interior optimum $(h^*, m^*) = (0.76, 0.73)$ with $\det H_W > 0$. (b) Observed distribution of $N = 873$ U.S. occupations (124.0 million BLS workers) across composite O*NET 30.1 human agency H_o and Anthropic Economic Index augmentative AI exposure M_o, colored by normalized meaningful work potential W_{occ} and sized proportionally to national employment.

Fig. 1(a) illustrates Theorem 2: unlike unbounded bilinear forms, the regulated reaction curves $\partial_h W = 0$ and $\partial_m W = 0$ intersect at a well-defined, globally stable interior peak (h^*, m^*).

B. Canonical Limiting Regimes and Adaptive Gradient-Ascent Load Balancing

Evaluating the anchored decomposition (3) and its regulated parameterization (11) across boundary and interior quadrants recovers four canonical socio-technical regimes:

1) *Case 1 — Purely Human Work ($m = m_0 = 0$):* By boundary annihilation (4), $W_M(0) = 0$ and $W_{\text{int}}(h, 0) = 0$, reducing the potential to $W(h, 0) = W_0 + W_H(h) = W_0 + \sum_{i=1}^{n_H} [\alpha_i \ln(1 + h_i) - \frac{\kappa_i}{2} h_i^2]$. Starting from forced idleness ($h = 0$), marginal human value $\partial_{h_i} W(0, 0) = \alpha_i > 0$ rises steeply as workers exercise craft, autonomy, and social purpose, attains a finite physiological peak $h^*_{i,\text{solo}} > 0$, and turns negative ($\partial_{h_i} W < 0$) when unassisted drudgery or excessive hours trigger cognitive exhaustion.
2) *Case 2 — Full Machine Automation ($h = h_0 = 0$):* Removing active human participation yields $W(0, m) = W_0 + W_M(m) = W_0 + \sum_{j=1}^{n_M} [\beta_j \ln(1 + m_j) - \frac{\nu_j}{2} m_j^2]$. Full automation generates bounded instrumental utility—which is socially desirable in hazardous, toxic, or purely menial tasks where human engagement carries disutility—yet lacks intrinsic eudaimonic fulfillment, empathetic care, and moral accountability.
3) *Case 3 — Synergistic Human–Machine Collaboration ($h > 0, m > 0, W_{\text{int}}(h, m) > 0$):* When collaborative interfaces and shared-autonomy robotics realize $h^\top \Gamma m > 0$, joint welfare is strictly superadditive ($W(h, m) - W_0 > W_H(h) + W_M(m)$). Machines absorb repetitive computational or biomechanical strain while amplifying human judgment, creativity, and interpersonal presence.
4) *Case 4 — Misaligned or Antagonistic Interaction ($h > 0, m > 0, W_{\text{int}}(h, m) < 0$):* When operators must continually monitor and override brittle, opaque automation ("babysitting the machine") or submit to rigid algorithmic pacing, negative cross-partials $h^\top \Gamma m < 0$ degrade situational awareness and autonomy. In severe misalignment, hybrid operation falls below standalone performance: $W(h, m) < \max\{W(h, 0), W(0, m)\}$.

Locally, the signs of the generalized forces $(\nabla_h W, \nabla_m W)$ provide a real-time diagnostic for adaptive task allocation. In an online shared-autonomy controller with adaptation gains $\eta_H, \eta_M > 0$, projected gradient ascent on X,

$$\begin{aligned} \dot{h}(t) &= \eta_H \nabla_h W\big(h(t), m(t)\big), \\ \dot{m}(t) &= \eta_M \nabla_m W\big(h(t), m(t)\big), \end{aligned} \tag{16}$$

dynamically rebalances workload: whenever an operator is under-challenged by excessive automation ($\partial_{h_i} W > 0, \partial_{m_j} W < 0$), machine intervention recedes to restore human engagement; conversely, when cognitive overload sets in ($\partial_{h_i} W < 0, \partial_{m_j} W > 0$), assistive automation steps up to relieve strain. Under the spectral concavity bound (12), $V_{\text{Lyap}}(h, m) = W(h^*, m^*) - W(h, m)$ serves as a strict Lyapunov function guaranteeing global asymptotic convergence of (16) to (h^*, m^*).

C. Supermodularity and Monotone Comparative Statics

Beyond guaranteeing stability, the cross-partial matrix $\Gamma_{ij} = \partial^2 W / \partial h_i \partial m_j$ provides an operational classification of socio-technical regimes:

1) *Augmentative Synergy ($\Gamma_{ij} > 0$):* Machine capability m_j raises the marginal socio-technical return to human engagement h_i (e.g., transparent diagnostic tools or interactive copilot feedback amplifying expert skill).
2) *Neutral Modularity ($\Gamma_{ij} = 0$):* Back-office infrastructure scales output independently without altering marginal human fulfillment.
3) *Substitutive Crowding-Out ($\Gamma_{ij} < 0$):* Machine capability m_j depresses the marginal return to human agency h_i (e.g., algorithmic conveyor pacing stifling creative discretion, or physical robotics displacing relational craft).

Let $c_M(m;\tau)$ denote the organizational cost of deploying machine vector m indexed by technological progress parameter $\tau \in \mathbb{R}$, where higher τ reduces marginal machine deployment cost ($\partial^2 c_M/\partial m_j \partial \tau < 0$).

Proposition 1 (Monotone Comparative Statics of Technological Progress)**.** *Suppose an organization chooses $(h,m) \in X$ on a product lattice $X = X_H \times X_M$ to maximize net socio-technical surplus $\Pi(h,m;\tau) = W(h,m) - c_H(h) - c_M(m;\tau)$, where c_H, c_M are separable across coordinates and $\partial^2 c_M/\partial m_j \partial \tau < 0$.*

1) (Pure Augmentation Lattice): *If $\Gamma_{ij} \geq 0$ for all (i,j) with strict inequality for at least one pair, then Π is supermodular in (h,m) and exhibits strictly increasing differences in $(m;\tau)$. By Topkis's monotonicity theorem [33], [34], an exogenous advance in AI/robotic capability ($\tau' > \tau$) induces nondecreasing optimal machine adoption $m^*(\tau') \geq m^*(\tau)$* and *nondecreasing optimal human engagement $h^*(\tau') \geq h^*(\tau)$ (crowding-in).*
2) (Pure Substitution Dual Lattice): *If $\Gamma_{ij} \leq 0$ for all (i,j) with strict inequality for (i_0,j_0) at an interior optimum, reversing the order on X_H ($h \mapsto -h$) renders $\Pi(-h,m;\tau)$ supermodular in $(-h,m)$. Consequently, $\tau' > \tau$ yields $m^*(\tau') \geq m^*(\tau)$ accompanied by $h^*(\tau') \leq h^*(\tau)$, strictly eroding human engagement $h^*_{i_0}(\tau') < h^*_{i_0}(\tau)$ (crowding-out).*

Proof. Because $c_H(h)$ and $c_M(m;\tau)$ are coordinate-separable, $\partial^2\Pi/\partial h_i \partial h_k = 0$ for $i \neq k$, $\partial^2\Pi/\partial m_j \partial m_\ell = 0$ for $j \neq \ell$, and $\partial^2\Pi/\partial h_i \partial m_j = \Gamma_{ij}$. In Case 1, all off-diagonal cross-partials of Π on $X_H \times X_M$ are nonnegative, so Π is supermodular on the Euclidean lattice $X_H \times X_M$, while $\partial^2\Pi/\partial m_j \partial \tau = -\partial^2 c_M/\partial m_j \partial \tau > 0$ and $\partial^2\Pi/\partial h_i \partial \tau = 0$ establish increasing differences in $((h,m);\tau)$. Topkis's theorem [33, Theorem 2.8.1] yields $(h^*(\tau'), m^*(\tau')) \geq (h^*(\tau), m^*(\tau))$. In Case 2, under the dual partial order $\tilde{h} = -h$, $\partial^2\Pi/\partial \tilde{h}_i \partial m_j = -\Gamma_{ij} \geq 0$, so $(\tilde{h}^*(\tau), m^*(\tau))$ is nondecreasing in τ, meaning $h^*(\tau)$ is nonincreasing; interiority and $\Gamma_{i_0 j_0} < 0$ imply $\partial_{h_{i_0}}\Pi$ shifts strictly downward when $m^*_{j_0}$ rises, forcing $h^*_{i_0}(\tau') < h^*_{i_0}(\tau)$. □

Proposition 1 formalizes why aggregate AI exposure indices alone cannot predict worker outcomes: identical underlying model capability τ generates opposite trajectories for human skill and autonomy depending on whether the deployment architecture realizes $\Gamma_{ij} > 0$ or $\Gamma_{ij} < 0$.

IV. Socio-Technical Dynamics, Optimal Control, and Fluctuation Theorems

A. Intertemporal Variational, HJB, and Pontryagin Optimal Control

In dynamic organizational environments, human mastery $h(t)$ and machine capability $m(t)$ evolve endogenously in response to training investment $u_H(t) \geq 0$, automation investment $u_M(t) \geq 0$, and interface augmentation design $\sigma_{\text{aug}}(t) \in [0,1]$ (the fraction of human–machine interactions structured as iterative collaborative feedback rather than directive execution). Let quality-adjusted economic output be $Y(h,m,\sigma_{\text{aug}}) = y_H^\top h + y_M^\top m + \sigma_{\text{aug}} h^\top Q m - \mathcal{D}_{\text{err}}(h,m)$, where $\mathcal{D}_{\text{err}}(h,m) = \sum_i \xi_i m_i \max(0, \bar{h}_i - h_i)$ captures supervisory error drag when high machine autonomy operates alongside degraded human vigilance ($h_i < \bar{h}_i$) [8].

The state dynamics follow coupled differential equations:

$$\dot{h}_i(t) = u_{H,i}(t) - \delta_{\text{atr},i}\big(1 - \sigma_{\text{aug}}(t)\big) m_i(t) h_i(t) - \delta_{\text{burn},i}\Phi_{\text{ov},i}\big(h(t), m(t)\big) h_i(t), \tag{17}$$

$$\dot{m}_j(t) = u_{M,j}(t) - \delta_{M,j} m_j(t), \tag{18}$$

where $\delta_{\text{atr},i} > 0$ is the rate of directive skill atrophy, $\delta_{\text{burn},i} > 0$ governs overload-induced attrition above threshold $\Phi_{\text{ov},i}$, and $\delta_{M,j} > 0$ is technical depreciation. An organization weighting socio-technical welfare flow W by $\theta \in [0,1]$ and net output Y by $(1-\theta)$ maximizes

$$J[u_H, u_M, \sigma_{\text{aug}}] = \int_0^T e^{-\rho t}\Big[\theta W\big(h(t), m(t)\big) + (1-\theta) Y\big(h(t), m(t), \sigma_{\text{aug}}(t)\big) - \frac{c_H}{2}\|u_H(t)\|_2^2 - \frac{c_M}{2}\|u_M(t)\|_2^2\Big] dt. \tag{19}$$

On an infinite or stationary horizon, let $\ell(h, m, u_H, u_M, \sigma_{\text{aug}})$ denote the bracketed instantaneous surplus in (19). Variational optimality over state trajectories $(h(\cdot), m(\cdot))$ is characterized equivalently by the Euler–Lagrange equations on the reduced Lagrangian or by the dynamic-programming Hamilton–Jacobi–Bellman (HJB) equation for the value function $V(h,m)$:

$$\rho V(h,m) = \max_{u_H, u_M, \sigma_{\text{aug}}} \Big\{ \ell(h, m, u_H, u_M, \sigma_{\text{aug}}) + \nabla_h V \cdot \dot{h} + \nabla_m V \cdot \dot{m} \Big\}. \tag{20}$$

Introducing current-value costates $\psi_H(t) = \nabla_h V \in \mathbb{R}^{n_H}$ and $\psi_M(t) = \nabla_m V \in \mathbb{R}^{n_M}$ (representing the shadow prices of human mastery and machine capital), Pontryagin's Maximum Principle [35] yields optimal investment rules $u^*_{H,i}(t) = \psi_{H,i}(t)/c_H$ and $u^*_{M,j}(t) = \psi_{M,j}(t)/c_M$, governed by costate evolution

$$\dot{\psi}_{H,i} = \big[\rho + \delta_{\text{atr},i}(1-\sigma_{\text{aug}}) m_i + \delta_{\text{burn},i}\tilde{\Phi}_i\big]\psi_{H,i} - \theta\frac{\partial W}{\partial h_i} - (1-\theta)\frac{\partial Y}{\partial h_i}. \tag{21}$$

Equation (21) exposes a fundamental dynamic trap: a naive throughput controller ($\theta = 0$) observing low short-run marginal human output $\partial_{h_i} Y$ when $h_i(0) > \bar{h}_i$ sets $\psi_{H,i}$

near zero, starving training u_H while accelerating directive automation ($\sigma_{\text{aug}} \to 0$). Once atrophy (17) pushes $h_i(t)$ below the vigilance threshold $\bar{h}_i$, supervisory error drag $\mathcal{D}_{\text{err}}$ activates, causing long-run output $Y(t)$ to fall below that of an MWF-aligned controller ($\theta > 0$) that preserves human mastery throughout the transition.

B. Overdamped Langevin Rollouts and Nonequilibrium Fluctuation Theorems

Deterministic trajectories ignore behavioral heterogeneity, bounded rationality, and stochastic organizational shocks during technological transitions. Note also that defining transition work via a deterministic state-gradient integral $\int_0^\tau \nabla_x W(x_t) \cdot \dot{x}_t\, dt$ would, by the Fundamental Theorem of Calculus, collapse identically to $W(x_\tau) - W(x_0)$ along every smooth curve—producing zero path dependence and zero rollout friction.

As Prigogine [36] established for open, far-from-equilibrium dissipative structures, sustaining coherent order across structural reorganizations entails irreversible entropy production. To formulate a mathematically consistent nonequilibrium theory [37], [39]–[41], let $\lambda(t) \in \mathbb{R}^p$ denote an externally controlled *organizational rollout protocol* over $t \in [0, \tau]$ (such as a phased enterprise deployment schedule shifting target automation allocation from $\lambda(0) = \lambda_A$ to $\lambda(\tau) = \lambda_B$). At fixed protocol λ, individual teams or workers $x_t = (h_t, m_t) \in \mathbb{R}^d$ ($d = n_H + n_M$) climb gradients of the socio-technical potential $W(x; \lambda)$ subject to idiosyncratic behavioral and operational noise:

$$dx_t = \mu \nabla_x W\big(x_t; \lambda(t)\big)\, dt + \sqrt{2K\mu}\, dB_t, \tag{22}$$

where $\mu > 0$ is organizational adaptation mobility, B_t is standard d-dimensional Brownian motion, and $K > 0$ is the *socio-technical fluctuation scale* (behavioral temperature [38], calibrated empirically from within-occupation dispersion in Section VII). Defining effective socio-technical energy $E(x; \lambda) \coloneqq -W(x; \lambda)$, when λ is held constant and $W(\cdot; \lambda)$ is coercive (Theorem 2), the Fokker–Planck equation associated with (22) relaxes to a unique Boltzmann–Gibbs stationary equilibrium:

$$\begin{aligned} p_{\text{eq}}(x; \lambda) &= \frac{1}{Z(\lambda)} \exp\!\left(\frac{W(x; \lambda)}{K}\right), \\ Z(\lambda) &= \int_{\mathbb{R}^d} \exp\!\left(\frac{W(x; \lambda)}{K}\right) dx. \end{aligned} \tag{23}$$

Definition 1 (Macroscopic Master Work Potential and Reorganization Cost)**.** The *Macroscopic Master Work Potential* at protocol state λ is the log-partition functional

$$\begin{aligned} \Phi(\lambda) &\coloneqq K \ln Z(\lambda) \\ &= \mathbb{E}_{p_{\text{eq}}(\cdot;\lambda)}\big[W(X; \lambda)\big] + K\mathcal{S}\big(p_{\text{eq}}(\cdot; \lambda)\big), \end{aligned} \tag{24}$$

where $\mathcal{S}(p) = -\int p(x) \ln p(x)\, dx$ is Shannon organizational diversity entropy (so $\Phi(\lambda) = -F(\lambda)$ is the exact dual of Helmholtz free energy). Along a finite-time stochastic trajectory $\gamma = \{x_t\}_{t \in [0,\tau]}$ driven by protocol $\lambda(t)$, the *Protocol Reorganization Cost* is the Sekimoto–Jarzynski path functional

$$\begin{aligned} C_{\text{reorg}}[\gamma; \lambda] &\coloneqq \int_0^\tau \frac{\partial E(x_t; \lambda_t)}{\partial \lambda} \cdot \dot{\lambda}_t\, dt \\ &= -\int_0^\tau \frac{\partial W(x_t; \lambda_t)}{\partial \lambda} \cdot \dot{\lambda}_t\, dt. \end{aligned} \tag{25}$$

By Itô's formula (Stratonovich chain rule for total energy differential $dE = \nabla_x E \circ dx_t + \partial_\lambda E \cdot \dot{\lambda}_t dt$), total welfare change decomposes into $W(x_\tau; \lambda_B) - W(x_0; \lambda_A) = -C_{\text{reorg}}[\gamma; \lambda] - Q_{\text{diss}}[\gamma]$, where Q_{diss} is organizational heat dissipated into friction and turnover during adaptation. Unlike the naive gradient integral, $C_{\text{reorg}}[\gamma; \lambda]$ depends genuinely on the rollout speed $\dot{\lambda}_t$ and finite-time lag between $p_t(x)$ and $p_{\text{eq}}(x; \lambda_t)$.

Theorem 3 (Socio-Technical Jarzynski Equality and Crooks Fluctuation Relation)**.** *Suppose an organizational ensemble is initialized in stationary equilibrium $p_{\text{eq}}(x; \lambda_A)$ at $t = 0$ and driven over $t \in [0, \tau]$ by an arbitrary smooth protocol $\lambda : [0, \tau] \to \mathbb{R}^p$ with $\lambda(0) = \lambda_A, \lambda(\tau) = \lambda_B$ under* (22)*. Let $\tilde{\lambda}(t) \coloneqq \lambda(\tau - t)$ denote the time-reversed rollback protocol initialized in $p_{\text{eq}}(x; \lambda_B)$, and let $P_{\text{fwd}}(c)$ and $P_{\text{rev}}(c)$ denote the probability densities of C_{reorg} under λ and $\tilde{\lambda}$, respectively. Then:*

1) (Crooks Symmetry): *For every reorganization cost level $c \in \mathbb{R}$,*

$$\frac{P_{\text{fwd}}(c)}{P_{\text{rev}}(-c)} = \exp\!\left(\frac{c + \Delta\Phi}{K}\right), \tag{26}$$

where $\Delta\Phi \coloneqq \Phi(\lambda_B) - \Phi(\lambda_A)$.

2) (Socio-Technical Jarzynski Equality): *Regardless of how rapidly $\lambda(t)$ is executed, the exponential average of reorganization cost recovers the exact equilibrium potential gain:*

$$\mathbb{E}_{\text{fwd}}\left[\exp\!\left(-\frac{C_{\text{reorg}}[\gamma; \lambda]}{K}\right)\right] = \exp\!\left(\frac{\Delta\Phi}{K}\right). \tag{27}$$

3) (Second-Law Reorganization Friction Bound): *Expected reorganization dissipation $W_{\text{diss}} \coloneqq \mathbb{E}_{\text{fwd}}[C_{\text{reorg}}] + \Delta\Phi$ satisfies*

$$W_{\text{diss}} = K\, D_{\text{KL}}\big(\mathcal{P}_{\text{fwd}} \,\|\, \mathcal{P}^\dagger_{\text{rev}}\big) \geq 0, \tag{28}$$

or equivalently $\mathbb{E}_{\text{fwd}}[C_{\text{reorg}}] \geq -\Delta\Phi$, with equality if and only if the rollout is quasi-static ($\tau \to \infty$, $\dot{\lambda} \to 0$).

Proof. Partition $[0, \tau]$ into M steps $t_k = k\Delta t$ ($\Delta t = \tau/M$) with $\lambda_k = \lambda(t_k)$, alternating discrete protocol updates $\lambda_k \to \lambda_{k+1}$ at fixed state x_k with Markov transition kernels $\pi_{\lambda_{k+1}}(x_k \to x_{k+1})$ satisfying detailed balance $p_{\text{eq}}(x; \lambda)\pi_\lambda(x \to x') = p_{\text{eq}}(x'; \lambda)\pi_\lambda(x' \to x)$ for (22). Along discrete path $\boldsymbol{x} = (x_0, \ldots, x_M)$, forward path probability $p_F(\boldsymbol{x}) = p_{\text{eq}}(x_0; \lambda_0) \prod_{k=0}^{M-1} \pi_{\lambda_{k+1}}(x_k \to$

x_{k+1}) and time-reversed path probability $p_R(\boldsymbol{x}^\dagger) = p_{\mathrm{eq}}(x_M;\lambda_M)\prod_{k=0}^{M-1}\pi_{\lambda_{k+1}}(x_{k+1}\to x_k)$ satisfy

$$\begin{aligned}\frac{p_F(\boldsymbol{x})}{p_R(\boldsymbol{x}^\dagger)} &= \frac{p_{\mathrm{eq}}(x_0;\lambda_0)}{p_{\mathrm{eq}}(x_M;\lambda_M)}\prod_{k=0}^{M-1}\frac{p_{\mathrm{eq}}(x_{k+1};\lambda_{k+1})}{p_{\mathrm{eq}}(x_k;\lambda_{k+1})}\\ &= \frac{Z(\lambda_M)}{Z(\lambda_0)}\exp\Bigg(-\frac{1}{K}\sum_{k=0}^{M-1}\big[W(x_k;\lambda_{k+1})\\ &\qquad - W(x_k;\lambda_k)\big]\Bigg)\\ &= \exp\left(\frac{\Delta\Phi + C_{\mathrm{reorg}}^{(M)}(\boldsymbol{x})}{K}\right).\end{aligned}$$

Taking the continuum limit $M\to\infty$ ($\Delta t\to 0$) and marginalizing over paths with $C_{\mathrm{reorg}}[\gamma;\lambda] = c$ yields (26). Multiplying (26) by $P_{\mathrm{rev}}(-c)e^{-c/K}$ and integrating over $c\in\mathbb{R}$ gives $\int P_{\mathrm{fwd}}(c)e^{-c/K}dc = e^{\Delta\Phi/K}\int P_{\mathrm{rev}}(-c)dc = e^{\Delta\Phi/K}$, proving (27). Finally, taking path-space expectations of $\ln[p_F(\gamma)/p_R(\gamma^\dagger)] = (\Delta\Phi + C_{\mathrm{reorg}})/K$ yields (28) by nonnegativity of Kullback–Leibler divergence (equivalently, Jensen's inequality applied to (27)). □

Corollary 1 (Sivak–Crooks Minimum-Dissipation Rollout Geodesic)**.** *For a scalar rollout schedule $\lambda:[0,\tau]\to[0,1]$ executed over a sufficiently slow but finite deployment horizon τ, linear-response theory [42] approximates mean reorganization dissipation by the Riemannian energy functional*

$$\begin{aligned}W_{\mathrm{diss}}[\lambda] &\approx \int_0^\tau \zeta\big(\lambda(t)\big)\dot\lambda(t)^2\,dt,\\ \zeta(\lambda) &= \frac{1}{K}\int_0^\infty \mathrm{Cov}_\lambda\big(\partial_\lambda W(x_0;\lambda),\\ &\qquad \partial_\lambda W(x_s;\lambda)\big)\,ds,\end{aligned}\tag{29}$$

where Cov_λ denotes stationary covariance under $p_{\mathrm{eq}}(\cdot;\lambda)$. By the Cauchy–Schwarz inequality, $W_{\mathrm{diss}}[\lambda]$ is minimized over boundary conditions $\lambda(0)=0,\lambda(\tau)=1$ when protocol velocity scales inversely with the square root of socio-technical friction:

$$\begin{aligned}\dot\lambda^*(t) &= \frac{\mathcal{L}_{\mathrm{thermo}}}{\tau\sqrt{\zeta\big(\lambda^*(t)\big)}},\\ \mathcal{L}_{\mathrm{thermo}} &\coloneqq \int_0^1\sqrt{\zeta(\lambda)}\,d\lambda,\end{aligned}\tag{30}$$

achieving minimum dissipation $W_{\mathrm{diss}}^* = \mathcal{L}_{\mathrm{thermo}}^2/\tau$.

Corollary 1 delivers an actionable engineering rule for enterprise AI and robotics transitions: rather than imposing an abrupt "big-bang" cutover or a rigid linear rollout, management should measure role variance and adaptation lag $\zeta(\lambda)$ across pilot phases and deliberately slow deployment speed $\dot\lambda^*(t)$ across high-friction retraining bottlenecks.

V. Normative Foundations, Philosophy of Technology, Equity, and Mechanism Design

A. Phenomenology, Eudaimonia, Unpaid Work, and AI Ethics

Each structural component of the Master Work Function formalizes a distinct tradition in moral philosophy, phenomenology, and the ethics of technology:

1) *Phenomenology, Lifeworld, and Enframing:* Husserl [50] warned against substituting formal technical abstractions for the lived lifeworld (*Lebenswelt*); within the MWF, the gradient $\nabla_h W$ represents an intentional horizon of felt skill and agency rather than external mechanical motion. Heidegger [5] diagnosed modern technological "enframing" (*Gestell*) as reducing human beings to interchangeable "standing-reserve" (*Bestand*)—which corresponds analytically to setting $\alpha_i\to 0$ and deploying directive automation with $\Gamma_{ij}<0$, treating workers merely as residual buffers around machine bottlenecks.
2) *Aristotelian Eudaimonia, Vita Activa, and Unpaid Purposeful Work:* Aristotle [51] distinguished bare subsistence (*zên*) from human flourishing (*eu zên* or *eudaimonia*) realized through the active exercise of practical reason and craft excellence, while Arendt [4] separated repetitive biological *labor* from durable *work* and plural civic *action*. Crucially, the human state space X_H is not restricted to waged market employment: a complete accounting of societal Master Work encompasses unpaid purposeful activities—child-rearing, family and elder caregiving, community volunteering, artistic creation, open-source software and scientific stewardship, and democratic jury or council deliberation—whose intrinsic and relational value $W_H(h)$ is invisible to gross domestic product.
3) *Existential Meaning Safeguard:* Synthesizing Marx's critique of estranged labor [3] with Frankl's logotherapy [52], meaningful agency $W_H(h)+W_{\mathrm{int}}(h,m)$ functions as a psychological buffer against existential vacuum and depressive anhedonia—a hypothesis we test directly using clinical PHQ-2 mental-health microdata in Section VII.
4) *Kantian Dignity and the Four Pillars of Beneficial AI:* Kant's Formula of Humanity [53] enjoins treating rational agents always as ends in themselves and never merely as instruments of automated efficiency. Combined with capability justice [44], [54] and Russell's framework for provably beneficial AI [55], the MWF operationalizes four core ethical pillars: *beneficence* ($\Delta W>0$), *non-maleficence* (bounding overload strain $\kappa_i h_i^2$ and deskilling), *distributive justice* (inequality-averse social aggregation across strata), and *human dignity* (preserving substantive autonomy $h_{\mathrm{aut}}>0$ in moral and civic decisions).

B. Governmentality, Algorithmic Bossware, and Industrial Democracy

Foucault's analyses of panoptic discipline [56] and biopolitical governmentality [57] illuminate how workplace instrumentation shapes power asymmetries. Contemporary "bossware"—continuous keystroke and gaze tracking, algorithmic warehouse scan-rate timers, and automated disciplinary scoring—induces workers to internalize relentless self-surveillance, simultaneously inflating cognitive burnout curvature $\kappa_i h_i^2$ and imposing negative cross-partials $\Gamma_{\mathrm{aut,pace}}<0$. By contrast, epistemic transparency (m_{transp}) and participatory industrial democracy (co-determination councils,

worker-in-the-loop tool co-design, and transparent algorithmic auditing) invert this power geometry: when workers help specify what machines automate and retain override authority, surveillance disutility recedes and collaborative synergy $\Gamma_{\text{skill,transp}} > 0$ emerges.

C. Atkinson–Rawls Social Master Work Aggregation

Consider an economy or enterprise comprising G occupational groups or workers with population weights $\omega_k > 0$ ($\sum_{k=1}^{G} \omega_k = 1$) and positive normalized Master Work realizations $W_k > 0$. Rawls [6], [7] identified self-respect—anchored in meaningful work and purposeful civic association—as "perhaps the most important primary good." Pure utilitarian averaging ($\sum_k \omega_k W_k$) is blind to a polarized "caste system of meaning" in which an elite cognitive stratum enjoys creative AI augmentation while routine clerical and manual workers suffer severe autonomy collapse.

Definition 2 (Atkinson–Rawls Social Master Work Functional)**.** Parameterized by inequality aversion $\varepsilon \geq 0$, the equally distributed equivalent *Social Master Work Functional* [43] is

$$W_{\text{soc}}(\varepsilon) = \begin{cases} \left(\sum_{k=1}^{G} \omega_k W_k^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}, & \varepsilon \geq 0,\ \varepsilon \neq 1, \\ \exp\left(\sum_{k=1}^{G} \omega_k \ln W_k \right), & \varepsilon = 1, \\ \min_{k:\omega_k > 0} W_k, & \varepsilon \to \infty. \end{cases} \tag{31}$$

Equation (31) spans a continuous normative spectrum: $\varepsilon = 0$ recovers utilitarian mean welfare; $\varepsilon = 1$ yields scale-invariant Nash bargaining welfare; $\varepsilon = 2$ places strong curvature weight on low-meaning occupations; and $\varepsilon \to \infty$ converges to the Rawlsian maximin floor (operationalized empirically via the bottom employment decile $P_{10}(W)$ against outlier noise).

D. Multimodal Work Metrics, Goodhart Proxy Distortion, and VCG Elicitation

Operationalizing $W(h, m)$ in practice requires triangulating three complementary measurement streams: (i) objective workflow and interface telemetry (collaborative iteration cadences, override rates, and error recovery); (ii) validated psychometric instruments capturing autonomy, engagement, and eudaimonic meaning (JCM [26], SDT [27], UWES [28], and WAMI [29]); and (iii) in high-hazard domains such as surgical suites or flight decks, strictly opt-in, worker-owned physiological cognitive-load indicators (e.g., heart-rate variability or pupillometry) used exclusively to trigger local assistive load-relief rather than managerial evaluation.

Crucially, organizations cannot observe latent eudaimonic fulfillment continuously and may be tempted to tie high-powered bonuses or KPIs to easily logged machine telemetry (e.g., lines of AI-generated code, tickets closed per hour, or prompt volume). By Goodhart's Law [46] and Holmstrom and Milgrom's multitask principal–agent theory [45], rewarding noisy proxies distorts effort away from unmeasured craft and relational quality.

Proposition 2 (Multitask Goodhart Distortion Bound)**.** *Suppose a worker allocates effort $(e_1, e_2) \in \mathbb{R}_+^2$ between observable quantitative throughput e_1 and unobservable intrinsic craft/relational care e_2 under convex coupled effort cost $C(e_1, e_2) = \frac{1}{2}(c_1 e_1^2 + 2c_{12} e_1 e_2 + c_2 e_2^2)$ with $c_1 c_2 > c_{12}^2$ and attention substitutability $c_{12} > 0$. True socio-technical value is $W(e_1, e_2) = v_1 e_1 + v_2 e_2$, with intrinsic motivation baseline $\underline{e}_2 = \eta v_2 / c_2 > 0$, whereas the principal observes only $s_1 = e_1 + \epsilon_1$ ($\text{Var}(\epsilon_1) = \sigma_1^2$) and offers linear incentive wage $w(s_1) = w_0 + b_1 s_1$ to a CARA-normal agent with risk aversion $r > 0$. Then:*

1) *Raising proxy incentive intensity b_1 strictly crowds out unmeasured craft effort at rate $\partial e_2^*/\partial b_1 = -c_{12}/(c_1 c_2 - c_{12}^2) < 0$, driving $e_2^*(b_1) = 0$ whenever $b_1 \geq \bar{b}_1 := \eta v_2 (c_1 c_2 - c_{12}^2)/(c_2 c_{12})$.*
2) *Whenever unmeasured relational/craft value v_2 satisfies $v_2/v_1 > c_2/c_{12}$, any positive proxy bonus $b_1 > 0$ strictly* reduces *true Master Work value $W\big(e_1^*(b_1), e_2^*(b_1)\big) < W\big(e_1^*(0), e_2^*(0)\big)$.*

Proof. Including intrinsic valuation $\eta v_2 e_2$ ($\eta \in (0, 1]$), the agent maximizes certainty-equivalent payoff $w_0 + b_1 e_1 + \eta v_2 e_2 - C(e_1, e_2) - \frac{1}{2} r b_1^2 \sigma_1^2$. First-order conditions $c_1 e_1 + c_{12} e_2 = b_1$ and $c_{12} e_1 + c_2 e_2 = \eta v_2$ invert on the interior region to give $e_1^*(b_1) = (c_2 b_1 - c_{12} \eta v_2)/\Delta$ and $e_2^*(b_1) = (c_1 \eta v_2 - c_{12} b_1)/\Delta$, where $\Delta := c_1 c_2 - c_{12}^2 > 0$. Differentiating with respect to b_1 yields $\partial e_1^*/\partial b_1 = c_2/\Delta > 0$ and $\partial e_2^*/\partial b_1 = -c_{12}/\Delta < 0$, hitting the corner $e_2^* = 0$ at $\bar{b}_1$ (adjusted by $\eta v_2 c_1 / c_{12}$). Along the interior response, $\frac{d}{db_1} W(e_1^*, e_2^*) = (v_1 c_2 - v_2 c_{12})/\Delta$, which is strictly negative whenever $v_2/v_1 > c_2/c_{12}$. □

To elicit truthful preferences over shared workplace automation configurations $a \in \mathcal{A}$ (such as selecting between assistive copilots vs. automated ticket routing) without Goodhart gaming, organizations can combine low-powered audit metrics with Vickrey–Clarke–Groves (VCG) internal mechanism design [58]–[60]. If team $k \in \{1, \ldots, G\}$ holds private socio-technical valuation $\theta_k(a)$ and reports $\hat{\theta}_k(a)$, selecting $a^*(\hat{\boldsymbol{\theta}}) \in \arg\max_{a \in \mathcal{A}} \sum_{k=1}^{G} \hat{\theta}_k(a)$ paired with Clarke pivot budget credits $\tau_k(\hat{\boldsymbol{\theta}}) = \sum_{\ell \neq k} \hat{\theta}_\ell(a^*) - \max_{a \in \mathcal{A}} \sum_{\ell \neq k} \hat{\theta}_\ell(a)$ makes truthful reporting $\hat{\theta}_k = \theta_k$ a dominant strategy while selecting the exact MWF-maximizing configuration.

VI. Domain Formulations: Robotics, Agentic AI, Pedagogy, and Cyber-Physical Infrastructure

A. Collaborative, Assistive, and Social Robotics

Across industrial manufacturing cells, surgical robotics, rehabilitation clinics, and eldercare facilities, let $h = (h_{\text{skill}}, h_{\text{vig}})$ denote human sensorimotor agency and supervisory vigilance, and let $m = (m_{\text{assist}}, m_{\text{pace}})$ denote shared-

control robotic assistance versus machine-dictated cycle pacing [9], [25]. Instantiating (11) gives

$$W_{\mathrm{CPS}}(h,m) = \alpha_1 \ln(1+h_{\mathrm{skill}}) + \alpha_2 \ln(1+h_{\mathrm{vig}}) - \frac{\kappa}{2}\|h\|_2^2 + \beta_1 \ln(1+m_{\mathrm{assist}}) - \frac{\nu}{2}\|m\|_2^2 + \Gamma_{11} h_{\mathrm{skill}} m_{\mathrm{assist}} - |\Gamma_{12}| h_{\mathrm{skill}} m_{\mathrm{pace}}, \tag{32}$$

subject to safety-critical vigilance bounds $h_{\mathrm{vig}} \in [h_{\min}, h_{\max}]$. Three robotics sub-domains illustrate how (32) governs controller synthesis:

1) *Industrial Cobots and Interaction Fluency:* In collaborative assembly and robotic surgery, Hoffman [61] demonstrated that human satisfaction and team performance hinge on *interaction fluency*—anticipatory timing, minimal concurrent idle delay, and seamless motion handovers—rather than raw actuator speed. Shifting control from rigid conveyor pacing m_{pace} ($-|\Gamma_{12}| < 0$) to impedance-controlled fluency m_{assist} ($\Gamma_{11} > 0$) moves the KKT equilibrium from deskilled machine tending to high-craft partnership.
2) *Rehabilitation and Assistive Robotics:* For patients using lower-limb exoskeletons, myoelectric prosthetics, or smart wheelchairs, excessive robotic takeover ($\partial_{m_{\mathrm{assist}}} W_{\mathrm{CPS}} < 0$) induces motor slacking, muscle atrophy, and learned helplessness. Assist-as-needed impedance controllers that track (16) supply just enough torque to prevent perilous falls while keeping the user actively exerting rehabilitative effort ($\partial_{h_{\mathrm{skill}}} W_{\mathrm{CPS}} > 0, W_{\mathrm{int}} > 0$).
3) *Social and Companion Robotics:* In pediatric tutoring and eldercare companionship, social robots maximize W_{CPS} when calibrated to scaffold human-to-human social engagement and autonomous daily routines rather than substituting synthetic interaction or intervening prematurely in tasks the user can accomplish independently.

B. Value-Aligned Agentic AI, Workforce Scheduling, and Regulatory Auditing

Modern enterprise AI increasingly deploys multi-step autonomous agents via programmatic APIs alongside interactive chat/artifact interfaces [21]–[23]. Standard reinforcement learning from human feedback (RLHF) optimizes immediate task-completion or helpfulness scores $R_{\mathrm{task}}(s,a)$, incentivizing the agent to take over cognitive decisions directively even when doing so deskills the human user and masks brittle out-of-distribution reasoning [62]. Modeling the joint human–agent interaction as a Markov Decision Process over state $s_t = (h_t, m_t, z_{\mathrm{env},t})$ with discount $\gamma_{\mathrm{disc}} \in (0,1)$, we define the *MWF-Shaped Alignment Reward*:

$$R_{\mathrm{MWF}}(s_t,a_t,s_{t+1}) = R_{\mathrm{task}}(s_t,a_t) + \lambda_{\mathrm{align}}\big[\gamma_{\mathrm{disc}} W(h_{t+1},m_{t+1}) - W(h_t,m_t)\big] + \mu_{\mathrm{int}} W_{\mathrm{int}}(h_{t+1},m_{t+1}). \tag{33}$$

When $\lambda_{\mathrm{align}}, \mu_{\mathrm{int}} > 0$, (33) penalizes opaque, directive actions that erode user comprehension h_{t+1} and rewards Socratic option-generation, uncertainty calibration, and verifiable intermediate artifacts that elevate $W_{\mathrm{int}}(h,m)$. Because individual users hold heterogeneous preferences over autonomy and cognitive load, (33) integrates naturally with Cooperative Inverse Reinforcement Learning (CIRL) [55], [63] and online actor–critic architectures: the AI agent treats the user's structural parameters (α, κ, Γ) as latent variables inferred through deferential clarification and shared decision checkpoints.

In enterprise operations, optimizing R_{MWF} transforms concrete management algorithms: (i) in *algorithmic workforce scheduling*, incorporating human schedule predictability and recovery bounds prevents fragmented "clopening" micro-shifts that inflate burnout $\kappa_i h_i^2$; (ii) in *customer-service triage*, automated agents absorb repetitive tier-1 status queries while routing complex, emotionally nuanced cases to empowered human specialists ($\Gamma_{ij} > 0$), rather than scripting human representatives like robotic voice menus; and (iii) in *AI governance and regulatory auditing*, evaluating candidate deployments against the criterion $\mathbb{E}[\Delta W] > 0$ and $\Delta P_{10}(W) \geq 0$ operationalizes a concrete human-impact audit: *does this system measurably improve the human condition inside the loop of the job?*

C. Triadic Pedagogy, Vocational Blended Learning, and Lifelong Retraining

Applying the multi-agent hypergraph expansion (10) to educational and workforce-training environments yields a three-agent state $(h_s, h_t, m_{\mathrm{AI}})$ spanning learner cognitive effort h_s, human educator/mentor guidance h_t, and AI instructional capability m_{AI}:

$$W_{\mathrm{edu}}(h_s,h_t,m_{\mathrm{AI}}) = W_S(h_s) + W_T(h_t) + W_M(m_{\mathrm{AI}}) + \Gamma_{st} h_s h_t + \Gamma_{sm} h_s m_{\mathrm{AI}} + \Gamma_{tm} h_t m_{\mathrm{AI}} + \Lambda_{stm} h_s h_t m_{\mathrm{AI}} - \frac{\kappa_s}{2}(h_s - h_{\mathrm{ZPD}})^2, \tag{34}$$

where h_{ZPD} denotes Vygotsky's Zone of Proximal Development [64]. Answer-generating chatbots exhibit $\Gamma_{sm} < 0$ (cognitive offloading), shrinking optimal learner effort h_s^*. Conversely, AI systems designed to relieve administrative grading drudgery ($\Gamma_{tm} > 0$) while prompting student self-explanation ($\Gamma_{sm} > 0, \Lambda_{stm} > 0$) simultaneously expand teacher mentoring capacity and learner mastery. In corporate and vocational retraining for mid-career workers, (34) explains why purely solitary e-learning modules suffer high attrition ($h_t = 0$, losing both $\Gamma_{st} h_s h_t$ and $\Lambda_{stm} h_s h_t m_{\mathrm{AI}}$): pairing adaptive simulation sandboxes with human-moderated peer cohorts and master-apprentice shop-floor coaching preserves social belonging while accelerating cross-occupational skill transitions.

D. Safety-Critical Engineering, Smart Grids, and Cyber-Physical Infrastructure

At the scale of critical public infrastructure, optimizing machine efficiency without state variables for human vigilance and civic agency creates systemic fragility:

1) *Safety-Critical Supervisory Systems:* In commercial aviation flight decks, air traffic control, nuclear power plants, and SAE Level 2–3 automated vehicles [8], [9], relegating operators to passive monitoring drives

$h_{\mathrm{vig}}(t) < h_{\min}$ via (17), inducing mode confusion and delayed takeover during rare tail-risk faults. Designing periodic interactive engagement loops, shared trajectory planning, and transparent diagnostic displays sustains $\partial_{h_{\mathrm{vig}}} W \geq 0$ within the safe operating envelope.

2) *Smart Electrical Grids and Prosumer Energy Management:* Fully opaque utility load-shedding algorithms provoke consumer overrides when households lose control over domestic comfort and backup storage ($\Gamma_{\mathrm{aut,grid}} < 0$). Conversely, interpretable demand-response interfaces that allow households and community microgrid operators to trade off tariff savings, resilience, and carbon intensity convert passive ratepayers into active grid participants ($\Gamma_{\mathrm{aut,grid}} > 0$).
3) *Participatory Smart Cities and Cyber-Physical Security:* Municipal analytics platforms and critical-infrastructure security operations centers (SOCs) achieve higher resilience when machine anomaly detectors filter high-volume telemetry noise while preserving human deliberation in urban planning trade-offs and strategic adversarial threat hunting.

VII. Multi-Scale Empirical Validation and U.S. Macroeconomic Counterfactuals

To test the structural, microeconometric, stochastic-dynamic, and macroeconomic predictions of the Master Work Function, we assembled a multi-scale empirical pipeline linking six major public datasets across U.S. occupations and individual workers.

A. Datasets, Crosswalks, and Construct Validity

Table I summarizes the empirical data architecture, sample sizes, and operationalized state coordinates. At the occupational scale ($N = 873$ complete O*NET-SOC occupations representing 124,009,411 U.S. workers), we extract task, skill, ability, work-context, and work-value descriptors from the U.S. Department of Labor's O*NET 30.1 database [47] and merge them via exact O*NET-SOC and 6-digit Standard Occupational Classification (SOC) crosswalks with: (i) observed generative AI job exposure, task-level augmentation versus directive telemetry, and April–May 2026 interactive (`claude_ai`) versus programmatic enterprise API (`1p_api`) transcripts from the Anthropic Economic Index (AEI) [20]–[22]; (ii) 923 occupation-level physical robotics exposure scores (`robot_exposure_index`) from Legate-Yang and Massenkoff [25]; (iii) LLM task exposure (β-ratings) from Eloundou *et al.* [17]; (iv) AI Occupational Exposure (AIOE) scores from Felten *et al.* [16]; and (v) national employment and wage distributions from the Bureau of Labor Statistics Occupational Employment and Wage Statistics (BLS OEWS) [49].

Following Self-Determination Theory [27] and psychological meaningful-work taxonomies [26], [29], we construct the occupational target $W_{\mathrm{occ}} \in [0, 1]$ from O*NET Work Values capturing intrinsic eudaimonic fulfillment (*Achievement*, *Independence*, and *Relationships*), which exhibits strong internal consistency (Cronbach $\alpha = 0.764$; six-item broad Work Values index $\alpha = 0.830$). Crucially, to verify that O*NET occupational W_{occ} reflects genuine lived worker experience rather than expert-rating artifacts, we cross-validate W_{occ} out-of-dataset against independent worker self-reports of *"I find my work meaningful"* (`wrkmeangfl`) in the 2022 NORC General Social Survey [48]. Across $N = 1{,}897$ linked workers, occupational W_{occ} significantly predicts individual worker meaningfulness ($r = 0.203, p = 4.76 \times 10^{-19}$); aggregating across $N = 115$ occupation cells with at least five survey respondents yields $r = 0.497$ ($p = 1.64 \times 10^{-8}$), establishing strong cross-dataset construct validity. Fig. 1(b) maps all 873 occupations across composite human agency H_o and augmentative AI exposure M_o.

B. Occupational Structural Estimation and the Interface Augmentation Gap

We estimate five nested econometric specifications of W_{occ} across the $N = 873$ occupations using exact ordinary least squares with MacKinnon–White HC3 heteroskedasticity-robust standard errors [65] and 10-fold out-of-sample cross-validation (CV R^2). All continuous human and machine coordinates are normalized to $[0, 1]$ and mean-centered in bilinear interaction products $h_{i,c} m_{j,c}$ so that main effects evaluate at the sample mean while interaction coefficients estimate exact cross-partials $\hat{\Gamma}_{ij} = \partial^2 W / \partial h_i \partial m_j$.

Table II reports model fit and structural parameters. Out-of-sample 10-fold CV R^2 rises monotonically from standalone human (0.735) and standalone machine (0.568) models to the separable additive model (0.769), unregulated log-bilinear form (0.776), and full burnout-regulated MWF specification ($R^2 = 0.798$, Adj. $R^2 = 0.792$, CV $R^2 = 0.782$). While Schwarz's BIC favors parsimony when all 16 interaction terms are entered simultaneously (BIC $= -1515.5$ additive vs. -1485.1 full 4×4 grid, since several off-diagonal cells are near zero), the joint nested F-test decisively rejects additive separability $\Gamma = 0$: $F(16, 844) = 4.56, p = 8.09 \times 10^{-9}$.

The estimated main effects and $\hat{\Gamma}_{ij}$ heatmap in Fig. 2(a) illuminate the structural mechanics of Theorem 1 and Proposition 1:

1) *Divergent Machine Main Effects:* Collaborative AI augmentation carries a positive and statistically significant main effect on occupational meaning ($\hat{\beta}_{\mathrm{aug}} = +0.092, p = 0.027$), whereas raw cognitive AI exposure without augmentation is insignificant ($\hat{\beta}_{\mathrm{ai_cap}} = -0.074, p = 0.179$) and physical robotics exposure is strongly associated with lower baseline occupational meaning ($\hat{\beta}_{\mathrm{robotics}} = -0.311, p = 1.26 \times 10^{-13}$).
2) *Selective Cross-Partial Synergies and Interference:* Technical skill pairs positively with structured equipment control ($\hat{\Gamma}_{\mathrm{skill,rigid}} = +0.628, p = 0.0032$, reflecting skilled precision operators in surgical, aviation, and laboratory instrumentation). Conversely, rigid machine-dictated pacing sharply penalizes human creative originality ($\hat{\Gamma}_{\mathrm{creat,rigid}} = -0.576, p = 0.0149$), and physical robotics exposure exhibits strong negative complementarity with relational care ($\hat{\Gamma}_{\mathrm{rel,robotics}} = -0.628, p = 5.1 \times 10^{-4}$).

TABLE I
SUMMARY OF LINKED PUBLIC DATASETS, SAMPLE COVERAGE, AND OPERATIONALIZED MASTER WORK FUNCTION VARIABLES

Dataset & Release	**Granularity**	**Sample Size**	**Operationalized MWF Coordinates & Role**
U.S. Dept. of Labor O*NET 30.1 [47]	8-digit O*NET-SOC	873 occs	Eudaimonic meaning target W_{occ} (*Achievement*, *Independence*, *Relationships*, $\alpha = 0.764$; 6-item $\alpha = 0.830$); human agency $h_{\text{skill}}, h_{\text{aut}}, h_{\text{creat}}, h_{\text{rel}}$; overload strain h_{strain}; equipment pacing m_{rigid}.
Anthropic Economic Index (2025–2026) [20], [22]	6-digit SOC & Task	696 SOCs / 1.06M rows	Observed AI exposure; interactive (`claude_ai`) vs. headless (`1p_api`) autonomy & collaborative augmentation share (m_{aug}, $N = 647$ matched SOCs); multi-wave task shifts.
Physical Robot Exposure (Sept. 2026) [25]	8-digit O*NET-SOC	923 occs	Time-weighted physical robotics capability m_{robotics} across 18,796 rated O*NET tasks (tiers E_0–E_3).
GPTs-are-GPTs [17] & AIOE [16]	O*NET-SOC / SOC	873 occs	Composite cognitive AI capability index $m_{\text{ai_cap}}$ combining AIOE, LLM β-exposure, and observed AEI penetration.
BLS National OEWS [49]	6-digit SOC	124.01M workers	Employment weights ω_k, median and mean annual wages (828 wage-complete occupations, 123.41M workers across 21 major SOC sectors).
NORC General Social Survey (2018, 2022, 2024) [48]	Individual Worker	10,483 respondents	Worker meaningful work (`wrkmeangfl`, $N = 1{,}880$), NIOSH Quality of Worklife fulfillment ($N = 3{,}301$), PHQ-2 depressive anhedonia ($N = 1{,}598$), 3-wave job satisfaction ($N = 7{,}474$), linked via Census `occ10`–SOC crosswalk.

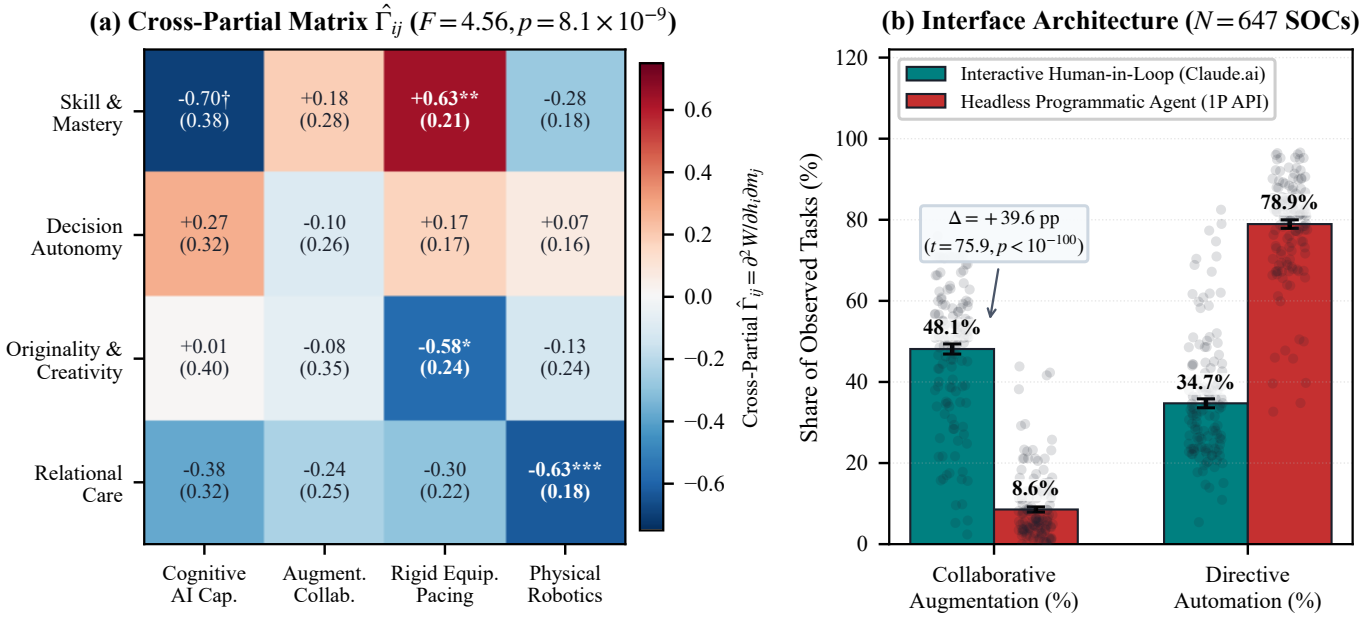

Fig. 2. Empirical socio-technical complementarity and deployment architecture contrast. (a) Estimated 4×4 cross-partial interaction matrix $\hat{\Gamma}_{ij} = \partial^2 W / \partial h_i \partial m_j$ across $N = 873$ U.S. occupations (HC3 robust standard errors in parentheses; $^{\dagger}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001$). (b) Paired within-occupation comparison across $N = 647$ 6-digit SOCs in May 2026 Anthropic Economic Index telemetry between interactive human-in-the-loop sessions (`claude_ai`) and headless enterprise API deployments (`1p_api`).

Why does raw cognitive AI capability exhibit mixed complementarity unless paired with collaborative interfaces? Fig. 2(b) provides direct behavioral telemetry evidence from the June 2026 Anthropic Economic Index release [22]. Comparing interactive human-in-the-loop usage (`claude_ai`: chat and collaborative artifacts) against programmatic first-party enterprise API traffic (`1p_api`) across the *exact same* $N = 647$ matched 6-digit SOC occupations, interactive workflows achieve a mean collaborative augmentation share of 48.1% (34.7% directive execution), whereas headless API deployments collapse to 8.6% augmentation and 78.9% directive automation, a paired within-occupation gap of +39.6 percent-

TABLE II
NESTED OCCUPATIONAL MWF MODEL COMPARISON AND SELECTED STRUCTURAL ESTIMATES ($N = 873$ OCCUPATIONS, HC3 ROBUST SEs)

Panel A: Nested Model	k	R^2	**Adj.** R^2	**CV** R^2
1. Pure Human $W_H(h)$	5	0.740	0.739	0.735
2. Pure Machine $W_M(m)$	5	0.575	0.573	0.568
3. Separable ($\Gamma = 0$)	10	0.774	0.772	0.769
4. Unregulated Log-Bilinear	25	0.792	0.786	0.776
5. Burnout-Regulated MWF	29	**0.798**	**0.792**	**0.782**
Joint test of $\Gamma = 0$: $F(16, 844) = 4.56$, $p = 8.09 \times 10^{-9}$				
Panel B: Structural Terms	**Coef.**	**SE**	t**-stat**	p**-val.**
$\ln(1 + h_{\text{skill}})$ (Mastery)	$+0.605$	0.053	$+11.32$	9.6×10^{-28}
$\ln(1 + h_{\text{creat}})$ (Originality)	$+0.184$	0.054	$+3.40$	7.0×10^{-4}
$\ln(1 + h_{\text{rel}})$ (Relational)	$+0.431$	0.057	$+7.59$	8.4×10^{-14}
$\ln(1 + m_{\text{aug}})$ (Augment. AI)	$+0.092$	0.042	$+2.21$	0.0273
$\ln(1 + m_{\text{robotics}})$ (Robotics)	-0.311	0.041	-7.54	1.3×10^{-13}
$\hat{\Gamma}_{\text{skill,rigid}}$ ($h_{\text{skill}} \times m_{\text{rigid}}$)	$+0.628$	0.213	$+2.95$	0.0032
$\hat{\Gamma}_{\text{creat,rigid}}$ ($h_{\text{creat}} \times m_{\text{rigid}}$)	-0.576	0.236	-2.44	0.0149
$\hat{\Gamma}_{\text{rel,robotics}}$ ($h_{\text{rel}} \times m_{\text{robotics}}$)	-0.628	0.180	-3.49	5.1×10^{-4}
$\hat{\Gamma}_{\text{skill,ai_cap}}$ ($h_{\text{skill}} \times m_{\text{ai_cap}}$)	-0.698	0.376	-1.86	0.0639

age points ($t = 75.88, p < 10^{-100}$; AI autonomy score 2.74 vs. 2.24 scale inversion across modes, $|t| = 55.11$). Thus, Γ_{ij} is not technologically predetermined by the foundation model; it is an architectural design choice governed by how human direction is interleaved into execution [23].

C. Individual Worker Microdata, Burnout Concavity, and Depressive Anhedonia

To test whether structural MWF relationships hold at the level of individual workers controlling for occupational clustering and income, we analyze 10,483 respondents across the 2018, 2022, and 2024 waves of the NORC General Social Survey (GSS) [48], computing Liang–Zeger standard errors clustered by Census occupation (`occ10`) [66]. Worker-level coordinates are constructed on $[0, 1]$ from validated NIOSH Quality of Worklife (QWL) items: skill mastery h_{skill} (`myskills`, `learnnew`, `opdevel`), task autonomy h_{aut} (`wkfreedm`, `wkdecide`), relational respect h_{rel} (`partteam`, `respect`, `trustman`), overload strain h_{overload} (`overwork`, `usedup`, `stress`), and tool/information transparency m_{transp} (`hlpequip`, `haveinfo`), merged with occupational m_{aug} and m_{rigid}.

Table III and Fig. 3 establish four microeconometric results:

1) *Individual Agency and Tool Transparency:* In Model (1) ($N = 1{,}880, R^2 = 0.381$), worker skill mastery ($\hat{\beta} = +0.504, t = 15.82, p = 4.36 \times 10^{-42}$), relational respect ($\hat{\beta} = +0.158, p = 3.46 \times 10^{-10}$), and transparent equipment/information support ($\hat{\beta} = +0.104, p = 5.36 \times 10^{-4}$) strongly elevate perceived meaningfulness, whereas working in an occupation characterized by rigid equipment-dictated automation significantly depresses worker meaning ($\hat{\beta} = -0.231, t = -4.99, p = 1.01 \times 10^{-6}$).
2) *Empirical Verification of Quadratic Burnout Concavity:* In the larger pooled 2018–2022 QWL panel (Model 2, $N = 3{,}301, G = 382, R^2 = 0.456$), both linear strain ($\hat{\beta}_1 = -0.105, p = 0.074$) and quadratic overload curvature ($\hat{\beta}_2 = -0.117, t = -2.08, p = 0.038$) are negative, confirming the accelerating marginal welfare loss postulated in (11) [Fig. 3(b)].
3) *Existential Vacuum and Mental-Health Safeguard:* Model (3) ($N = 1{,}598, G = 308$) regresses the validated two-item Patient Health Questionnaire (PHQ-2) depressive anhedonia index (`feeldown` and `nointerest`) [67] on worker meaningful work, task autonomy, overload strain, and log real personal income. Holding real income ($\hat{\beta} = -0.039, p = 8.46 \times 10^{-14}$) and workload exhaustion ($\hat{\beta} = +0.346, p = 2.49 \times 10^{-26}$) constant, meaningful work exerts a large, highly significant protective effect against depressive anhedonia ($\hat{\beta} = -0.135, t = -5.05, p = 7.58 \times 10^{-7}$). As shown in Fig. 3(c), high-meaning workers below median income exhibit lower depressive anhedonia than low-meaning workers above median income, empirically substantiating Frankl's existential thesis [52].
4) *Post-2018 Aggregate Fulfillment Drift:* Across all three GSS waves ($N = 7{,}474$ workers reporting `satjob`), normalized U.S. job satisfaction declined monotonically from 0.770 in 2018 ($N = 1{,}739$) to 0.744 in 2022 ($N = 2{,}953$) and 0.738 in 2024 ($N = 2{,}782$). Controlling for real log income and weekly hours ($N = 4{,}573, G = 412$), the 2024 wave fixed effect remains significantly negative ($\hat{\beta}_{2024} = -0.029, t = -3.13, p = 0.0019$), underscoring that rising aggregate technology adoption has not automatically translated into higher lived work

TABLE III
NORC GSS INDIVIDUAL WORKER MICRODATA REGRESSIONS WITH OCCUPATION-CLUSTERED STANDARD ERRORS

Explanatory Variable	**(1) Meaning (2022)**	**(2) QWL Fulfillment**	**(3) PHQ-2 Anhedonia**
Skill Mastery h_{skill}	$+0.504^{***}$	$+0.362^{***}$	—
	(0.032)	(0.017)	
Task Autonomy h_{aut}	$+0.049^{\dagger}$	$+0.078^{***}$	-0.029
	(0.027)	(0.014)	(0.027)
Relational Respect h_{rel}	$+0.158^{***}$	$+0.140^{***}$	—
	(0.024)	(0.014)	
Tool Transparency m_{transp}	$+0.104^{***}$	$+0.177^{***}$	—
	(0.030)	(0.020)	
Overload Strain h_{overload}	-0.234^{**}	$-0.105^{\dagger}$	$+0.346^{***}$
	(0.083)	(0.059)	(0.030)
Quadratic Burnout h^2_{overload}	$+0.111$	$\mathbf{-0.117}^{*}$	—
	(0.088)	(0.056)	
Rigid Automation m_{rigid}	-0.231^{***}	-0.055^{*}	—
	(0.046)	(0.027)	
Meaning (`wrkmeangfl`)	—	—	$\mathbf{-0.135}^{***}$
			(0.027)
Log Real Income ln(inc)	—	—	-0.039^{***}
			(0.005)
Observations N	1,880	3,301	1,598
Occupation Clusters G	322	382	308
R^2 (Adj. R^2)	0.381 (0.378)	0.456 (0.454)	0.164 (0.162)

Clustered SEs in parentheses ($^{\dagger}p < 0.10, ^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001$). Models (1)–(2) include centered interactions and wave fixed effects.

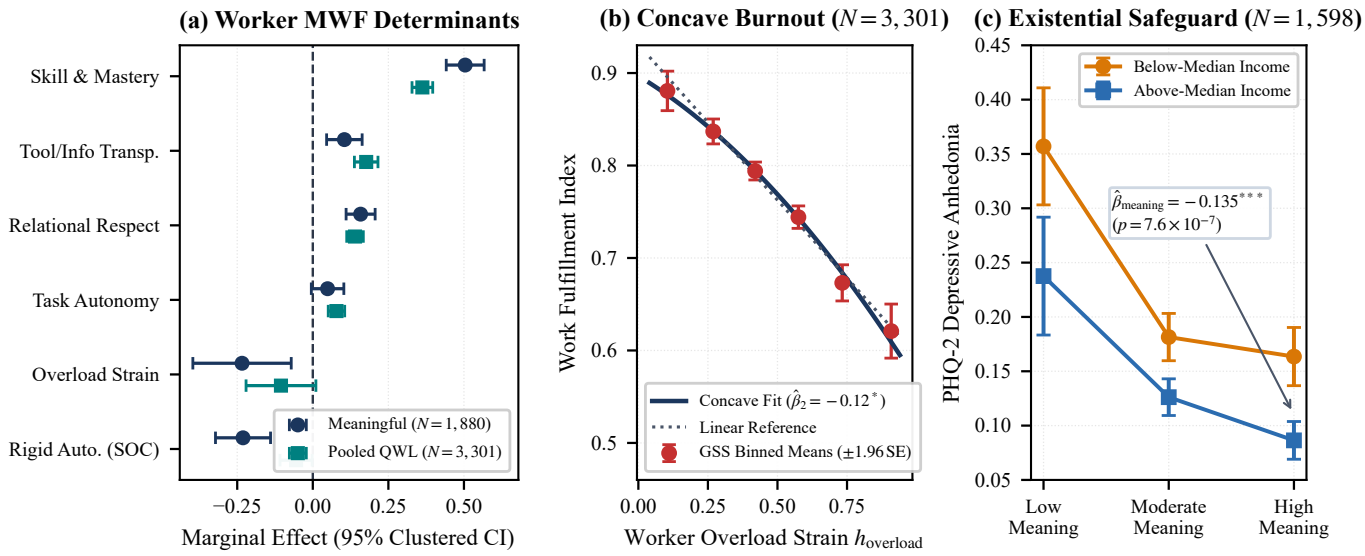

Fig. 3. Individual worker microdata evidence from the NORC General Social Survey (2018–2024). (a) Marginal effects and 95% occupation-clustered confidence intervals for direct meaningful work ($N = 1{,}880, G = 322$) and pooled QWL work fulfillment ($N = 3{,}301, G = 382$). (b) Binned scatterplot confirming accelerating concave burnout regulation ($\hat{\beta}_{\text{ov}^2} = -0.117^*, p = 0.038$). (c) Mean PHQ-2 depressive anhedonia [67] across meaningful-work tiers stratified by below- versus above-median real income ($N = 1{,}598, G = 308$).

TABLE IV
STOCHASTIC LANGEVIN VALIDATION OF JARZYNSKI EQUALITY AND SIVAK–CROOKS ROLLOUT PROTOCOLS ($M = 15{,}000$ TRAJECTORIES, $\hat{K} = 0.085$)

Protocol $\lambda(t)$	$\mathbb{E}[C_{\text{reorg}}]$	**Est.** $\hat{\Delta\Phi}$	$\lvert\text{Err}\rvert$	W_{diss}
Abrupt Shock ($\tau = 0.8$)	$+0.0202$	0.2963	0.0107	0.3272
Uniform Linear ($\tau = 0.8$)	-0.0761	0.3150	0.0080	0.2309
Sivak–Crooks Geodesic	-0.0854	0.3025	**0.0045**	**0.2216**
Quasi-Reversible ($\tau = 4.0$)	-0.2038	0.3003	0.0067	0.1032

Exact quadrature $\Delta\Phi_{\text{exact}} = 0.3070$ ($\tau = 0.8$ for shock, linear, and geodesic). Sivak–Crooks geodesic saves 32.3% dissipation vs. shock.

quality.

D. Stochastic Rollout Validation and Optimal Control Trajectories

To validate Theorem 3 and Corollary 1, we calibrate the socio-technical fluctuation scale $\hat{K} = 0.085$ from observed within-occupation residual variance in GSS meaningfulness and simulate ensembles of $M = 15{,}000$ overdamped Langevin trajectories (22) transitioning target machine allocation from $m_0 = 0.15$ ($\lambda = 0$) to $m_1 = 0.80$ ($\lambda = 1$) across a retraining bottleneck where curvature stiffness dips by 55% at $\lambda = 0.5$ (elevating relaxation lag $\tau_{\text{rel}}(\lambda)$ and friction $\zeta(\lambda)$ mid-transition). Exact 2D numerical quadrature of (24) yields equilibrium potential gain $\Delta\Phi_{\text{exact}} = 0.3070$.

Table IV and Fig. 4 confirm all three claims of Theorem 3:

1) *Exact Crooks Intersection and Jarzynski Recovery:* Forward linear $P_{\text{fwd}}(c)$ and reversed $P_{\text{rev}}(-c)$ cross precisely at $c = -\Delta\Phi_{\text{exact}} = -0.3070$ [Fig. 4(a)], and the exponential Jarzynski estimator recovers $\Delta\Phi_{\text{exact}}$ within 0.0045–0.0107 across all protocols despite up to $3\times$ variation in mean reorganization dissipation.
2) *Geodesic Dissipation Reduction:* Within the identical finite budget horizon $\tau = 0.8$, replacing an abrupt enterprise cutover ($W_{\text{diss}} = 0.3272$) with the Sivak–Crooks geodesic schedule $\lambda^*(t/\tau)$ ($W_{\text{diss}} = 0.2216$) reduces wasted organizational reorganization friction by 32.3% (and outperforms uniform linear rollout by 4.0%) by automatically slowing deployment across the high-variance retraining zone $\lambda \in [0.35, 0.65]$ [Fig. 4(b)].

Complementarily, Fig. 5 displays 20-quarter trajectories of the deterministic optimal control system (17)–(21). Naive throughput maximization ($\theta = 0, \sigma_{\text{aug}} = 0.14$) enjoys a brief initial output surge over quarters $t \in [0, 3.5]$, but directive skill atrophy drives $h(t)$ below the supervisory vigilance threshold by $t = 5$, causing quality-adjusted output $Y(t)$ to stagnate and decline while MWF-aligned Pontryagin co-investment ($\theta > 0, \sigma_{\text{aug}} = 0.82$) achieves strictly higher sustained output and nearly $3\times$ higher eudaimonic flow $W(t)$.

E. U.S. Macroeconomic Wage–Meaning Landscape and Policy Counterfactuals

Aggregating across $N = 828$ wage-complete occupations weighted by BLS national employment (123,410,998 U.S. workers across 21 major SOC groups), we examine how market wages align with—or diverge from—socio-technical meaning, and evaluate four economy-wide automation counterfactuals.

At the macro level, log median annual wage correlates positively with occupational W_{occ} ($r = 0.617, p = 4.90 \times 10^{-88}$), with between-occupation wage Gini $= 0.269$ and baseline meaning Gini $= 0.237$. Yet Fig. 6(a) reveals profound sectoral divergences:

1) *The Care and Education Wage–Meaning Gap: Education & Library* (5.89M workers) ranks **#1** in the U.S. economy in meaningful work ($W_{\text{occ}} = 0.784$) despite a

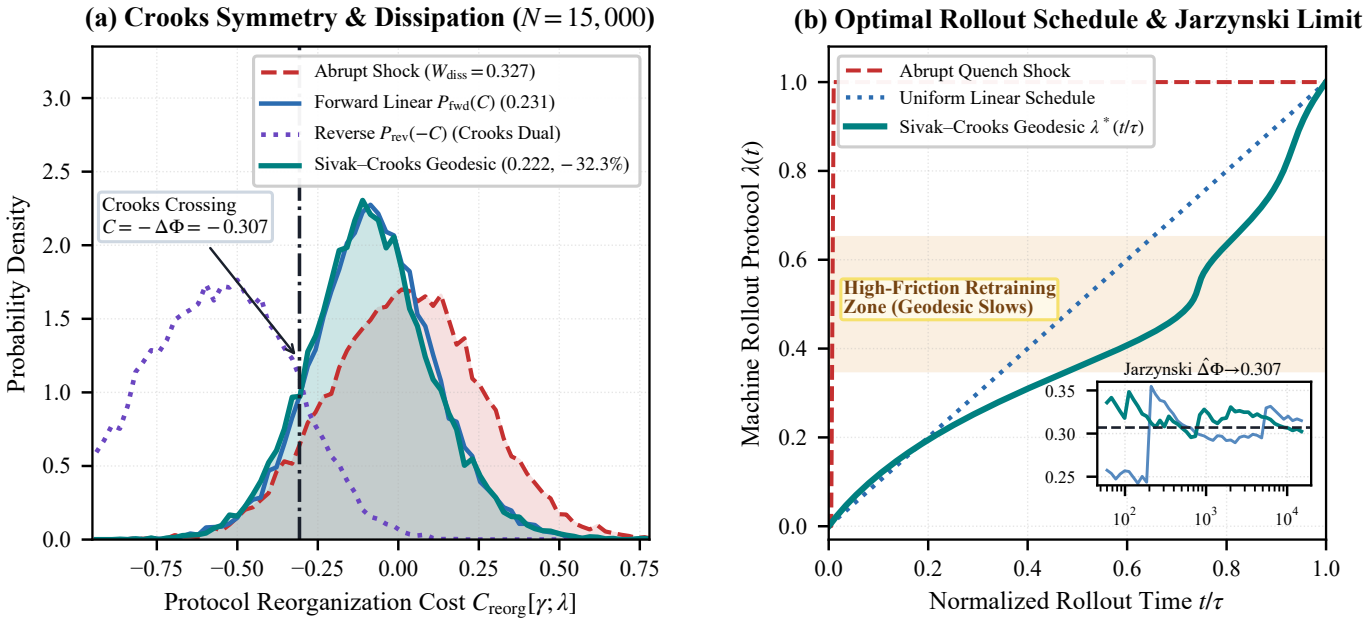

Fig. 4. Stochastic socio-technical thermodynamics across 15,000 Langevin rollout trajectories ($\hat{K} = 0.085, \Delta\Phi_{exact} = 0.3070$). (a) Forward and time-reversed reorganization cost densities $P_{fwd}(C)$ and $P_{rev}(-C)$ intersecting sharply at $-\Delta\Phi = -0.307$, alongside abrupt quench shock and Sivak–Crooks geodesic distributions. (b) Optimal minimum-dissipation rollout schedule $\lambda^*(t/\tau)$ decelerating through the mid-rollout retraining barrier, with inset showing Monte Carlo convergence of the Jarzynski estimator $\hat{K}\ln\langle e^{-C/\hat{K}}\rangle \to 0.307$.

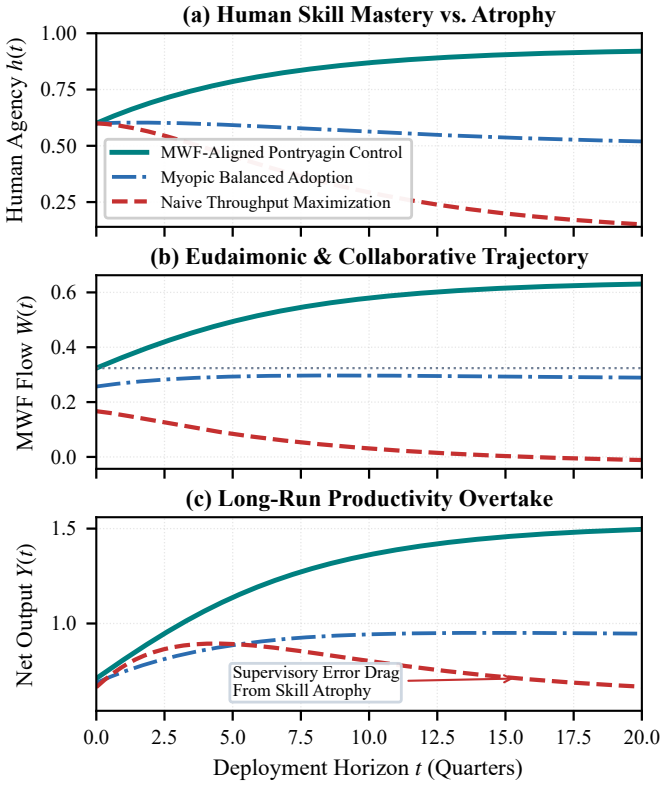

Fig. 5. Pontryagin dynamic control trajectories over a 20-quarter enterprise deployment horizon comparing MWF-aligned co-investment, myopic balanced adoption, and naive throughput maximization: (a) human skill/agency state $h(t)$, (b) instantaneous MWF flow $W(t)$, and (c) net quality-adjusted economic output $Y(t)$ exhibiting long-run productivity overtake as skill atrophy triggers supervisory error drag under naive directive automation.

modest median annual wage of \$61,465, and *Community & Social Service* (1.75M workers) ranks **#3** ($W_{occ} = 0.766$) at \$49,935—surpassing much higher-paid sectors including *Legal* ($W_{occ} = 0.644$, \$101,701), *Architecture & Engineering* ($W_{occ} = 0.598$, \$87,440), and *Computer & Mathematical* ($W_{occ} = 0.599$, \$84,474). Evaluating labor-market transitions solely through pecuniary wages systematically undervalues care and educational institutions.

2) *Dual Bifurcated Vulnerability Zones:* Cognitive AI exposure concentrates heavily in *Office & Administrative Support* (17.97M workers; highest sectoral AI exposure 0.634, moderate $W_{occ} = 0.447$, median wage \$41,360) and *Sales & Related* (11.99M workers; AI exposure 0.480, $W_{occ} = 0.452$). Conversely, physical robotics exposure concentrates in *Transportation & Material Moving* (11.44M workers; robot exposure 0.676, $W_{occ} = 0.272$), *Food Preparation & Serving* (11.10M workers; robot exposure 0.410, $W_{occ} = 0.335$), and *Production* (6.51M workers; robot exposure 0.403, $W_{occ} = 0.294$). Together, these routine clerical and physical-manual clusters encompass nearly 59 million U.S. workers.

Using our structurally calibrated parameters, Table V and Fig. 6(b) evaluate four national automation regimes:

1) *Baseline U.S. Economy:* Normalized productivity $Y = 1.000$, utilitarian $W_{soc}(0) = 0.496$, strong inequality-averse $W_{soc}(2) = 0.400$, Rawlsian bottom decile $P_{10}(W) = 0.222$, and Meaning Gini $= 0.237$.
2) *Directive Substitution ("So-So" Headless Automation):* Deploying AI and robotics primarily via directive substitution raises output modestly ($Y = 1.060$, $+6.0\%$) while

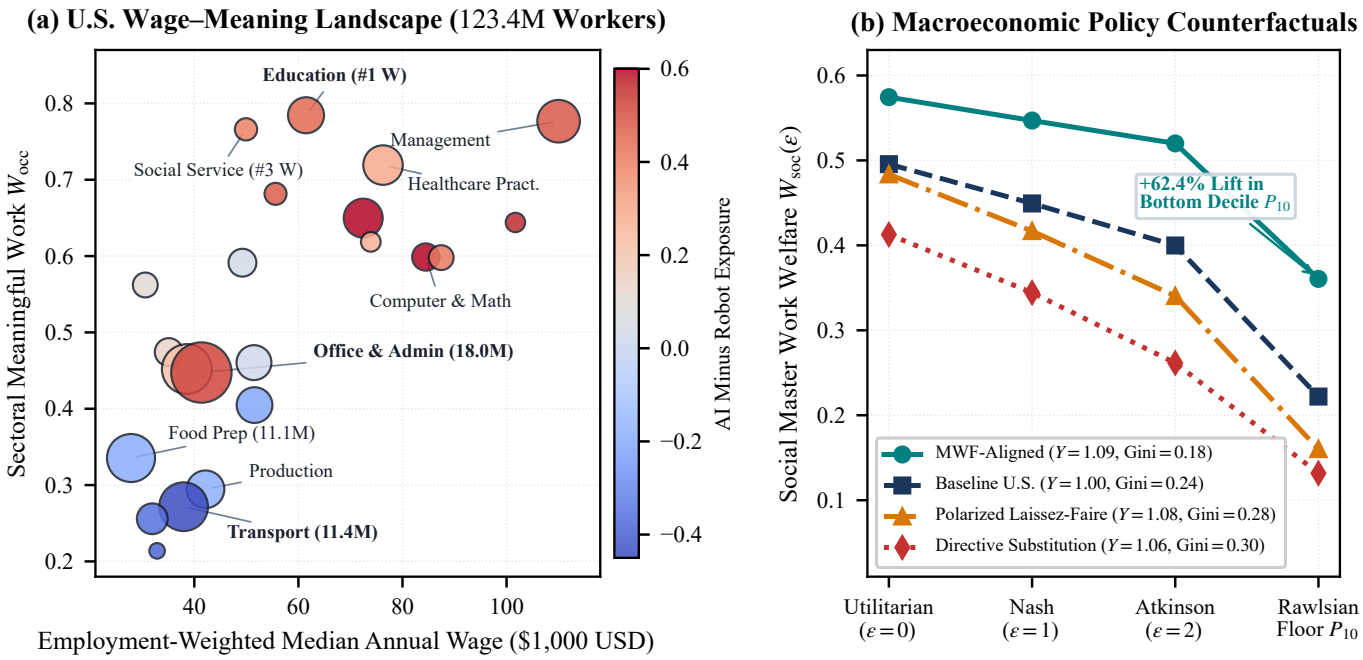

Fig. 6. U.S. macroeconomic socio-technical welfare across 123.41 million workers (21 major SOC sectors). (a) Sectoral median annual wage versus employment-weighted meaningful work W_{occ}, with bubble area proportional to BLS employment and color indicating cognitive AI minus physical robot exposure. (b) Inequality-averse Atkinson–Rawls welfare $W_{\text{soc}}(\varepsilon)$ and bottom-decile floor $P_{10}(W)$ across four national automation policy counterfactuals.

TABLE V
U.S. MACROECONOMIC POLICY COUNTERFACTUALS ACROSS 123.41 MILLION BLS WORKERS (828 OCCUPATIONS)

Welfare Metric	Base.	Direct.	Polar.	Aligned
Output Index Y	1.000	1.060	1.083	**1.093**
Utilitarian $W_{\text{soc}}(0)$	0.496	0.413	0.484	**0.574**
Atkinson $W_{\text{soc}}(0.5)$	0.473	0.381	0.452	**0.561**
Nash Welfare $W_{\text{soc}}(1)$	0.449	0.344	0.417	**0.547**
Inequality-Averse $W_{\text{soc}}(2)$	0.400	0.262	0.340	**0.520**
Rawlsian Floor $P_{10}(W)$	0.222	0.132	0.160	**0.361**
Meaning Gini Coefficient	0.237	0.302	0.279	**0.177**

eroding utilitarian meaning by -16.7% (0.413), collapsing the Rawlsian bottom decile by -40.6% ($P_{10} = 0.132$), and inflating Meaning Gini to 0.302.

3) *Polarized Laissez-Faire Adoption:* Allowing high-skill professionals to capture collaborative AI synergies while subjecting routine clerical and manual workers to rigid algorithmic pacing raises output to $Y = 1.083$ ($+8.3\%$) and leaves mean utilitarian welfare nearly flat (0.484, -2.4%), masking a -14.9% drop in $W_{\text{soc}}(2)$ (0.340) and a -27.9% drop in bottom-decile meaning ($P_{10} = 0.160$). This demonstrates why aggregate productivity and even unweighted mean satisfaction fail to detect distributional agency erosion.
4) *MWF-Aligned Complementarity and Targeted Upskilling:* Steering enterprise deployments toward interactive human-in-the-loop architectures ($\Gamma_{ij} > 0$), relieving hazardous physical drudgery via shared-control cobots, and co-investing in bottom-tercile task autonomy yields the highest productivity ($Y = 1.093$, $+9.3\%$), lifts utilitarian meaning to 0.574 ($+15.9\%$), raises $W_{\text{soc}}(2)$ to 0.520 ($+30.1\%$), boosts the Rawlsian floor by $+62.4\%$ ($P_{10} = 0.361$), and compresses Meaning Gini to 0.177 (-25.5%).

F. Macroeconomic Governance and Post-Scarcity Policy Instruments

Realizing the MWF-aligned frontier in Table V requires institutional instruments that correct market externalities between pecuniary output and socio-technical meaning:

1) *Care and Education Compensation Rebalancing:* Because *Education & Library* ($W_{\text{occ}} = 0.784$) and *Community & Social Service* ($W_{\text{occ}} = 0.766$) deliver the highest intrinsic and relational welfare in the economy despite modest market wages [Fig. 6(a)], public wage supplements, care-sector tax credits, and AI administrative-relief grants expand high-meaning employment insulated from substitutive crowd-out.
2) *Work-Time Reduction and Shared Productivity Dividends:* Revisiting Keynes's projection of progressively shorter working hours under technological abundance [68], channeling AI productivity gains into structured work-time reduction (such as four-day workweeks at constant real pay) distributes meaningful engagement across the workforce—preventing a polarized equilibrium in which retained specialists suffer quadratic burnout ($\kappa_i h_i^2$) while displaced workers confront forced idleness ($h = 0$).
3) *Public-Interest Job Guarantees:* Operationalizing Rawls's conception of constitutional democracy acting as an employer of last resort [7], voluntary

public-service employment in ecological restoration, community health, eldercare, and local cultural infrastructure establishes a direct institutional floor beneath $P_{10}(W)$ across regions undergoing rapid industrial automation.

4) *Universal Basic Income Paired with Civic Meaning Infrastructure:* Unconditional income floors protect bargaining power against demeaning labor, yet pecuniary transfers alone cannot prevent existential vacuum and depressive anhedonia [Table III, Model (3)] [52] if recipients lack recognized avenues for mastery and belonging. Pairing basic income security with community maker spaces, caregiving credits, artistic fellowships, and open-source stewardship—explicitly accounted within broad societal $W_{\mathrm{soc}}(\varepsilon)$—ensures material security reinforces purposeful human agency.
5) *Asymmetric Automation Tax Reform and Craft Differentiation:* Eliminating tax asymmetries that subsidize directive labor-replacing capital over human training [13], encouraging full machine substitution in hazardous low-meaning drudgery (Case 2), and supporting transparent provenance for human-centric relational and artisanal services preserve plural pathways to high-agency work (Case 3).

VIII. Limitations, Open Frontiers, Experimental Protocol, and Conclusion

A. Methodological Limitations

Several empirical limitations warrant transparent discussion. First, although linking O*NET 30.1, Anthropic Economic Index telemetry, and repeated NORC GSS cross-sections enables multi-scale estimation across 873 occupations and $10{,}483$ respondents, cross-sectional regressions cannot completely rule out unobserved worker sorting across occupations (e.g., intrinsically motivated individuals selecting into high-skill roles). Second, 6-digit SOC crosswalks aggregate heterogeneous firm-level workflows within an occupation; two software engineers or warehouse associates sharing an SOC code may experience markedly different managerial pacing regimes. Third, while AEI telemetry captures over one million real-world interactions across `claude_ai` and `1p_api`, single-provider LLM telemetry under-represents proprietary vertical software and legacy industrial programmable logic controllers.

B. Mathematical Extensions and Open Theoretical Frontiers

Beyond the globally concave regulated family (11), four mathematical extensions merit dedicated investigation:

1) *Morse Theory, Bifurcations, and Persistent Homology:* In multi-modal industries where network externalities or institutional lock-in violate (12), $W(x)$ forms a non-convex Morse landscape with multiple local maxima separated by index-1 saddles. Tracking saddle-node bifurcations under exogenous AI cost shocks and computing persistent homology over superlevel filtrations $\{x \in X : W(x) \geq c\}$ can quantify the topological barriers separating low-autonomy traps from high-synergy equilibria.
2) *Multi-Agent Non-Zero-Sum Game Theory:* Whereas Section II aggregates team interactions via hypergraph potential (10), strategic workplaces involve workers, managers, and autonomous AI vendors optimizing distinct payoffs $U_a(\boldsymbol{z})$. Characterizing the price of anarchy between decentralized Nash equilibria and Pareto-optimal MWF states—and designing cooperative Shapley or Nash bargaining contracts—connects socio-technical potential theory to strategic industrial organization.
3) *Risk-Sensitive Stochastic Control and Contextual Bandits:* In safety-critical operations where tail drops in vigilance are catastrophic, replacing expectation-neutral objectives (19) with entropic risk-sensitive HJB control $(-\gamma_{\mathrm{risk}}^{-1} \ln \mathbb{E}[\exp(-\gamma_{\mathrm{risk}} \mathcal{V})])$ enforces high-probability barrier constraints $W(x_t) \geq W_{\min}$, while contextual multi-armed bandit algorithms enable online exploration of unknown person-specific complementarities Γ_{ij} with bounded regret.
4) *Non-Commutative Operator Generalizations:* Psychometric and managerial observations are often context- and order-dependent: auditing quantitative output before soliciting creative input alters a worker's cognitive framing differently than reversing the sequence. Representing incompatible socio-technical observables via non-commutative operators $[\hat{\mathcal{O}}_H, \hat{\mathcal{O}}_M] \neq 0$ on cognitive state spaces offers a formal language for measurement back-action in human-in-the-loop governance.

C. Pre-Registered Organizational Field-Trial Protocol

To establish causal identification of both Γ_{ij} and Sivak–Crooks geodesic rollout savings inside operating enterprises, our framework prescribes a two-by-two stepped-wedge cluster-randomized trial across $J \geq 40$ organizational teams over 24 weeks:

1) *Factorial Treatment Arms:* Randomize teams along two orthogonal axes: *Interface Architecture* (interactive copilot with human override and transparent rationale vs. directive automated queue assignment) $\times$ *Rollout Pacing* (abrupt week-1 cutover vs. variance-adaptive Sivak–Crooks geodesic schedule (30) calibrated from biweekly pilot friction $\hat{\zeta}(\lambda)$).
2) *Pre-Registered Outcomes:* Track quality-adjusted output Y_{jt}, supervisory error rates $\mathcal{D}_{\mathrm{err}}$, biweekly validated short-form meaningful work and burnout scales (WAMI [29], UWES-9 [28], PHQ-2 [67]), and voluntary turnover, avoiding high-powered incentives on raw telemetry counts in accordance with Proposition 2.

D. Conclusion

Technological progress does not possess a predetermined sign for human flourishing. By formulating the Master Work Function on rigorous axiomatic, differential-geometric, and stochastic-thermodynamic foundations, this work demonstrates that human agency and machine capability are coupled through a measurable cross-partial interaction structure Γ_{ij} and finite-time transition friction $\zeta(\lambda)$. Across 873 U.S. occupations, 124.0 million workers, and three waves of national

worker microdata, separability between human fulfillment and machine design is decisively rejected, interactive human-in-the-loop architectures sustain nearly fivefold higher collaborative augmentation than headless directive APIs, meaningful work provides a robust buffer against depressive anhedonia independent of income, geodesic deployment schedules eliminate nearly one-third of reorganization dissipation, and MWF-aligned governance expands both economic output and distributional dignity. Treating the Master Work Function as a first-class engineering and policy design objective offers a practical, mathematically principled path toward socio-technical systems that amplify rather than diminish the human condition.

Data Availability and Acknowledgments

All empirical datasets analyzed in this study are publicly available from their respective custodians: the O*NET 30.1 Database by the U.S. Department of Labor, Employment and Training Administration (USDOL/ETA, used under CC BY 4.0; O*NET® is a trademark of USDOL/ETA, which has not endorsed or tested modifications reported here); the Anthropic Economic Index and Robot Exposure datasets (released under CC-BY 4.0); the NORC General Social Survey at the University of Chicago; Eloundou *et al.* [17] and Felten *et al.* [16] public repositories; and the U.S. Bureau of Labor Statistics OEWS program. Complete Python data-retrieval, econometric estimation, and figure-generation scripts are archived in the companion repository directory.

References

[1] N. Wiener, *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA, USA: MIT Press, 1948.

[2] ——, *The Human Use of Human Beings: Cybernetics and Society*. Boston, MA, USA: Houghton Mifflin, 1950.

[3] K. Marx, *Economic and Philosophic Manuscripts of 1844*. Moscow: Progress Publishers, 1959.

[4] H. Arendt, *The Human Condition*. Chicago, IL, USA: University of Chicago Press, 1958.

[5] M. Heidegger, *The Question Concerning Technology and Other Essays*. New York, NY, USA: Harper & Row, 1977.

[6] J. Rawls, *A Theory of Justice*. Cambridge, MA, USA: Harvard University Press, 1971.

[7] ——, *The Law of Peoples*. Cambridge, MA, USA: Harvard University Press, 1999.

[8] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[9] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.

[10] C. B. Frey and M. A. Osborne, "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, vol. 114, pp. 254–280, 2017.

[11] D. H. Autor, "Why are there still so many jobs? The history and future of workplace automation," *Journal of Economic Perspectives*, vol. 29, no. 3, pp. 3–30, 2015.

[12] D. Acemoglu and P. Restrepo, "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, vol. 108, no. 6, pp. 1488–1542, 2018.

[13] ——, "Automation and new tasks: How technology displaces and reinstates labor," *Journal of Economic Perspectives*, vol. 33, no. 2, pp. 3–30, 2019.

[14] ——, "Robots and jobs: Evidence from US labor markets," *Journal of Political Economy*, vol. 128, no. 6, pp. 2188–2244, 2020.

[15] E. Brynjolfsson and T. Mitchell, "What can machine learning do? Workforce implications," *Science*, vol. 358, no. 6370, pp. 1530–1534, 2017.

[16] E. W. Felten, M. Raj, and R. Seamans, "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses," *Strategic Management Journal*, vol. 42, no. 12, pp. 2195–2217, 2021.

[17] T. Eloundou, S. Manning, P. Mishkin, and D. Rock, "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, vol. 384, no. 6702, pp. 1306–1308, 2024.

[18] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023.

[19] E. Brynjolfsson, D. Li, and L. Raymond, "Generative AI at work," *The Quarterly Journal of Economics*, vol. 140, no. 2, pp. 889–942, 2025.

[20] K. Handa, A. Tamkin, M. McCain, S. Huang, E. Durmus, S. Heck, J. Mueller, J. Hong, S. Ritchie, T. Belonax, K. K. Troy, D. Amodei, J. Kaplan, J. Clark, and D. Ganguli, "Which economic tasks are performed with AI? Evidence from millions of Claude conversations," arXiv preprint arXiv:2503.04761, 2025, https://arxiv.org/abs/2503.04761.

[21] R. Appel, P. McCrory, A. Tamkin, M. McCain, T. Neylon, and M. Stern, "Anthropic Economic Index report: Uneven geographic and enterprise AI adoption," arXiv preprint arXiv:2511.15080, 2025, https://arxiv.org/abs/2511.15080.

[22] M. Massenkoff, E. Lyubich, S. Sacher, Z. Hitzig, S. Zhang, R. Heller, and P. McCrory, "Anthropic Economic Index report: Cadences," Anthropic Economic Research Report, 2026, June 2026 Release, https://www.anthropic.com/research/economic-index-june-2026-report.

[23] J. Fábrega, "Governing delegation to generative artificial intelligence: Human direction, work-related orientation, and modes of use," arXiv preprint arXiv:2608.17624, 2026, https://arxiv.org/abs/2608.17624.

[24] R. Jadhav and J. Danve, "The AI skills shift: Mapping skill obsolescence, emergence, and transition pathways in the LLM era," arXiv preprint arXiv:2604.06906, 2026, https://arxiv.org/abs/2604.06906.

[25] R. Legate-Yang and M. Massenkoff, "What work can robots do?" Anthropic Economic Research Report, 2026, September 2026 Release, https://www.anthropic.com/research/what-work-can-robots-do.

[26] J. R. Hackman and G. R. Oldham, "Motivation through the design of work: Test of a theory," *Organizational Behavior and Human Performance*, vol. 16, no. 2, pp. 250–279, 1976.

[27] E. L. Deci and R. M. Ryan, "The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior," *Psychological Inquiry*, vol. 11, no. 4, pp. 227–268, 2000.

[28] W. B. Schaufeli, A. B. Bakker, and M. Salanova, "The measurement of work engagement with a short questionnaire: A cross-national study," *Educational and Psychological Measurement*, vol. 66, no. 4, pp. 701–716, 2006.

[29] M. F. Steger, B. J. Dik, and R. D. Duffy, "Measuring meaningful work: The Work and Meaning Inventory (WAMI)," *Journal of Career Assessment*, vol. 20, no. 3, pp. 322–337, 2012.

[30] G. Debreu, "Representation of a preference ordering by a numerical function," in *Decision Processes*, R. M. Thrall, C. H. Coombs, and R. L. Davis, Eds. New York, NY, USA: John Wiley & Sons, 1954, pp. 159–165.

[31] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton, NJ, USA: Princeton University Press, 1944.

[32] F. Y. Kuo, I. H. Sloan, G. W. Wasilkowski, and H. Woźniakowski, "On decompositions of multivariate functions," *Mathematics of Computation*, vol. 79, no. 270, pp. 953–966, 2010.

[33] D. M. Topkis, *Supermodularity and Complementarity*. Princeton, NJ, USA: Princeton University Press, 1998.

[34] P. Milgrom and C. Shannon, "Monotone comparative statics," *Econometrica*, vol. 62, no. 1, pp. 157–180, 1994.

[35] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko, *The Mathematical Theory of Optimal Processes*. New York, NY, USA: Interscience Publishers, 1962.

[36] I. Prigogine, *The End of Certainty: Time, Chaos, and the New Laws of Nature*. New York, NY, USA: Free Press, 1997.

[37] U. Seifert, "Stochastic thermodynamics, fluctuation theorems and molecular machines," *Reports on Progress in Physics*, vol. 75, no. 12, p. 126001, 2012.

[38] P. A. Ortega and D. A. Braun, "Thermodynamics as a theory of decision-making with information-processing costs," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 469, no. 2153, p. 20120683, 2013.

[39] P. Hack, S. Gottwald, and D. A. Braun, "Jarzynski's equality and Crooks' fluctuation theorem for general Markov chains with application to decision-making systems," *Entropy*, vol. 24, no. 12, p. 1731, 2022.

[40] C. Jarzynski, "Nonequilibrium equality for free energy differences," *Physical Review Letters*, vol. 78, no. 14, pp. 2690–2693, 1997.

[41] G. E. Crooks, “Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences,” *Physical Review E*, vol. 60, no. 3, pp. 2721–2726, 1999.

[42] D. A. Sivak and G. E. Crooks, “Thermodynamic metrics and optimal paths,” *Physical Review Letters*, vol. 108, no. 19, p. 190602, 2012.

[43] A. B. Atkinson, “On the measurement of inequality,” *Journal of Economic Theory*, vol. 2, no. 3, pp. 244–263, 1970.

[44] A. Sen, *Development as Freedom*. Oxford, UK: Oxford University Press, 1999.

[45] B. Holmstrom and P. Milgrom, “Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design,” *The Journal of Law, Economics, and Organization*, vol. 7, no. Special Issue, pp. 24–52, 1991.

[46] C. A. E. Goodhart, “Problems of monetary management: The UK experience,” in *Monetary Theory and Practice*. London, UK: Macmillan, 1984, pp. 91–121.

[47] National Center for O*NET Development, “O*NET 30.1 database,” U.S. Department of Labor, Employment and Training Administration (USDOL/ETA), 2025, CC BY 4.0 License, https://www.onetcenter.org/database.html.

[48] M. Davern, R. Bautista, J. Freese, P. Herd, and S. L. Morgan, “General Social Survey 1972–2024 cross-sectional cumulative data,” NORC at the University of Chicago, 2025, https://gss.norc.org/.

[49] U.S. Bureau of Labor Statistics, “Occupational Employment and Wage Statistics (OEWS) national research estimates,” U.S. Department of Labor, 2023, https://www.bls.gov/oes/.

[50] E. Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*. Evanston, IL, USA: Northwestern University Press, 1970.

[51] Aristotle, *Nicomachean Ethics*, 2nd ed. Indianapolis, IN, USA: Hackett Publishing, 1999.

[52] V. E. Frankl, *Man’s Search for Meaning*. Boston, MA, USA: Beacon Press, 1959.

[53] I. Kant, *Groundwork of the Metaphysics of Morals*, M. Gregor, Ed. Cambridge, UK: Cambridge University Press, 1998.

[54] M. C. Nussbaum, *Creating Capabilities: The Human Development Approach*. Cambridge, MA, USA: Harvard University Press, 2011.

[55] S. Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*. New York, NY, USA: Viking, 2019.

[56] M. Foucault, *Discipline and Punish: The Birth of the Prison*. New York, NY, USA: Pantheon Books, 1977.

[57] ——, *The Birth of Biopolitics: Lectures at the Collège de France, 1978–1979*. Basingstoke, UK: Palgrave Macmillan, 2008.

[58] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.

[59] E. H. Clarke, “Multipart pricing of public goods,” *Public Choice*, vol. 11, no. 1, pp. 17–33, 1971.

[60] T. Groves, “Incentives in teams,” *Econometrica*, vol. 41, no. 4, pp. 617–631, 1973.

[61] G. Hoffman, “Evaluating fluency in human–robot collaboration,” *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 3, pp. 209–218, 2019.

[62] M. Mitchell, “How do we know how smart AI systems are?” *Science*, vol. 381, no. 6654, p. adj5957, 2023.

[63] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, “Cooperative inverse reinforcement learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 29, 2016, pp. 3909–3917.

[64] L. S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA, USA: Harvard University Press, 1978.

[65] J. G. MacKinnon and H. White, “Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties,” *Journal of Econometrics*, vol. 29, no. 3, pp. 305–325, 1985.

[66] K.-Y. Liang and S. L. Zeger, “Longitudinal data analysis using generalized linear models,” *Biometrika*, vol. 73, no. 1, pp. 13–22, 1986.

[67] K. Kroenke, R. L. Spitzer, and J. B. W. Williams, “The Patient Health Questionnaire-2: Validity of a two-item depression screener,” *Medical Care*, vol. 41, no. 11, pp. 1284–1292, 2003.

[68] J. M. Keynes, “Economic possibilities for our grandchildren,” in *Essays in Persuasion*. London, UK: Macmillan, 1931, pp. 358–373.