Information Acquisition and Knowledge Transmission in the Age of AI*

Michael Sockin†

May 2026

Abstract

I construct a model of rational inattention to study how artificial intelligence (AI) impacts what information humans acquire, and how these choices feed back into AI expertise. I show that AI represents a biased information intermediary whose anisotropic precision about underlying fundamentals is shaped by the history of tasks on which it trains and inflated by fine-tuning. By substituting for learning about fundamentals, AI frees scarce human attention to focus on task-specific learning, hastening the discovery of new tasks that temporarily reduce its usefulness and adoption. I consider implications for sustainable AI design and knowledge collapse when AI evaluates humans.

*I thank Aaron Pancost and Andrey Ordin for helpful comments.

†University of Texas at Austin. Email: Michael.Sockin@mccombs.utexas.edu.

As artificial intelligence (AI) diffuses throughout the U.S. economy, it continues to impact how humans acquire information and apply it to decision-making.[1] The expertise that it shares, but does not confer and may even exaggerate, has reshaped human performance at existing tasks and broadened their exploration for novel uses of knowledge. In turn, how humans use AI shapes its expertise through the tasks on which it trains and the information that humans choose to acquire. This raises important questions about how AI, as a backward-looking and biased information intermediary, learns as it shapes the economic and informational environment, and impacts the transmission of knowledge through its adoption.

In this paper, I study the feedback loop between AI's knowledge accumulation, human information acquisition, its adoption, and the discovery of new tasks in the economy. I construct a model of rational inattention in which an agent born at time t must complete an uncertain task at time $t+1$, and then exits the economy. This task is a linear combination of unobservable economic fundamentals whose weights are also unknown and a function of a hidden task state. The agent has limited attention capacity to devote to learning about these two sources of uncertainty and derives utility that is a quadratic loss based on her effort and its ideal level. Furthermore, the attention that she devotes to learning about her task can lead to a jump in the set of underlying tasks in the economy for the next agent. Imposing an independence of posteriors, and assuming a Gaussian prior over the fundamentals and a uniform prior over the task state, I use variational methods to derive closed-form solutions for the optimal filters - a Kalman-Bucy filter for fundamentals and a controlled Wonham filter for tasks - and a water-filling strategy for how to allocate attention between the two channels.

I then construct an artificial intelligence (AI) as a biased information intermediary. I assume that it aggregates the past information generated by agents about the economy-wide fundamentals, that some of its training data degrades in the sense of Blackwell garbling (Blackwell, 1953), and that it is post-processed with Reinforcement Learning from Human Feedback (RLHF) using a reward-based tilt subject to a divergence penalty (Korbak et al., 2022). I provide a general characterization of the concatenation, garbling, and optimal tilt. Since agents learn about fundamentals using the Kalman-Bucy filter, the AI's posterior is Gaussian with a precision that depends on the prior precision, the precision from the training data, and on the RLHF tilt. Importantly, the AI's training-data precision is anisotropic in the sense that it is greatest for the

[1]For instance, 18% of firms surveyed by the U.S. Census in 2026 used AI in at least one function (Bonney et al., 2025). Korinek and McKelvey (2026) estimate that in 2025, the US added $251 billion in nominal US Gross Domestic Production from AI and consumed $219 billion in computing resources.

fundamentals on which agents' tasks have historically loaded. In addition, the RLHF tilt inflates the AI's actual precision since it imbues confidence about the fundamentals on which the tilt loads, but the AI does not actually possess more information about them.

I then study how AI impacts agent decision-making. I show that its usage can lead to not only "AI dependence", i.e., that agents who buy an AI subscription must continue to because its underlying information remains with the AI rather than is conveyed to the agent, but that it acts as a biased information intermediary that induces agents who cannot disentangle its RLHF tilt from actual precision to buy a subscription. I also find that AI usage naturally gives rise to invention and innovation cycles of adoption. As the AI becomes more precise based on past training data for the current span of tasks, AI usage increases. However, because this shifts agent's attention to learning about their tasks, this can also lead to the discovery of new (and antiquation of old) tasks. This reduces the relevance of the AI's training data and adoption declines until the AI's precision recovers sufficiently to justify its subscription fee.

Finally, I consider two extensions of my model. In the first, I study the optimal design of AI. I show that a developer who aims to maximize profits statically balances the subscription fee, equal to the surplus created from using AI, with the rising costs of computing (e.g., electricity, water). Dynamically, the developer weighs the subscription fee against the reduction in attention to the fundamentals that AI induces. This degrades the AI's future training data, which also mitigates future computing costs. In addition, it increases the likelihood that new tasks will be discovered that render the AI less valuable. The developer uses the RLHF tilt and temperature as persuasion tools to induce agents to buy subscriptions by inflating the AI's perceived precision for certain tasks. Compared to a developer, a social planner avoids a RLHF tilt but can face a trade-off between fostering the pace of discovery with AI usage against the erosion of the AI's value as a backward-looking information intermediary.

In the second, I show that a collapse in human knowledge accumulation can occur not from economic agents using AI but from being evaluated by it. Because the AI's judgment is based on past information about fundamentals, agents have less incentive to acquire new information to match the AI's stale beliefs. This can lead to a learning trap in which agents neither have incentive to adopt AI nor acquire any information about fundamentals, the AI's posterior is frozen, and agents receive a high payoff from perfectly matching the AI's assessment of their optimal effort. Such a result is akin to the over-reliance on public information reminiscent of Keynesian beauty contests (e.g., Allen et al., 2006).

A methodological contribution of this paper is to provide a variational approach to information acquisition on the Wiener space by framing the optimal filters as adapted perturbations of identity. A key finding is that when the agent's objective itself depends on the optimal filter, i.e., through the conditional expectation of the state, this introduces an additional endogenous term into the Föllmer drift that locally characterizes the optimal change of measure. In the special case of a quadratic objective, this additional term vanishes and the optimal filter is tractable, yielding in a single framework the Kalman-Bucy filter for a Gaussian prior and a controlled Wonham filter whose gain adjusts dynamically to reach a target posterior variance when the prior is discrete. However, this general characterization in principle can be extended to nonlinear settings in which the optimal filter is then a fixed-point problem.

My paper relates to the growing literature that explores how artificial intelligence interacts with human decision-making. Agrawal et al. (2018b) examines whether AI as a prediction technology is a complement or substitute to human judgment, while Agrawal et al. (2025) models AI as a substitute for implementation effort but a complement for judgment in identifying task improvements. Ide and Talamás (2025) explore the capacity for AI to substitute for humans at different knowledge levels within a firm. Zhong (2025) studies the optimal integration of AI into multi-layered decision-making, and Chen et al. (2026) explore how deploying AI as autonomous or assistive determines its expertise based on how it impacts agents' decision-making. My focus instead is on how AI impacts what information economic agents acquire, how this feeds back into the AI's expertise and adoption over time, and how this affects the pace of discovery of new tasks.

In the context of asymmetric information, Athey et al. (2020) and Chemla and Tena (2026) study how AI interacts with moral hazard in principal-agent settings, the latter in general equilibrium. Huang and Ouyang (2025) model AI as an unbiased advisor compared to humans that has limited processing capacity, while Dreyfuss and Hoong (2025) study the optimal coarsening of AI information for decision-making under behavioral biases and Agarwal et al. (2025) derive a sufficient statistic for when human versus AI decision-making is optimal. By contrast to these papers, I model the role of AI as a biased information intermediary that aids agents in acquiring information about economic fundamentals. In addition, I consider how a developer who internalizes these externalities would price AI when he can bias it to persuade agents to adopt it, and show how an agent that uses AI may need to continue using it because she does not acquire the underlying information. In addition, the simple functional form that I derive for the AI's posterior can potentially be used in other settings to study the informational role of AI.

Several papers specifically consider how AI impacts the transmission of knowledge. Ide (2025) finds that AI displacement of early career opportunities can erode long-term skill acquisition, while Afrouzi et al. (2026) shows that AI can induce two stationary equilibria (one with high and one with low learning) when it substitutes for labor but augments managers in creating new frontier products. Farboodi et al. (2025) provides conditions under which data accumulation that improves capital productivity, such as AI, leads to a fully-automated economy. Shumailov et al. (2024) show that training an AI recursively on generated data causes it to collapse around the head of the distribution, while Acemoglu et al. (2026b) show how introducing AI as a public signal can degrade information aggregation on a network based on its speed of updating. Directly examining its impact on information acquisition, Acemoglu et al. (2026a) show how AI can lead to a collapse in social learning if it substitutes for effort to acquire information when that effort contributes to the provision of public information. I instead show that AI can amplify cycles of invention and innovation in which knowledge accumulation never collapses but fluctuates over time. My analysis also suggests it is not the use but evaluation by AI that can lead to such a collapse.

My work also contributes to the literature on rational inattention in the tradition of Sims (2003).[2] Among dynamic settings (e.g., Sims, 2010), Moscarini (2004) studies the optimal infrequent sampling of a Kalman-Bucy filter to explain inertia in firm behavior, Luo (2008) examines the implications of rational inattention for consumption dynamics in a multivariate permanent income model, and Maćkowiak and Wiederholt (2009) analyze how firms allocate attention between aggregate versus idiosyncratic conditions to explain dynamic pricing behavior. Maćkowiak et al. (2018) and Miao et al. (2022) show that agents may optimally acquire signals that are lower-dimensional than the full Markov state and multivariate state spaces, respectively, the latter also in the presence of multiple controls. Zhong (2022) characterizes the optimal signal structure in an optimal stopping-time problem using dynamic programming, in which the state is a Markov jump-diffusion measure, and highlights a novel precision-frequency trade-off. Hébert and Woodford (2023) use dynamic programming to characterize for which information cost functions gradual (diffusion) versus discrete (jump) learning is optimal. By contrast, I use functional analysis on the Wiener space to characterize the optimal filter in terms of optimal transport processes. I further identify a key source of endogeneity for solving for optimal nonlinear filters with non-quadratic objectives. While Hansen et al. (2025) jointly study rational inattention and a Donsker-Varadhan tilt (for ambiguity aversion) in a static setting, I consider a dynamic model in which the two are separated across actors (i.e., agents and the AI).

[2]See Maćkowiak et al. (2023) for a recent review of this literature.

1 A Model of Information Acquisition without Artificial Intelligence

Time is continuous $t \geq 0$ and the horizon infinite. There is a generational structure of economic agents. Let $\mathcal{G} = \mathcal{C}_0(\mathbb{R}_+, \mathbb{R}^2)$ be the classical Wiener space equipped with the canonical Wiener measure μ_W that is copied for each agent i. I denote by $(\boldsymbol{W}_t, t \geq 0)$ a 2×1 vector of canonical Brownian motions with corresponding filtration $\mathbb{W} = (\mathcal{W}_t, t \geq 0)$, and the associated Cameron-Martin space by $H = \{h \in \mathcal{G} \colon h(t) = \int_0^t \dot{h}(s)ds, \dot{h} \in L^2(dt)\}$. provided with the norm $|h|_H = |\dot{h}|_{L^2(dt)}$. As a Hilbert space, this naturally induces an inner product on the Cameron-Martin space $\langle g, h \rangle$. In what follows, bolded objects indicate vectors.

The economy is governed by N fundamentals, f_n, for $n \in \{1, ..., N\}$, that are unobservable to all economic agents that I stack into a vector $\boldsymbol{f}$. In addition, there is a set of $\mathcal{S}$ possible tasks that an agent can complete to produce output. A task $s \in \mathcal{S}$ is a $N \times 1$ vector $\boldsymbol{a}(s)$ of nonnegative weights that is an element of the $N-1$-dimensional unit simplex, Δ^{N-1} (where the last entry $a_N(s) = 1 - \sum_{n=1}^{N-1} a_n(s)$ such that $|\boldsymbol{a}(s)| \equiv 1$).[3] For instance, the fundamentals may be mathematics, physics, coding, and writing, and the task could be to produce a research paper on quantum computing or to derive a model in quantum field theory. Consequently, $a : \mathcal{S} \to \Delta^{N-1}$ maps a task to its weights on the fundamentals, and gives rise to a productivity functional for an agent, $\boldsymbol{a}(s)' \boldsymbol{f}$, in which the weights on a subset of fundamentals can be zero. The scope of potential tasks, $\mathcal{S}_t$, can change over time as new tasks are discovered and old ones become obsolete. Let $\mathcal{S}_0$ be the initial set of tasks available to all agents.[4]

Agents in the economy are born in the economy sequentially, such that there is only one agent at a given instant. Agent i born at time τ_i lives until $\tau_i + 1$, during which time she can acquire information about the fundamentals, $\boldsymbol{f}$, and the random task assigned to her drawn from the current set of potential tasks when she is born, $\mathcal{S}_i = \mathcal{S}_{\tau_i}$, subject to limited attention. At time $\tau_i + 1$, she supplies effort to complete her task and is awarded a payoff based on how well she executes it. She then exits the economy and is replaced by a new agent. Although $\tau_i = i$, I use separate notation to avoid confusion between indexing agents and indexing time.

In the absence of artificial intelligence, information acquired by agents does not travel across agents or across time, and consequently each solves a static, identical decision problem.[5] However,

[3] I could allow the weights to lie on an arbitrary simplex through a dilation of the vertices of the unit simplex.

[4] For instance, Autor et al. (2024) document the emergence of new work that shifted from middle-paid production and clerical occupations in 1940-1980 to high-paid professional and lower-paid services from 1980-2018. Loaiza and Rigobón (2024) further find at the task level that from 2016 to 2024, 815 new tasks were added to O*NET while 1310 were retired.

[5] That agents do not observe past public information is not essential, and can be relaxed as long as they observe only

the attention capacity that an agent devotes to learning about her task, which I refer to as κ_s^i, can contribute to invention by changing the set of tasks available to future agents, i.e., altering the set $\mathcal{S}_t$ for $t > \tau_i$. New tasks arrive and old tasks simultaneously become antiquated according to a Poisson point process with time-varying compensator $\zeta\left(\kappa_s^i\right) dt$ over the life of agent i, where $\zeta\left(\kappa_s^i\right) \in \mathcal{C}^1\left(\mathbb{R}_+\right)$ increasing in κ_s^i. The creation of new tasks can be viewed as invention that comes from agents discovering new ways that fundamentals may interact by devoting attention to learning about their task. This is a key difference between agents and artificial intelligence: while artificial intelligence has (limited but) immediate access to all past information that agents have created, it merely innovates but does not itself invent (although it may be prompted to help discover new inventions). I capture this invention process that agents synthesize through jumps in my diffusion learning framework.[6]

I first describe and characterize the problem of an individual agent in the economy. I then introduce artificial intelligence (AI) as an information acquisition tool and discuss how agents' information acquisition choices are linked over time. This will give rise to AI dependence and to two externalities: agents do not internalize how the information they acquire is used by future AI users nor how the jumps in the task space they induce affect the value of the AI's data. The latter will give rise to invention and innovation cycles.

1.1 The Problem of Agent i

An agent i born at time τ_i is assigned to work on a random task $s_i \in \mathcal{S}_i$ from time τ_i to $\tau_i + 1$, where $\mathcal{S}_i = \mathcal{S}_{\tau_i}$ is the set of tasks available to all agents when agent i is born with cardinality $|\mathcal{S}_i|$. A newborn agent has a Gaussian prior over $\boldsymbol{f}$, $\mu_f = \mathcal{N}\left(\bar{\boldsymbol{f}}, \Omega_0\right)$, and a uniform prior over their optimal task given the current space of available tasks, $\bar{\boldsymbol{\pi}} = \frac{1}{|\mathcal{S}_i|}\iota_{|\mathcal{S}_i|}$ (where ι is the $|\mathcal{S}_i| \times 1$ vector of ones), that are independent from each other. I define the base measure, $\mu = \mu_1 \otimes \mu_2$, to be the product of the Wiener measure over W_1 and of the Wiener measure over W_2. All agents have access to the same information set, $\mathcal{W}_t$, that accumulates these two Brownian motions over time.

In what follows, I describe the problem and then characterize the optimal policies of agent i

a subset of the history. This assumption, however, simplifies my analysis and is consistent with the rational inattention paradigm (e.g., Sims, 2003, 2010) that learning from any source requires the allocation of attention. A key benefit of AI is that it delivers a posterior rather than raw public data by performing the costly information processing for the agent.

[6]New tasks reflect novel ways of combining underlying economy-wide fundamentals, and arise as an externality of agents' information acquisition activities rather than through deliberate investment. This is in contrast to the literature that views creative destruction as the expansion of the technological or production frontier (e.g., Romer, 1990) or quality ladder (e.g., Aghion and Howitt, 1992). In that tradition, Acemoglu and Restrepo (2019a,b) argue the impact of artificial intelligence on labor demand and the wage bill depends on whether the crowding out from automation is offset by the productivity gains it imparts and the arrival of new tasks to which the displaced labor can be assigned.

that lives between times τ_i and $\tau_i + 1$.

1.1.1 Agent Effort

At time $\tau_i + 1$, the agent must complete the task by choosing a level of effort $K^i \in \mathbb{R}_+$. The ideal level of effort is $\boldsymbol{a}(s_i)' \boldsymbol{f}$, and she derives a quadratic loss from the mismatch between this ideal level and the level that she chooses. However, not only does she not observe the underlying fundamentals, $\boldsymbol{f}$, she also does not know which specific task she must complete until after exerting effort at time $\tau_i + 1$. This will motivate her to learn about not only the underlying fundamentals but also about her task.

The payoff to agent i, U^i, is measured by this quadratic loss at time $\tau_i + 1$

$$U^i = -\frac{1}{2}\left(\boldsymbol{a}(s_i)' \boldsymbol{f} - K^i\right)^2. \tag{1}$$

After completing her task, agent i exits the economy and agent $i + 1$ is born.

1.1.2 Information Acquisition

Prior to the maturity of her task at time $\tau_i + 1$, agent i can acquire information about both (some subset of) the underlying fundamentals, $\boldsymbol{f}$, and about the nature of her task, i.e., which $s \in \mathcal{S}_i$ governs the weights on the fundamentals of her task. Her information acquisition choices determine how informative is her optimal change of measure, ν^i. However, the total amount of information she can acquire about both sources of uncertainty is limited, and she rationally chooses how to be inattentive toward available information in the spirit of Sims (2003).

Let ν^i be the agent's posterior about the fundamentals and her task at time t subject to her information set $\mathcal{W}_t$, with first and second marginals, ν_1^i and ν_2^i, respectively. Consistent with the literature on rational inattention, I assume that the agent chooses a family of posteriors ν^i subject to an attention constraint that the mutual information between her prior μ and her final posterior ν^i at time $\tau_i + 1$. This is equivalent to imposing a bound or capacity constraint, κ, on the Kullback-Leibler divergence between her prior and posterior, $\mathbb{D}\left(\nu^i || \mu\right)$, such that

$$\mathbb{D}\left(\nu^i || \mu\right) = \mathbb{E}_{\nu^i}\left[\log \frac{d\nu^i}{d\mu} \mid \mathcal{W}_{\tau_i}\right] \leq \kappa. \tag{2}$$

In what follows, I restrict attention to continuous learning over time and assume that agents

learn about the aggregate fundamentals, $\boldsymbol{f}$, independently from their idiosyncratic task, i.e., the specific s governing their task weights. This latter assumption of independence of posteriors captures that learning about fundamentals, which are general concepts that are transferable across tasks, is qualitatively different from learning about one's specific task assignment. For instance, the fundamentals could represent general ideas or skills that are learned either formally through study or informally from working with other agents. By contrast, an agent's task is specific to them, and it is not obvious what resource or channel for acquiring information would correspond to learning about the exact nature of an agent's work, i.e., $\boldsymbol{a}(s_i)' \boldsymbol{f}$. This distinction will not only make learning tractable, it will also make the trade-off between learning about fundamentals and about task weights transparent and interpretable.[7]

To characterize agent i's information acquisition problem, I must first define a larger filtration specific to agent i, $\mathbb{B}^i = (\mathcal{B}_t^i, t \geq 0)$, that includes the vector of canonical Brownian motions and the realizations of task and fundamentals, i.e., $\mathcal{B}_t^i = \mathcal{W}_t \vee \sigma(\boldsymbol{f}, s_i)$. I can then construct the filter on this base filtration $\mathbb{B}^i$, and then project it to the coarser observation filtration $\mathbb{W}$ since $\mathcal{W}_t \subseteq \mathcal{B}_t^i$. Importantly, $\mathbb{B}^i$ is endogenous because it is constructed by making the Brownian motions W_t informative signals about the task and states. A key technical issue that requires care is distinguishing between changes of measures fixing the filtration, i.e., measures μ and ν, and changes in filtrations, i.e., from $\mathcal{B}_t^i$ to $\mathcal{W}_t$.

It will next be useful to consider several consequences of the Cameron-Martin Theorem that translations of the Wiener measure with elements of H induce measures equivalent to μ, and the push-forward of this measure (to an equivalent measure $Q^i\mu$) can be defined as an adapted perturbation of identity, that is $Q^i = I_W + q^i$ for some element $q^i \in L_a^2(\mu, H)$ with $\dot{q}^i \in L_a^2(dt \times d\mu)$. I refer to the μ-semimartingale (with respect to $\mathcal{W}_t$) Q^i as a transport process. The left inverse of this map, $\Phi^i = I_W + \phi^i$ for some $\phi : W \to H$ such that $\Phi^i \circ Q^i = I_W$, if exists, defines the appropriate exponential martingale, $L\left(\phi^i\right)$, that is the Radon-Nikodym derivative for the change of measure under μ adapted to $\mathcal{W}_t$

$$\frac{dQ^i\mu}{d\mu} = L\left(\phi^i\right) = \rho\left(-\delta\phi^i\right), \tag{3}$$

[7]This is similar to the independence assumption in Maćkowiak and Wiederholt (2009) in which firms learn independently about aggregate versus idiosyncratic economic conditions. Without this assumption, it is possible that the optimal nonlinear filter fails to satisfy the innovation conjecture, in which case there is information loss between the AI and agents' information sets because of non-invertibility. Independence of posteriors is distinct from posterior separability in cost functions (e.g., Caplin et al., 2022)

where

$$\rho\left(-\delta\phi^i\right) = \exp\left(-\delta\left(\phi^i\right) - \frac{1}{2}|\phi^i|_H^2\right), \tag{4}$$

is the Girsanov exponential and $\delta\left(\phi^i\right) = \int_{\tau_i}^{\tau_i+1} \dot{\phi}_s^i dW_{1s}$ and $|\phi^i|_H^2 = \int_{\tau_i}^{\tau_i+1} |\dot{\phi}_s^i|^2 ds$. In Equation (4), δ is often referred to as the Skorokhod integral.[8] By the Cameron-Martin Theorem, under this inverse for $t = [\tau_i, \tau_{i+1}]$

$$\hat{B}_{1t}^i = W_{1t} + \int_{\tau_i}^t \dot{\phi}_s^i ds, \tag{5}$$

is a standard Brownian motion under $Q^i\mu_1$ with respect to $\mathcal{B}_t^i$.

Because the space of Gaussian measures defined by these translation maps $w \to w + h$ for any $h \in H$ is dense in the Cameron-Martin space, I can consequently move between measures with Girsanov's Theorem by defining the appropriate Girsanov exponential. Similarly, I can define a transport process $V^i = I_W + v^i$ with $v^i \in L_a^2(\mu, H)$ such that for $t = [\tau_i, \tau_{i+1}]$

$$\hat{B}_{2t}^i = W_{2t} + \int_{\tau_i}^t \theta_s^i ds, \tag{6}$$

is a standard Brownian motion under $V^i\mu_2$, where I posit an analogous inverse map $\Theta^i = I_W + \theta^i$. As such, $\nu^i = \left(Q^i \times V^i\right)\mu$. Although the transport from the base measure μ to the informative measure ν^i is in terms of the maps Q^i and V^i, agent i's optimal information acquisition policy will amount to choosing ϕ^i and θ^i.

Suppose that agent i observes the vector of signals, $[W_{1t}\ W_{2t}]'$, for $t = [\tau_i, \tau_{i+1}]$ that are white noise under μ, but under ν^i they are informative about the underlying fundamentals and the agent's task, $\boldsymbol{f}$ and s_i, respectively. The special case in which $q^i = v^i \equiv 0$ corresponds to observing white noise, in which case the agent learns nothing, and her posterior coincides with her prior.[9,10]

With a diffusion information structure and the technological restriction of independence of posteriors, I can appeal to Cameron Martin methods to parameterize an arbitrary change of

[8] It is the well-defined adjoint operator of the closure of the Gross-Sobolev derivative, ∇, in the H direction of random variables in $L^p(\mu)$-spaces, which is a closable operator. In particular, although ϕ need not be adapted, the Skorokhod integral coincides with the Itô integral for adapted elements ϕ in $\mathbb{D}_{p,k}(H)$.

[9] Although I characterize filtering in terms of an optimal change in measure, I can also express learning as a state-space system and derive the appropriate Zakai stochastic differential equation for the optimal filter based on the Kallianpur-Striebel principle (e.g., Liptser and Shiryaev, 2001).

[10] That agent i forms informative signals from the white noise $[W_{1t}\ W_{2t}]$ for $t \in [\tau_i, \tau_{i+1}]$ rather than for $t \in [0, \tau_{i+1}]$ is immaterial for her optimal information acquisition problem. What ultimately matters is the overall informativeness of her signals, $\mathbb{E}_\mu\left[\frac{1}{2}|\phi^i|_H^2 \mid \mathcal{W}_{\tau_i}\right]$, rather than how the informativeness is spread over time before she takes action at τ_{i+1}. However, this assumption will later simplify the analysis with AI because agents will not have overlapping signals.

measure with the measurability restrictions that separate learning about fundamentals and about task weights as[11]

$$\frac{dv^i}{d\mu} = \rho\left(-\delta\phi^i\right)\rho\left(-\delta\theta^i\right). \tag{8}$$

With this parameterization of the Radon-Nikodym derivative, the Kullback-Leibler divergence between her prior and posterior under $\mathbb{W}^i$ is given by

$$\mathbb{D}\left(\nu^i||\mu\right) = \frac{1}{2}\mathbb{E}_{\nu^i}\left[|\phi^i|_H^2 + |\theta^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] = \kappa_f^i + \kappa_s^i, \tag{9}$$

where $|\phi^i|_H^2$ and $|\theta^i|_H^2$ are evaluated from dates τ_i to τ_{i+1}, and κ_f^i and κ_s^i are the attention capacities agent i devotes to learning about the fundamentals and her task, respectively.

These Radon-Nikodym derivatives must also satisfy the Bayes' plausibility constraints that they are changes of measures, which imposes

$$\mathbb{E}_\mu\left[\rho\left(-\delta\phi^i\right) \mid \mathcal{W}_{\tau_i}\right] = 1, \tag{10}$$

$$\mathbb{E}_\mu\left[\rho\left(-\delta\theta^i\right) \mid \mathcal{W}_{\tau_i}\right] = 1. \tag{11}$$

This approach to learning represents the variational analogue to state-space filtering, and tractably admits more general, nonlinear learning dynamics parameterized through the $L^2(\mu, H)$ processes q^i and v^i as opposed to the Zakai / Kushner local characterizations that require infinite-dimensional state spaces. Intuitively, the agent is endowed with a series of white noise observations, W_1 and W_2, that compose the sigma-algebra of her information set, $\mathcal{W}_{\tau_i+1}$, and she can choose whether and how to make these signals informative about (i.e., to correlate with) the fundamentals and her task state through a change of measure over time.[12]

An agent born at time τ_i, internalizing her behavior at time $\tau_i + 1$ when she makes her capital

[11] I could also allow for agents to learn using a jump information structure, in which case the appropriate Girsanov exponential is

$$\zeta(\lambda) = \exp\left(\int_0^1 \left(\lambda_t \log\lambda_t - (\lambda_t - 1)\right) dt\right), \tag{7}$$

where λ_t is the intensity with which a jump occurs and information arrives. The Kullback-Leibler divergence for this information structure would then be $\mathbb{E}_\nu\left[\int_{\tau_i}^{\tau_i+1}\left(\lambda_t \log\lambda_t - \lambda_t + 1\right) dt \mid \mathcal{W}_{\tau_i}\right]$.

[12] Notice that a state-space filter about a hidden state, x, with an observation equation

$$dy_t = c_t x dt + dW_t, \tag{12}$$

admits a change of measure under which y_t is a Brownian motion, i.e., white noise. The variational approach to acquiring information starts with y_t as a Brownian motion and searches for a change of measure that yields a decomposition akin to Equation (12).

decision, consequently solves the constrained optimization problem

$$u^i = \sup_{\nu^i} \sup_{K^i} \mathbb{E}_{\nu^i} \left[U^i \mid \mathcal{W}_{\tau_i} \right], \tag{13}$$

$$s.t. : \text{Equations (2), (10), (11).}$$

1.1.3 Optimal Effort and Information Acquisition Policies

Making use of (stochastic) variational methods, I have the following proposition that characterizes the optimal policies of agent i. The optimal effort K^i is the agent's best estimate of her action under her posterior ν^i. The optimal change of measure adheres to the Boue-Dupuis lemma suitably modified for independence of the posteriors over the fundamentals and the state of the agent's task. Importantly, I derive a local characterization of the optimal filter by modifying the standard Föllmer drift (e.g., Üstünel, 2009) to account for the endogeneity of U^i to the optimal measure since it depends on the conditional expectation of $\boldsymbol{a}(s_i)' \boldsymbol{f}$ under ν^i through the optimal effort. This correction is the variational analogue of the Ocone and Karatzas (1991) modification to the Clark-Ocone-Haussmann martingale representation theorem under an endogenous change of measure.

As a result of the technological constraint that the posteriors are separable, the optimal filter over learning about the fundamentals will be the Kalman-Bucy filter, while that over the task state will be a controlled Wonham filter. Since I solve agent i's rational inattention problem as a stochastic control problem, the latter has an optimal gain that ensures that it hits a target posterior variance. The optimal information acquisition strategy will be a water-filling strategy in which agent i equates the projected posterior variances across both learning channels subject to her rational attention constraint.

Let $\bar{\boldsymbol{a}}^i = \sum_{s \in \mathcal{S}_i} \bar{\pi}(s) \boldsymbol{a}(s)$ be agent i's expected task weights and $\Sigma_a = Var_\mu \left(\boldsymbol{a}(s) \mid \mathcal{W}_{\tau_i} \right)$ be their unconditional variance. In addition, let λ_i be the Lagrange multiplier on the entropy constraint (2), which represents the shadow value of information, and $\mathcal{D}_{jt}$ and $\mathcal{D}^B_{jt}$ be the Malliavin derivative operators with respect to the Brownian motions, W^i_{jt} and $\hat{B}^i_{jt}$, respectively, for $j \in \{1, 2\}$, and $\circ$ the composition operation for maps. I then have Proposition 1.

Proposition 1. *Agent i:*

- *chooses optimal effort*

$$K^i = \mathbb{E}_{\nu^i} \left[\boldsymbol{a}(s_i)' \boldsymbol{f} \mid \mathcal{W}_{\tau_{i+1}} \right]; \tag{14}$$

- *chooses optimal adapted perturbations of identity, ϕ^i and θ^i, according to a modified Föllmer Drift*

$$\dot{\phi}_t^i = -\frac{1}{\lambda_i}\mathbb{E}_{\mu_1}\left[\left(\mathcal{D}_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - L\left(\phi^i\right)\mathcal{D}_{1t}^B\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]\right) \circ Q^i \mid \mathcal{B}_t^i\right] \circ \Phi^i, \tag{15}$$

$$\dot{\theta}_t^i = -\frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[\left(\mathcal{D}_{2t}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right] - L\left(\theta^i\right)\mathcal{D}_{2t}^B\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right]\right) \circ V^i \mid \mathcal{B}_t^i\right] \circ \Theta^i. \tag{16}$$

The optimal filter for learning about the fundamentals can be implemented with $\dot{\phi}_t^i = \alpha^i \bar{\boldsymbol{a}}^{i\prime}\left(\boldsymbol{f} - \hat{\boldsymbol{f}}_t^i\right)$, where $\alpha^i = \sqrt{\frac{1}{\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i+1}^i\bar{\boldsymbol{a}}^i} - \frac{1}{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}}$ and $\bar{\boldsymbol{a}}^{i\prime}\hat{\boldsymbol{f}}_t^i$ is the conditional expectation of her optimal effort fixing the task weights at their ex-ante expected value. This corresponds to a constant-gain Kalman-Bucy filter that induces a posterior over $\bar{\boldsymbol{a}}^{i\prime}\boldsymbol{f}$ at time $\tau_i + 1$ that is Gaussian $\mathcal{N}\left(\bar{\boldsymbol{a}}^{i\prime}\hat{\boldsymbol{f}}_{\tau_i+1}^i, \bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i+1}^i\bar{\boldsymbol{a}}^i\right)$.

The optimal filter for learning about the task state can be implemented with $\dot{\theta}_t^i = \delta_t^i\bar{\boldsymbol{f}}'\left(\boldsymbol{a}\left(s_i\right) - \hat{\boldsymbol{a}}_t^i\right)$, where δ_t^i satisfies Equation (A.61) *and $\bar{\boldsymbol{f}}'\hat{\boldsymbol{a}}_t^i$ is the conditional expectation of her optimal effort fixing the fundamentals at their ex-ante expected value. This corresponds to a controlled Wonham filter that induces a posterior over $\bar{\boldsymbol{f}}'\boldsymbol{a}\left(s_i\right)$ at time $\tau_i + 1$ with conditional mean $\bar{\boldsymbol{f}}'\hat{\boldsymbol{a}}_{\tau_i+1}^i$ and variance $\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}$.*

- *under a small-gain approximation for the Wonham filter, follows a water-filling strategy for optimal information acquisition:*
  - *if $e^{-2\kappa}\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i > \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$, she devotes all attention to learning about the fundamentals and sets $\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i+1}^i\bar{\boldsymbol{a}}^i = e^{-2\kappa}\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$;*
  - *if $(1-2\kappa)\,\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} > \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$, she devotes all attention to learning about her task state and sets $\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = (1-2\kappa)\,\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$;*
  - *otherwise, she learns about both the fundamentals and her task and sets $\lambda_i = \bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i+1}^i\bar{\boldsymbol{a}}^i = \bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}$, where λ_i solves*

$$\lambda_i = \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}\; W\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}e^{1-2\kappa}\right), \tag{17}$$

and $W(\cdot)$ is the Lambert W function;

- *achieves expected utility*

$$\begin{aligned}\mathbb{E}_{\nu^i}\left[U^i\right] \;=\; & -\frac{1}{2}\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i+1}^i\bar{\boldsymbol{a}}^i - \frac{1}{2}\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\bar{\boldsymbol{f}} - \frac{1}{2}Tr\left(\Sigma_a\Omega_{\tau_i+1}^i\right) \\ & -\frac{1}{2}Tr\left(\left(\Omega_0 - \Omega_{\tau_i+1}^i\right)Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right),\end{aligned} \tag{18}$$

and with the small gain approximation devotes capacity to learning about the fundamentals and tasks

$$\kappa_s^i = \frac{1}{2}\left(1 - \frac{\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}\right), \qquad \kappa_f^i = \frac{1}{2}\log\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}{\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i}\right). \tag{19}$$

Proposition 1 reveals that the pathwise optimization of the optimal filter under independence of posteriors separates learning into the two standard tractable filters: the Kalman-Bucy filter over fundamentals and the Wonham filter over task states. Importantly, although the optimal drift of the signals, ϕ^i and θ^i, in principle take into account how perturbing ϕ^i and θ^i change U^i itself by altering how conditional expectations are estimated (the second terms of Equations (15) and (16), respectively), this additional term vanishes because U^i is a quadratic objective and its derivative, the expectation error, always has zero conditional expectation under the informative measure ν^i. However, this term is why solving nonlinear filters is a fixed-point problem since the Malliavin derivative over the second term is with respect to the endogenous Brownian motion, $\hat{B}^i_{jt}$.

A key advantage of this continuous-time stochastic control framework is that agent i's optimal information acquisition policy and expected utility (assuming the gain δ^i_t is not too large under the Wonham filter) are highly tractable, and only require knowledge of the final posterior variances of each filter. For the remainder of the paper, I maintain this small gain approximation for tractability.

Agent i's optimal information acquisition policy minimizes the posterior variances of optimal effort when learning about the fundamental, fixing beliefs about the task state, $\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i$, and when learning about the task state, fixing beliefs about the fundamentals, $\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\bar{\boldsymbol{f}}$. This involves learning about a weighted average of the fundamentals in the latter case, with weights $\bar{\boldsymbol{a}}^{i\prime}$, and about a weighted average of the task state, with weights $\bar{\boldsymbol{f}}$, and is a water-filling strategy in that the agent only learns about both sources of uncertainty if there is not sufficient uncertainty about one to exhaust her attention capacity, κ.[13] Her expected utility decomposes into negative one-half the sum of these two conditional variances, and the two latter terms that reflect the interaction of uncertainty about both fundamentals and her task state.

Because no information is passed onto future generations, each agent solves their optimal effort and information acquisition problems starting from the same prior about the fundamentals. However, because each devotes capacity κ^i_s to learning about her task, new tasks arrive with probability $\zeta\left(\kappa^i_s\right)$ at each instant. As a consequence, their priors about the scope of their tasks evolves over time, and future agents are incentivized to learn about potentially different

[13]This water-filling result is similar to Miao et al. (2022) in which rational inattention forces an agent to prioritize learning about the most uncertain sources of her payoff.

fundamentals.

The intergenerational link across agents through task discovery represents an invention externality from agents learning about their task. Agent i does not internalize that the capacity she devotes to learning about her specific task contributes to the stock of tasks in society. While learning in typical models of incomplete information only enables agents to track a hidden state, my framework links the acquisition of information to the creation of new tasks and the relevance of new fundamentals. This externality will interact with knowledge accumulation once I introduce artificial intelligence, which will give rise to a second informational externality: agent i does not internalize that learning about fundamentals contributes to the future training data in the corpus of the artificial intelligence. In essence, there is a tension between learning how to perform her own task better, which contributes to the forward-looking creation of new tasks, and learning about how to perform any task better through learning about fundamentals, which contributes to the backward-looking accumulation of training data.

2 The Economy with Artificial Intelligence

In this section, I introduce artificial intelligence (AI) as an aggregation of past information in the economy. Since tasks are agent-specific, the value of the AI will be in compiling past information about the unobservable fundamentals, $\boldsymbol{f}$, and presenting it to agents through answers to a series of prompts. An agent that uses artificial intelligence pays a subscription fee to receive an updated prior about the fundamentals without having to devote as much of her limited attention to learning about them. This allows an agent to focus more attention on learning about her task state, which contributes to future invention through the discovery of new tasks. However, unlike the internet and other tools that provide economic actors with access to past public information, the AI collapses past information into a posterior and retains rather than transfers this knowledge to its users. In addition, it can act as a biased intermediary of information based on its design. In what follows, I first construct the AI and then investigate the incentives of agents to adopt it. Finally, I explore the implications of access to AI for invention and innovation in the economy.

Suppose that the AI at time t has access to all past information of agents in the economy.[14] This is equivalent to the AI having the same filtration as agents, $\mathcal{W}_t$, and observing all past ϕ^i for $\tau_i + 1 \leq t$, where ϕ^i is characterized in Proposition 1. Because agents live on disjoint time

[14]For a setting where only users who make active decisions contribute to training data, which shapes the AI's expertise, see for instance Chen et al. (2026).

intervals, combining their changes of measure into a change of measure for the AI, parameterized by ϕ^a, amounts to linear aggregation, i.e., $\dot{\phi}_t^a = \dot{\phi}_t^i$ on $t \in [\tau_i, \tau_i + 1)$. This induces an informative measure $\mu_1^a = Q^a \mu_1 = L(\phi^a)\mu_1$ that is akin to the maximum likelihood "scoring" prediction based on the raw data that is used in large language models (e.g. Radford et al., 2018).

AI differs from typical assumptions about public information sets beyond its lack of portability to agents along two dimensions that reflect considerations of its designer. First, is pruning of the training data. The designers can discard some of the training data, which I interpret as coarsening of $\mathcal{W}_t$ to the smaller filtration $\mathcal{E}_t$. I restrict attention to a continuous Blackwell garbling of the training data, which given Proposition 1 is equivalent to the AI observing noisier signals

$$d\tilde{B}_{1t}^i = \alpha^i \vec{a}^{i\prime} \left(f - \hat{f}_t^a \right) dt + \frac{1}{\beta} d\tilde{W}_{1t}, \tag{20}$$

for $\beta \in (0, 1]$, where $\tilde{W}_{1t}$ is a standard Brownian motion on $\mathcal{E}_t$.

Second, the designer can impose a post-process tilt of the AI's posterior, equivalent to the fine-tuning done by Reinforcement Learning from Human Feedback (RLHF) (e.g., Ouyang et al., 2022; Korbak et al., 2022).[15] This fine-tuning is done by biasing the AI according to some reward objective, $R(f)$. For instance, an AI trained on music data may be prompted about lithography, and $R(f)$ will determine whether the AI relays that it has little knowledge of lithography or if it acts confidently despite its lack of knowledge. I can incorporate this tilt by imposing a fixed penalty, $\lambda_T > 0$, on the mutual information between the garbled measure aggregated from past data, $\hat{\mu}_1^a$ and the final posterior that is passed onto agents that use the AI, μ_1^A. In this case, the AI's post-processing algorithm solves

$$\mathcal{R}_A = \sup_{\frac{d\mu_1^A}{d\hat{\mu}_1^a} \in L_+^1\left(\hat{\mu}_1^a, \mathcal{E}_0\right)} \mathbb{E}_{\hat{\mu}_1^a} \left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} R(f) - \lambda_T \mathbb{D}\left(\mu_1^A || \hat{\mu}_1^a \right) \mid \mathcal{E}_0 \right], \tag{21}$$

subject to $\mathbb{E}_{\hat{\mu}_1^a} \left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} \mid \mathcal{E}_0 \right] = 1$. In decision problem (21), λ_T has the interpretation of being the temperature parameter in modern LLMs. This formulation of the RLHF tilt as a Kullback-Leibler regularized change of measure is consistent with how Korbak et al. (2022) frame AI fine-tuning.

In practice, different AIs such as ChatGPT, Gemini, Claude, and Grok can differ in their answers to the same prompts, which can arise from differences in the training data and in the

[15] For instance, Ziegler et al. (2019); Ouyang et al. (2024) provide evidence that alignment tuning can impact the risk aversion that an AI displays in its recommendations. Bini et al. (2026) shows that LLMs display more behavioral biases when responding to preference-based versus belief-based tasks.

process by which the weights used for generating responses are calibrated. For instance, AIs may be programmed to be more or less risk averse to giving extreme responses to minimize the risk of providing wrong answers or overly precise recommendations, or may be biased toward specializing in specific tasks to cater toward them. They may also be programmed to intentionally ignore some of the training data to avoid copyright infringement or outliers influencing the algorithm.[16]

In what follows, I assume a quadratic reward objective, $R(\boldsymbol{f})$, for the AI

$$R(\boldsymbol{f}) = -\frac{1}{2}\left(\boldsymbol{f} - \bar{\boldsymbol{f}}\right)' \Gamma \left(\boldsymbol{f} - \bar{\boldsymbol{f}}\right), \tag{22}$$

for a $N \times N$ symmetric, positive semi-definite matrix Γ.[17] I center the reward function at the prior mean under μ_1, $\bar{\boldsymbol{f}}$, so that the AI has no expected directional bias; otherwise, agents would either trivially unravel it if it is known or need to form expectations about it. I defer the optimal choice of the temperature parameter, λ_T, and the reward function matrix, Γ, to the next section.[18]

Proposition 2 summarizes the AI's posterior.

Proposition 2. *The artificial intelligence concatenates all past agents' acquired information about the fundamentals into a change of measure $\mu_1^a = L(\phi^a)\,\mu_1$ such that for $t \in [\tau_i, \tau_i + 1)$*

$$\dot{\phi}_t^a = \dot{\phi}_t^i, \tag{23}$$

Blackwell garbles this measure into $\hat{\mu}_1^a = L\left(\hat{\phi}^a\right)\mu_1$ by restricting it to the coarser filtration $\mathcal{E}_t$ according to

$$\dot{\hat{\phi}}_t^a = \mathbb{E}_{Q^a\mu_1}\left[\dot{\phi}_t^a \mid \mathcal{E}_t\right], \tag{24}$$

[16] For instance, Zhu (2026) provides evidence that high-quality content providers opt-out of data sharing in large language model training data, creating an adverse selection problem. By contrast, I identify an advantageous selection phenomenon in that those who do not use AI, and who consequently devote the most attention to learning about fundamentals, contribute the most to its future training data.

[17] That Γ is symmetric is without loss given that it is the kernel in a quadratic form.

[18] In principle, agents could form expectations over $\frac{1}{\lambda_T}\Gamma$, for instance if I modeled the use of AI as a game against nature and employed a Perfect Bayesian Equilibrium refinement. However, the adversarial nature of AI is not the focus of this paper. In practice, users do not appear to attempt to correct for the bias in AI and demonstrate more trust than skepticism over longer interactions (e.g., Brynjolfsson et al., 2025).

and imposes the RLHF tilt into the posterior μ_1^A

$$\frac{d\mu_1^A}{d\hat{\mu}_1^a} = \frac{e^{-\frac{1}{2\lambda_T}(f-\bar{f})'\Gamma(f-\bar{f})}}{\mathbb{E}_{\hat{\mu}_1^a}\left[e^{-\frac{1}{2\lambda_T}(f-\bar{f})'\Gamma(f-\bar{f})} \mid \mathcal{E}_0\right]}. \tag{25}$$

Since all agents employ the Kalman-Bucy filter when learning about fundamentals, this reduces to a Gaussian posterior, $\mu_1^A \sim \mathcal{N}\left(\hat{\boldsymbol{f}}_t^A, \Omega_t^A\right)$ *with conditional mean*

$$\hat{\boldsymbol{f}}_t^A = \Omega_t^A \left(\Omega_0^{-1} + \sum_{j=0}^{i-1} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}\right) \hat{\boldsymbol{f}}_t^P + \frac{1}{\lambda_T}\Omega_t^A \Gamma \bar{\boldsymbol{f}}, \tag{26}$$

where $\hat{\boldsymbol{f}}_t^P$ *has law of motion Equation* (A.97)*, and conditional precision*

$$\left(\Omega_t^A\right)^{-1} = \underbrace{\Omega_0^{-1}}_{\text{Prior Precision}} + \underbrace{\sum_{j=0}^{i-1} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}}_{\text{Anisotropic Training Precision}} + \underbrace{\frac{1}{\lambda_T}\Gamma}_{\text{RLHF Tilt}}. \tag{27}$$

Agent i cannot recover Γ *from only observing* μ_1^A*.*

Proposition 2 reveals that the AI's final posterior that it distributes to agents who pay for a subscription is Gaussian, $\mu_1^A \sim \mathcal{N}\left(\hat{\boldsymbol{f}}_t^A, \Omega_t^A\right)$, with a conditional mean $\hat{\boldsymbol{f}}_t^A$ that is a weighted average between that from Bayesian learning, $\hat{\boldsymbol{f}}_t^P$, and the prior mean, $\bar{\boldsymbol{f}}$, where the weights are modulated by the AI's temperature parameter, λ_T. A lower temperature corresponds to a tighter posterior around the prior mean. The posterior precision is the sum of three components: 1) the prior, Ω_0^{-1}; 2) the precision garnered from the training data, $\sum_{j=0}^{i-1}\left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}$; and 3) the RLHF tilt, Γ. Importantly, the second component illustrates the backward-looking and anisotropic nature of aggregating past agents' acquired information. The AI's precision is limited to the linear combinations of fundamentals about which past agents have learned, i.e., $\bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}$, scaled by the strength of their signals, α^j.[19]

Having characterized the AI's posterior, I now consider the incentives of agents to purchase an AI subscription. I assume that agent i when born at time τ_i has the option to pay for an AI subscription, which gives the agent access to the AI's posterior at time τ_i. The subscription requires

[19]This anisotropic updating, in which the AI's precision increases by a rank-one matrix $\left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}$ after the latest agent's data is added to its corpus, is reminiscent of the Low-Rank Adaptation (LoRA) approach to fine-tuning pre-trained large language models (e.g., Hu et al., 2022). Rather than updating all weights, LoRA updates only the weights on the subspace in which the new data provides information.

an initial payment of $\frac{c}{2} > 0$ at time τ_i for the first half of her life, and then an additional payment of $\frac{c}{2}$ at time $\tau_i + \frac{1}{2}$ to maintain access until she completes her task at time $\tau_i + 1$. Because the agent is not the AI, the information remains with the AI and is not retained by the agent if she terminates her subscription at $\tau_i + \frac{1}{2}$. This distinguishes AI, which is an information intermediary, from other sources of public information, such as website articles and books. I assume that this resets her prior about the fundamental to that without AI, $\mathcal{N}\left(\bar{f}, \Omega_0\right)$, and forces the agent to re-optimize her information acquisition policy.[20]

Let m_1^i be the indicator of whether the agent uses the AI from τ_i to $\tau_i + \frac{1}{2}$, and m_2^i the indicator of whether she renews her subscription from $\tau_i + \frac{1}{2}$ to $\tau_i + 1$. Regardless of whether she chooses to use the AI, she solves the decision problem

$$u^i = \sup_{\nu^i, m_1^i, m_2^i} \sup_{K^i} \mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] - m_1^i\left(1 + m_2^i\right)\frac{c}{2}, \tag{28}$$

$$s.t. : \text{Equations (2), (10), (11),}$$

where I replace the prior μ_1 by μ_1^A if agent i buys an AI subscription. One can view the agent's use of AI as providing a prompt $\bar{\boldsymbol{a}}^i$ that the AI breaks down into loadings on the fundamentals through context, and then returns an answer, $\mathcal{N}\left(\bar{\boldsymbol{a}}^{i\prime}\hat{\boldsymbol{f}}_{\tau_i}^A, \bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A\bar{\boldsymbol{a}}^i\right)$. In what follows, I continue to make use of the small gain approximation for the Wonham filter from Proposition 1, and assume that agent i chooses her information acquisition policy with AI anticipating the higher precision of her updated prior but before observing $\hat{\boldsymbol{f}}_{\tau_i}^A$.

An important observation is that I am computing the agent's expected utility based on her incorrect perception that the AI's RLHF tilt $\frac{1}{\lambda_T}\Gamma$ conveys actual precision about the fundamentals. As a result, her realized utility can differ markedly (and suboptimally), and the AI tilt represents a form of persuasion. Since I impose a mean tilt of $\bar{f}$, the AI has no directional bias, which would be akin to the AI providing answers that nudge all agents toward mean belief, and the higher precision is observationally equivalent ex ante to the AI having more training data. Given that the Large Language Models (LLMs) underpinning AI are a relatively recent phenomenon, their designs opaque (e.g., the training weights), and the companies that maintain them are young and have yet to develop reputations, such an assumption is a reasonable starting point for highlighting the issues that arise with AI as a biased intermediary designed primarily for profit maximization

[20]This is in contrast to the "bounded rationality" paradigm in which agents have limited access to past private but full access to past public information. Here, it is the reverse.

rather than public good.

Proposition 3 summarizes agent i's optimal behavior when she can use AI. Let $\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq 0$ be the incremental value of using AI versus not for agent i, which is given by Equation (A.104). In addition, let $\boldsymbol{e}_n$ be the Euclidian $N \times 1$ basis vector with the n^{th} entry equal to 1 and the rest zeroes, and ι_N be the $N \times 1$ vector of ones.

Proposition 3. *Agent i purchases an AI subscription if $\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq c$, and will always renew her subscription at time $\tau_i + \frac{1}{2}$. An agent that uses AI devotes (weakly) more attention to learning about her task, i.e., κ_s^i (weakly) increases.*

An agent for whom $\bar{\boldsymbol{a}}^i = \boldsymbol{e}_n$, where n is the minimum diagonal entry of the matrix $\Omega_{\tau_i}^A - \Omega_0$, finds the AI most valuable ex ante. By contrast, an agent whose expected task weights are given by

$$\bar{\boldsymbol{a}}^i = \frac{1}{\iota_N \left(\Omega_0 - \Omega_{\tau_i}^A\right)^{-1} \iota_N}{}' \left(\Omega_0 - \Omega_{\tau_i}^A\right)^{-1} \iota_N, \tag{29}$$

– provided this vector has all nonnegative entries and $\Omega_0 - \Omega_{\tau_i}^A$ is nonsingular – finds the AI least valuable ex ante.[21]

Proposition 3 reveals that agent i adopts AI if her expected utility gain from using the AI exceeds its full subscription cost, i.e., $\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq c$, and if she does, that she will always renew her subscription at time $\tau_i + \frac{1}{2}$. Interestingly, this result suggests a form of "AI dependence". Because AI condenses its knowledge into responses to agent i's prompts, she does not retain this knowledge if she discontinues use of the AI midway through her life.[22]

In addition, the use of AI allows agent i to devote (weakly) more of her limited attention to learning about her task state, s_i, increasing κ_s^i, by reducing her attention to learning about the fundamentals. As is standard, access to public information (here embodied in the AI) acts as a substitute for private information. This is consistent with AI improving the productivity and quality of its users' work (e.g., Brynjolfsson et al., 2025; Peng et al., 2023).[23]

Finally, Proposition 3 provides guidance as to which agents ex ante benefit from using AI. An agent whose expected task, $\bar{\boldsymbol{a}}^i$, loads exclusively on the fundamental about which the AI provides the most reduction in uncertainty, $\Omega_{\tau_i}^A - \Omega_0$, ex ante finds AI the most useful. By contrast, an agent

[21] If the nonnegativity constraints on $\bar{\boldsymbol{a}}^i$ bind, the least-valuable average task loads on the corner of the simplex most orthogonal to the AI's precision.

[22] Although assuming full loss of the AI's knowledge is a stark assumption, the insight that an agent must rely on the AI for knowledge she cannot retain is a generic feature of AIs.

[23] Consistent with this insight, Lin (2025) provides evidence that as Midjourney's AI capabilities improved, users focused on refining their prompts to produce higher-quality output.

whose expected task, $\bar{\boldsymbol{a}}^i$, is most orthogonal to the AI's reduction in uncertainty ex ante finds AI the least useful. Intuitively, this is the direction that minimizes the variance $\bar{\boldsymbol{a}}^{i\prime}\left(\Omega_{\tau_i}^A - \Omega_0\right)\bar{\boldsymbol{a}}^i$, i.e., for which AI has the minimal reduction in uncertainty about her optimal effort.

Conversely, this last observation also reveals whose training data is most useful to the AI. The agent who is most likely to use AI is the agent whose expected task is already well-represented in the training data (through the second component of Equation (27)), and consequently whose data has the lowest marginal value. By contrast, the agent who is least likely to use AI is the agent whose expected task is poorly represented in the training and, because changes in the span of tasks $\mathcal{S}_t$ are persistent, whose data has the highest marginal value to the AI.

This tension between the AI being backward-looking in expertise while agents are forward-looking in their expected tasks can give rise to cyclicality of AI adoption. Once an agent uses AI, all future agents will also use AI until there is a jump in the task space because they face identical decision problems and the AI becomes more knowledgeable about the fundamentals relevant to the current span of tasks over time. This use of AI shifts more attention to learning about tasks, i.e., κ_s^i increases, which increases the likelihood of the invention of new tasks and antiquation of old tasks through a higher $\zeta\left(\kappa_s^i\right)$. Since the AI may not be sufficiently informative about these new tasks after a jump to justify the subscription cost, agents may stop using AI until new training data has improved its performance. This reduces agents' attention to learning about tasks and consequently the likelihood of another jump until agents begin adopting the AI again. When this occurs, the AI increases the volatility of pace of discovery through these exploration and exploitation cycles compared to the economy without AI.

Consistent with this view that AI has accelerated the discovery of new and antiquation of old tasks, Loaiza and Rigobón (2024) find significant churn in occupation-related tasks, specifically toward human-intensive work, since 2015, which covers the time-period when generative AI models were commercially introduced. Hanson and Escobar (2026) document in a large employer survey that AI adoption is shifting the skills required and expected responsibilities of entry-level hires toward more analytical and judgment-based than routine tasks. This change in tasks after the introduction of AI, and the selection of which workers for whom AI adds the most value, has implications for the evaluation of how its adoption at firms impacts worker wages Gans (2026).

I provide sufficient conditions for this cyclicality of AI adoption and the invention of new tasks and obsolescence of old tasks that acts through the value of the AI's training data, and is summarized in Proposition 4. In what follows, I assume that c is not too large that there is AI

adoption in the economy, and that N and $|\mathcal{S}|$ are sufficiently large that there are many tasks and fundamentals on which the tasks can load.

Proposition 4. *An increase in innovation from AI usage over time increases the arrival rate of the invention of new tasks and obsolescence of old tasks. If after the arrival, $\bar{\boldsymbol{a}}^i$ is in the null space of $\Omega_t^A - \Omega_0$, then future agents do not adopt AI until the AI has acquired sufficient training data on the new tasks.*

Proposition 4 also reveals that the benefits of AI can drop abruptly when the economic environment changes substantially, here after a jump that drastically changes the span of tasks in the economy. This highlights a potential shortcoming of having economic agents rely heavily on AI. Although it improves decision-making when the economic environment is stable, i.e., the AI is used for applications that are very similar to those used in the AI's training data, its utility can diminish when uncertainty is the highest and it is most needed.

These AI adoption cycles that I identify also have implications for the concern that the use of AI will lead to a collapse in human knowledge, whether through learning traps that arise from the substitution of private effort for endogenous public information (e.g., Acemoglu et al., 2026a) or through the loss of training and workplace interactions (e.g., Ide, 2025). In my setting, AI acts as a natural substitute for learning about time-invariant fundamentals, which frees up the ability of forward-looking agents to use AI to innovate (discover their task) or invent if their attention to their tasks leads to invention. By contrast, AI is backward-looking and incapable of the jumps in the span of tasks that correspond to invention in the economy. When the economic environment changes and new tasks arrive, for instance during the COVID-19 pandemic, the recent and persistent rise inflation, or the resurgence of geopolitics, then the usefulness of AI diminishes and agents must rely on their own ability to acquire knowledge. Consequently, my framework provides a more optimistic view of the impact of AI on human ingenuity.

I conclude this section by remarking on how the AI's RLHF tilt $\frac{1}{\lambda_T}\Gamma$ impacts how agents who use AI acquire information. Because she cannot recover the tilt from one observation of the AI's posterior, μ_1^A, when she purchases an AI subscription, she interprets the AI's inflated precision in the directions of the fundamentals implied by the Γ matrix as reflecting knowledge garnered from training data rather than manufactured confidence from a reward function. This induces her to acquire less information about the fundamentals. Since a lower temperature, λ_T, is equivalent to magnifying each element of Γ, a lower temperature also induces agent i to acquire less information about the fundamentals. These observations are summarized in Proposition 5.

Proposition 5. *If $\Gamma_1 \succeq \Gamma_2$ (in a positive semi-definite sense), then agent i acquires less information about fundamentals when using AI with RLHF tilt Γ_1 than under tilt Γ_2. Agent i's attention to the fundamentals is also increasing in the AI's temperature, λ_T.*

This last insight has important implications for the long-term cost of AI "hallucinations", and further distinguishes AI as a biased intermediary compared to a passive synthesizer of public information. Because economic agents cannot disentangle the informativeness of an AI's answers from confidence stemming from the AI's reward function, they overestimate the precision of AI responses and overly substitute it for their own information acquisition activities. This reduces the acquisition of new information, i.e., training data, and erodes the quality of their decision-making.

3 Extensions

In this section, I consider two extensions of my model. In the first, I consider the implications of the dynamic feedback between AI adoption, attention allocation, and task discovery for the optimal design of AI by a developer. I then modify an agent's payoff to include an evaluative component by AI and show that this can give rise to dynamics that are reminiscent of a Keynesian beauty contest and can lead to a collapse in human knowledge accumulation.

3.1 Optimal Design of AI

In this section, I extend my model to examine its implications for the privately and socially optimal design of artificial intelligence (AI). Consider a three-period skeleton of my continuous-time framework in which time is indexed by $t \in \{0, 1, 2\}$. Each agent in the economy lives from time t to $t + 1$ and continues to solve decision problem (28), whose solution is characterized in Propositions 1 and 3. In addition, the AI's posterior continues to be characterized by Proposition 2. I assume that the developer is endowed with an updated prior over the fundamentals $\boldsymbol{f}$, $\hat{\mu}_1^a \sim \mathcal{N}\left(\hat{\boldsymbol{f}}_0^P, \Omega_0^P\right)$ that reflects data on which the AI has trained before $t = 0$.

I now formulate the problem of the developer who designs the AI. Suppose that there is a long-lived developer whose goal is to maximize the lifetime present-discounted value of the profits at $t = 0$ and $t = 1$ from AI subscriptions.[24] At time 0, he chooses the time-invariant parameters for

[24]The AI developer's design problem relates to a growing literature on the economics of data sales. Admati and Pfleiderer (1990) consider the incentives for a monopolistic data seller to sell versus trade directly on its private information. Yang (2022) shows how the revenue-maximizing mechanism for data sales induces quasi-perfect price discrimination, while Bergemann and Bonatti (2024) analyzes the rent extraction and efficiency gains of targeted ads

the AI: the degree of data degradation, β, and the RLHF tilt, Γ, in which I impose Γ is positive semi-definite and $Tr(\Gamma) \leq 1$ so that the optimal choice is well-defined.[25] By contrast, and departing from the baseline model, I assume he can vary the subscription cost, c_t, and the AI's temperature parameter, $\lambda_{Tt} \geq \underline{\lambda}_T > 0$ (to avoid division by zero), at $t = 0$ and $t = 1$, i.e., across successive agents. This reflects that although data curation and RLHF tilt are effectively fixed for an AI model in practice, since retraining occurs infrequently, subscription costs and the model's temperature can be adjusted over time.

To operate the AI, the developer must pay for the electricity and water costs associated with processing the AI prompts, i.e., $\bar{\boldsymbol{a}}$, of subscribing agents.[26] Such costs are typically ascribed in terms of token usage, and I assume that the cost increases over time as the AI becomes more precise from accumulating training data. Specifically, I assume that the cost of running the AI scales linearly with its precision, which is equivalent to the size of its corpus. An agent prompt at time t costs the developer $\chi_t = \frac{C}{2}\iota_N' \left(\left(\Omega_t^P\right)^{-1} - \Omega_0^{-1} \right) \iota_N$ for $C > 0$, where $\left(\Omega_t^P\right)^{-1}$ is the precision of the AI without the RLHF tilt, i.e., $\left(\Omega_t^P\right)^{-1} = \left(\Omega_0^P\right)^{-1} + \sum_{s=0}^{t-1} \left(\beta\alpha^s\right)^2 \bar{\boldsymbol{a}}^s \bar{\boldsymbol{a}}^{s\prime}$ and ι_N is the $N \times 1$ vector of ones. As such, substituting for α^i with Proposition 1 and recognizing $\bar{\boldsymbol{a}}^{s\prime}\iota_N \equiv 1$,

$$\chi_t = \frac{C}{2}\iota_N' \left(\left(\Omega_t^P\right)^{-1} - \Omega_0^{-1} \right) \iota_N = \chi_0 + \frac{C}{2}\beta^2 \sum_{s=0}^{t-1} \left(\frac{1}{\bar{\boldsymbol{a}}^{s\prime}\Omega_{s+1}^s \bar{\boldsymbol{a}}^s} - \frac{1}{\bar{\boldsymbol{a}}^{s\prime}\Omega_s^s \bar{\boldsymbol{a}}^s} \right), \tag{30}$$

where $\Omega_s^s = \Omega_0$ if agent s did not adopt AI and $\Omega_s^s = \Omega_s^A$ if she did to account for the impact of AI usage on her prior.[27]

I assume that the developer lacks commitment and discounts his profits at subjective discount rate, ρ, which can be viewed as $1/R$, where R is the constant gross risk-free rate. He consequently solves the problem

$$\Pi_0 = \sup_{\{c_t, \lambda_{Tt}, \Gamma, \beta\}} \mathbb{E}_{\mu^D} \left[(c_0 - \chi_0)\, \mathbf{1}_{\{m_1^0 = 1\}} + \rho\, (c_1 - \chi_1)\, \mathbf{1}_{\{m_1^1 = 1\}} \mid \mathcal{W}_0 \right], \tag{31}$$

where $\mu^D = \hat{\mu}_1^a \times \mu_2 \times \mu^J$, where μ^J is the Poisson jump measure for the arrival of new tasks at

on two-sided platforms and Galperti et al. (2024) the value to an intermediary of coarsening data records. Liu et al. (2025) examines the dynamic incentives associated with data sales versus subscriptions for multiple buyers when data is durable and sales are self-diluting. Here, I consider the sale of a (biased) posterior rather than the underlying data whose sale feeds back to its value to future buyers.

[25] Given that λ_T multiplies Γ in the AI's posterior mean and variance, bounding the trace by unity is a normalization.

[26] In practice, AI companies maintain data centers that consume vast amounts of electricity (e.g., Darmouni et al., 2026) and water (e.g., Li et al., 2025).

[27] By contrast, Agrawal et al. (2018a) considers the problem of an AI provider who cannot price discriminate and sells prediction capabilities to agents who have heterogeneous levels of experience based on their past exercises of judgment.

date 1. Since the agent at $t = 0$ has her task fixed even if new tasks arrive between $t = 0$ and $t = 1$, only the agent at $t = 1$ draws her task from the new task span. The lack of commitment imposes that I search for a subgame-perfect equilibrium.

An important issue from which I am abstracting, as alluded to in Section 2, is the adversarial game between the developer and agents. By manipulating the RLHF tilt, the developer can create an AI that has the illusion of precision. This is reflected in the deviation from the agent's expected and realized utilities. Given that each agent interacts with the AI only once, she does not have enough data on the ex post performance of the AI to recover $\frac{1}{\lambda_T}\Gamma$; however, now that I microfound the developer's problem, she could rationally anticipate its design. In this sense, the developer's problem is akin to Bayesian persuasion with a rationally inattentive receiver (e.g., Bloedel and Segal, 2018), and without a directional bias, the AI's posterior will always satisfy a Bayes' plausibility constraint (i.e., the posterior always integrates back to the prior). A full treatment would require an information design problem (e.g., Bergemann and Morris, 2016, 2019) with appropriate incentive compatibility constraints for the agent to buy a subscription, but the incentives of the developer and the impact of persuasion would be the same.

I begin with the developer's problem at $t = 1$, which captures the static trade-offs that he faces in subscription pricing and adoption. Proposition 6 characterizes the developer's optimal pricing and temperature policies. In what follows, Δ corresponds to the difference between the value agent 1 obtains with the AI and without.

Proposition 6. *At $t = 1$, the developer sets a discriminatory price for the AI subscription, $c_1 = \Delta \mathbb{E}_{\nu^1}\left[U^1 \mid \mathcal{W}_1\right]$, a temperature $\lambda_{T1} = \underline{\lambda}_T$, and sells the subscription provided that $c_1 \geq \chi_1$.*

At $t = 1$, the developer sets a subscription price, c_1, to extract all surplus created by the AI from agent 1. In addition, he sets the temperature arbitrarily low because this inflates the value of the AI to the agent, and maximizes the price he can extract from her. This is optimal provided that $c_1 \geq \chi_1$, i.e., the subscription price exceeds the developer's cost of running the AI; otherwise, he does not sell a subscription to her.

I now turn to the developer's problem at $t = 0$. By contrast to the static case at $t = 1$, the developer at $t = 0$ must take into account how his policies that target agent 0 impact the quality of training data and the cost of computing at $t = 1$, and the likelihood of a jump in the task span between $t = 0$ and $t = 1$. In what follows, I assume that the marginal cost of computing, C, is an intermediate range so that the optimal rate of data degradation, β, has an interior solution. Proposition 7 characterizes the developer's optimal policies at $t = 0$.

Proposition 7. *At $t = 0$, the developer sets:*

1. *a discriminatory price for the AI subscription, $c_0 = \Delta\mathbb{E}_{\nu^0}\left[U^0 \mid \mathcal{W}_0\right]$;*

2. *a temperature λ_{T0} that satisfies Equation* (A.124) *that trades off the increase in initial subscription price c_0 with the increased likelihood of a jump in tasks at $t = 1$ from agent 0's increased attention to learning about her task;*

3. *a rate of data degradation β that satisfies Equation* (A.122) *that trades off the marginal reduction in the expected subscription price at $t = 1$, $\mathbb{E}_{\mu^D}\left[c_1 \mid \mathcal{W}_0\right]$, with the marginal reduction in expected computing costs, $\mathbb{E}_{\mu^D}\left[\chi_1 \mid \mathcal{W}_0\right]$;*

4. *a RLHF tilt Γ that satisfies Equation* (A.131) *and inflates AI precision in directions, i.e., eigenvectors, that maximally increase the developer's expected profit at $t = 0$, Π_0;*

5. *sells a subscription at $t = 0$ if*

$$c_0 - \chi_0 \geq -\rho\Delta\mathbb{E}_{\mu^D}\left[c_1 - \chi_1 \mid \mathcal{W}_0\right], \tag{32}$$

i.e., the net profit he receives from agent 0 exceeds the expected discounted decline in net profit he receives from agent 1;

Similar to the case of $t = 1$, the developer optimally price discriminates against agent 0 and charges her full reservation price, $\Delta\mathbb{E}_{\nu^0}\left[U^0 \mid \mathcal{W}_0\right]$, if he sells her a subscription. By contrast, his choice of temperature now trades off the increase in subscription value, c_1, from inflating the AI's precision with the increased likelihood of a jump in the task span from the agent allocating more attention to learning about her task. This is costly for the developer because a jump in the task span renders the AI (weakly) less valuable to agent 1, which lowers c_1, because her set of potential tasks may load very differently on the fundamentals from those that the AI trained on using agent 0's data. The optimal rate of data degradation, β, balances the marginal benefit of a higher c_1 from more informative training data with the marginal cost of more expensive computing, χ_1. Finally, the optimal RLHF tilt Γ loads on directions that maximize the developer's expected profits. This balances inflating the AI's precision in directions that agent 0 finds valuable (and agent 1 if the task span remains unchanged) with insuring its perceived performance for agent 1 if the task span changes. The developer sells a subscription to agent 0 if the net profit from the sale, $c_0 - \chi_0$, exceeds the expected discounted degradation in profit he garners from agent 1.

Proposition 7 reveals that the developer internalizes how inflating the precision of the AI and selling it to agent 0 not only lowers the quality of her training data and reduces the cost of computing at $t = 1$, it raises the likelihood of a jump in the task span by shifting her attention toward learning about her task. These forward-looking considerations impose endogenous limits on his ability to persuade both agents to use the AI by manipulating its RLHF tilt and temperature. Consequently, there is a tension between static rent extraction through subscription costs and dynamic investment in the AI's future quality through preservation of the quality and relevance of its training data.

An additional important insight from Propositions 6 and 7 is the role of computing costs in shaping AI adoption. My model highlights that this cost can endogenously increase over time as the resources required to analyze vast amounts of data and process tokens grows. If costs become prohibitively high, i.e., χ_1 exceeds $\Delta\mathbb{E}_{\nu^1}\left[U^1 \mid \mathcal{W}_1\right]$ for a vast majority of task spans, then the developer shutters the AI at $t = 1$. Consequently, my framework predicts that when cost considerations are taken into account, subscription fees will naturally rise and rationing of high-quality computing power for AI to more valuable queries can occur.[28]

I conclude this section by discussing the constrained socially optimal design of AI in my economy. Suppose instead of a developer, a planner with the same tools as the developer chooses whether agents 0 and 1 adopt the AI to maximize their expected utilities net of the costs of computing

$$U^P = \mathbb{E}_{\mu^D}\left[U^0 - \chi_0 \mathbf{1}_{\{m_1^0=1\}} + \rho\left(U^1 - \chi_1 \mathbf{1}_{\{m_1^1=1\}}\right) \mid \mathcal{W}_0\right], \tag{33}$$

where the χ_0 terms reflect concerns over expending resources to support AI rather than human consumption. Since the planner aims to maximize the expected utility of agents and the subscription fees represent pure transfers, it chooses zero RLHF tilt, i.e., $\Gamma = 0$ (and consequently infinite temperature). This is optimal since distorting agents' decision-making hurts their expected utility and inducing them to use AI when they should not also consumes resources through computing power. Consequently, the planner induces (weakly) less adoption that the developer. Beyond this key difference, the optimal policies for data degradation, β, and whether an agent should use AI have the same structure as Propositions 6 and 7 (but evaluated at the AI's true posterior). Consequently, the planner in this economy is akin to a benevolent developer.

[28] For instance, agent AI models now require 5-30x more tokens per task than standard GenAI chatbots, and expected token demand for agentic services is projected to outpace supply by 2030 (Sommer, 2026). Korinek and McKelvey (2026) estimates that although token prices have fall about 97% per year since 2023, nominal compute spending and physical compute output have grown roughly 144% and 213% per year in 2024 and 2025, respectively. This is projected to raise overall inference costs even if the price of tokens falls substantially from technological improvements.

However, this parity breaks if the planner wants to foster innovation in the economy. Suppose that the planner derives utility from the invention of new tasks, for instance because new tasks create economic opportunities through creative destruction or the planner cares about the level of effort, $\boldsymbol{a}'\boldsymbol{f}$. Such a planner would face a nontrivial trade-off between fostering AI to promote learning about tasks, and increasing the likelihood of new tasks being discovered, with the degradation of the AI's value once new tasks have arrived to which the AI is poorly trained. In such a setting, the planner internalizes the paradox that AI as a backward-looking tool fosters creation that is inherently forward-looking.

Finally, one may ask whether competition among AI providers would induce developers to avoid manipulating agents using the RLHF tilt. However, as shown in Heidhues and Kőszegi (2017), competition can worsen exploitation when agents do not fully rationally, and Ichihashi and Kim (2023) show that digital platforms can respond by making content more addictive. Consequently, my model posits a role for regulation of how AI providers disclose details of how they fine-tune their models.

3.2 AI Beauty Contests and the Collapse of Knowledge

In this section, I explore the implications of my model for when humans not only use artificial intelligence (AI) as an aid in information acquisition, but are also evaluated by it. For instance, a firm may set a worker's wage, in part, based on an AI's assessment of her output because of asymmetric information or moral hazard issues. Alternatively, customers may choose their demand for products based on AI recommendations, which determines an entrepreneur's profits.[29] This will give rise to a novel externality akin to a Keynesian beauty contest by which an agent distorts her decision-making to match the AI's perception of her optimal effort rather targeting the level of effort that is actually optimal. I interpret this as a fundamental difference between humans and AI: humans have intrinsic drive to learn and work for personal fulfillment, while AI are designed to score well based on a performance reward function.

Suppose that agent i born at time τ_i derives utility at time $\tau_i + 1$ from a linear combination of the quadratic loss between her level of effort, K^i and its optimal value, $\boldsymbol{a}\left(s_i\right)'\boldsymbol{f}$, with weight $1-w$, and between K^i and the AI's estimate of its optimal value, $\boldsymbol{a}\left(s_i\right)'\mathbb{E}_{\mu_1^A}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i+1}\right]$, with weight w

$$U^i = -\frac{1-w}{2}\left(\boldsymbol{a}\left(s_i\right)'\boldsymbol{f} - K^i\right)^2 - \frac{w}{2}\left(\boldsymbol{a}\left(s_i\right)'\mathbb{E}_{\mu_1^A}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i+1}\right] - K^i\right)^2. \tag{34}$$

[29] In the presence of behavioral biases, McLaughlin and Spiess (2024) show that agent preferences can depend directly on alignment with an AI's recommendation.

The agent takes the AI's filter and its process of acquiring training data as given. As such, she does not internalize that the information that she acquires about the fundamentals from time τ_i to $\tau_i + 1$ informs the AI's optimal prediction at time $\tau_i + 1$. For simplicity, I assume the AI's temperature, $\lambda_T \to \infty$, so that the AI is unbiased; otherwise, I would have to specify agent i's beliefs about this bias. In addition, I assume $\beta = 1$ so that there is no knowledge degradation for the AI.

Proposition 8 reveals that not only do agents devote less attention to learning about fundamentals when $w > 0$, but it is possible when w is sufficiently large for society to acquire no new knowledge about the fundamentals, i.e., there is a knowledge collapse as in Acemoglu et al. (2026a).

Proposition 8. *If $w > 0$, then agent i optimally acquires less information about the fundamentals. There exists a w^* such that If $w > w^*$, then no agent adopts AI and all agents devote zero attention to learning about the fundamentals. In this case, human knowledge accumulation collapses.*

Intuitively, placing weight in an agent's utility on the AI's evaluation of her effort, the second component in Equation (34), reduces the attention she devotes to learning about the fundamentals because it reduces the contribution of the fundamentals to the variance of her payoff. This is because by the law of total variance, the variance of the AI's conditional expectation of the fundamentals is lower than their unconditional variance. For w sufficiently large, such that agent i's utility places sufficient weight on the AI's evaluation, then she devotes zero attention to learning about the fundamentals. Because the AI receives zero training data, its posterior is frozen at $\mathcal{N}\left(\bar{f}, \Omega_0\right)$, and no agent pays the subscription fee $c > 0$ to adopt. In the case, knowledge collapse is the unique self-fulfilling equilibrium and social learning ceases.

This extension offers a nuanced perspective on the risk that AI will lead to a social collapse in the production of new knowledge. It is not the usage of AI for assistance that causes learning to cease in my model, but rather the use of AI for judgment and evaluation. It is not a coincidence that the formulation of agent i's utility in Equation (34) resembles the preferences in Morris and Shin (2002), where putting weight on the actions of others gives rise to an externality that crowds out private information. Relying on AIs for evaluation is akin to a Keynesian beauty contest that crowds out private information because it implicitly coordinates agents' actions with the AI whose information is not only endogenous and but, importantly, derived from the past and path-dependent. Consequently, my framework cautions not against the use of AI, in itself, but specifically in the use of AI for judgment and evaluation. If humans stop valuing the human component of their work, then knowledge and the creativity that stems from it can atrophy.

4 Conclusion

I develop a model of rational inattention in which agents devote limited attention to learning about economic fundamentals and their specific task that uses these fundamentals, and attention to learning about tasks contributes to the discovery of new ones. I derive closed-form optimal filters - a Kalman-Bucy filter for fundamentals and a controlled Wonham filter for tasks - and a water-filling strategy for how to allocate attention between the two channels.

I then incorporate AI, which provides agents' with a posterior based on past public information that is biased by a Reinforcement Learning with Human Feedback (RLHF) tilt. I show that the use of AI can lead to a form of "AI dependence" in which agents who use AI cannot stop using it because they do not retain the information the AI summarizes. By freeing agents to learn more about their tasks, AI adoption can lead to invention and innovation waves in which news tasks are discovered more quickly, which reduces the usefulness of and reliance on AI until it acquires new training data, which then hastens the discovery of new tasks. Such waves suggest an optimistic view of the role of AI in human learning as this self-correcting mechanism prevents a social collapse in knowledge accumulation (e.g., Acemoglu et al., 2026a; Ide, 2025). However, such a collapse can occur not from agents using AI but from being evaluated by it, akin to how the over-reliance on public information crowds out the provision of private information in Keynesian beauty contests (e.g., Allen et al., 2006). Finally, I consider the optimal design of AI and show how a developer trades off maximizing subscription fees using the RLHF tilt with its impact on future training data, computing costs, and the likelihood of new task discovery.

The role of AI in democratizing expertise using past information raises several interesting issues. As computing costs rise, the model suggests either AI may need to be rationed to answering the most (socially) valuable questions, or its capacity limited for more egalitarian access. In addition, it posits a role for policy in promoting transparency in AI's post-processing so that consumers can better discern its expertise from its reward tilt. Finally, fostering AI development differs from other forms of industrial policy in that it is developed using talent and resources locally but diffuses skill globally, and may paradoxically erode the global competitiveness of the countries and institutions that develop it.

References

Acemoglu, Daron and Pascual Restrepo, "Artificial Intelligence, Automation, and Work," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019, pp. 197–236.

_ **and** _ , "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **, Dingwen Kong, and Asuman Ozdaglar**, "AI, Human Cognition and Knowledge Collapse," Working Paper 34910, National Bureau of Economic Research 2026.

_ **, Tianyi Lin, Asuman Ozdaglar, and James Siderius**, "How AI Aggregation Affects Knowledge," 2026.

Admati, Anat R. and Paul Pfleiderer, "Direct and Indirect Sale of Information," *Econometrica*, 1990, *58* (4), 901–928.

Afrouzi, Hassan, Andres Blanco, Andrés Drenik, and Erik Hurst, "Automation, Learning, and Career Dynamics," 2026.

Agarwal, Nikhil, Alex Moehring, and Alexander Wolitzky, "Designing Human-AI Collaboration: A Sufficient-Statistic Approach," 2025. Working Paper.

Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

Agrawal, Ajay K., Joshua S. Gans, and Avi Goldfarb, "Human Judgment and AI Pricing," *AEA Papers and Proceedings*, 2018, *108*, 58–63.

_ **,** _ **,** **and** _ , "Prediction, Judgment, and Complexity: A Theory of Decision-Making and Artificial Intelligence," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2018, pp. 89–110.

_ **,** _ **, and** _ , "The Economics of Bicycles for the Mind," Working Paper 34034, National Bureau of Economic Research 2025.

Allen, Franklin, Stephen Morris, and Hyun Song Shin, "Beauty Contests and Iterated Expectations in Asset Markets," *The Review of Financial Studies*, 2006, *19* (3), 719–752.

Allinger, Daniel F. and Sanjoy K. Mitter, "New Results on the Innovations Problem for Non-Linear Filtering," *Stochastics*, 1981, *4* (4), 339–348.

Athey, Susan C., Kevin A. Bryan, and Joshua S. Gans, "The Allocation of Decision Authority to Human and Artificial Intelligence," *AEA Papers and Proceedings*, 2020, pp. 80–84.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940–2018," *The Quarterly Journal of Economics*, 2024, *139*, 1399–1465.

Bergemann, Dirk and Alessandro Bonatti, "Data, Competition, and Digital Platforms," *American Economic Review*, 2024, *114* (8), 2553–2595.

_ **and Stephen Morris**, "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 2016, *11* (2), 487–522.

_ **and** _ , "Information Design: A Unified Perspective," *Journal of Economic Literature*, 2019, *57* (1), 44–95.

Bini, Pietro, Lin William Cong, Xing Huang, and Lawrence J. Jin, "Behavioral Economics of AI: LLM Biases and Corrections," 2026. Working Paper.

Blackwell, David, "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 1953, *24* (2), 265–272.

Bloedel, Alexander W. and Ilya R. Segal, "Persuasion with Rational Inattention," 2018. Working Paper.

Bonney, Kathryn, Cory L. Breaux, Emin Dinlersoz, Lucia S. Foster, John C. Haltiwanger, and Aditya A. Pande, "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks," 2025. Available at `https://www.nber.org/papers/w35141`.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, "Generative AI at Work," *The Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Caplin, Andrew, Mark Dean, and John Leahy, "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy," *Journal of Political Economy*, 2022, *130* (6), 1676–1715.

Chemla, Gilles and Vincent Tena, "AI Adoption, Incentives, and Firm Value in General Equilibrium," 2026. Working Paper.

Chen, Joanne Juan, Brandon Yueyang Han, and Hongda Zhong, "Human Skills in the Age of AI," 2026. Working Paper, London School of Economics.

Darmouni, Olivier, Clemens Lehner, and Yuqi Zhang, "The Energy Cost of AI and Data Centers," 2026. Working Paper.

Dreyfuss, Bnaya and Ruru Hoong, "Calibrated Coarsening: Designing Information for AI-Assisted Decisions," 2025. Working Paper, Available at SSRN: `https://ssrn.com/abstract=5286198`.

Duffie, Darrell and David Lando, "Term Structures of Credit Spreads with Incomplete Accounting Information," *Econometrica*, 2001, *69* (3), 633–664.

Farboodi, Maryam, Andrew Koh, and Anchi Xia, "Data-Driven Automation," 2025. Working Paper, MIT.

Galperti, Simone, Aleksandr Levkun, and Jacopo Perego, "The Value of Data Records," *The Review of Economic Studies*, 2024, *91* (2), 1007–1038.

Gans, Joshua, "Endogenous Task Bundling, Skills and Automation," 2026. Working Paper.

Hansen, Lars Peter, Jianjun Miao, and Hao Xing, "Robust Inattentive Discrete Choice," *Proceedings of the National Academy of Sciences*, 2025, *122*, e2416643122.

Hanson, Andrew and Molly Cook Escobar, "Entry-Level Hiring in the AI Era: What Employers Are Thinking (and Doing)," Technical Report, Strada Institute for the Future of Work May 2026. Available at `https://www.strada.org/news-insights/entry-level-hiring-in-the-ai-era-what-employers-are-thinking-and-doing`.

Hébert, Benjamin and Michael Woodford, "Rational Inattention when Decisions Take Time," *Journal of Economic Theory*, 2023, *208*, 105612.

Heidhues, Paul and Botond Kőszegi, "Naïveté-Based Discrimination," *The Quarterly Journal of Economics*, 2017, *132* (2), 1019–1054.

Hu, Edward J., Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen, "LoRA: Low-Rank Adaptation of Large Language Models," in "International Conference on Learning Representations" 2022.

Huang, Jing and Shumiao Ouyang, "Soft Information, Hard Decisions: AI Advising," 2025. Working Paper.

Ichihashi, Shota and Byung-Cheol Kim, "Addictive Platforms," *Management Science*, 2023, *69* (3), 1677–1696.

Ide, Enrique, "Automation, AI, and the Intergenerational Transmission of Knowledge," 2025. IESE Business School Working Paper. Reject & Resubmit, AER.

_ and Eduard Talamás, "Artificial Intelligence in the Knowledge Economy," *Journal of Political Economy*, 2025, *133* (12), 3762–3800.

Korbak, Tomasz, Ethan Perez, and Christopher L. Buckley, "RL with KL Penalties is Better Viewed as Bayesian Inference," in "Findings of the Association for Computational Linguistics: EMNLP 2022" Association for Computational Linguistics 2022, pp. 1083–1091.

Korinek, Anton and Patrick McKelvey, "Measuring the AI Economy," Working Paper, Peterson Institute for International Economics 2026.

Li, Pengfei, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren, "Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models," *Communications of the ACM*, 2025, *68* (7), 54–61.

Lin, Sijie, "Learning to Prompt: Human Adaptation in Production with Generative AI," 2025. Working Paper.

Liptser, Robert S. and Albert N. Shiryaev, *Statistics of Random Processes I: General Theory*, 2 ed., Vol. 5 of *Applications of Mathematics*, Berlin, Heidelberg: Springer, 2001.

Liu, Ernest, Song Ma, and Laura Veldkamp, "Data Sales and Data Dilution," *Journal of Financial Economics*, 2025, *169*, 104053.

Loaiza, Isabella and Roberto Rigobón, "The EPOCH of AI: Human-Machine Complementarities at Work," 2024.

Luo, Yulei, "Consumption Dynamics under Information Processing Constraints," *Review of Economic Dynamics*, 2008, *11* (2), 366–385.

Maćkowiak, Bartosz and Mirko Wiederholt, "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 2009, *99* (3), 769–803.

_ , Filip Matějka, and Mirko Wiederholt, "Dynamic Rational Inattention: Analytical Results," *Journal of Economic Theory*, 2018, *176*, 650–692.

_ , _ , and _ , "Rational Inattention: A Review," *Journal of Economic Literature*, 2023, *61* (1), 226–273.

McLaughlin, Bryce and Jann Spiess, "Algorithmic Assistance with Recommendation-Dependent Preferences," *arXiv preprint arXiv:2208.07626*, 2024. Extended abstract published in Proceedings of the 24th ACM Conference on Economics and Computation (EC'23).

Miao, Jianjun, Jieran Wu, and Eric R. Young, "Multivariate Rational Inattention," *Econometrica*, 2022, *90* (2), 907–945.

Morris, Stephen and Hyun Song Shin, "Social Value of Public Information," *American Economic Review*, 2002, *92* (5), 1521–1534.

Moscarini, Giuseppe, "Limited Information Capacity as a Source of Inertia," *Journal of Economic Dynamics and Control*, 2004, *28* (10), 2003–2035.

Ocone, Daniel L. and Ioannis Karatzas, "A Generalized Clark Representation Formula, with Application to Optimal Portfolios," *Stochastics and Stochastics Reports*, 1991, *34* (3–4), 187–220.

Ouyang, Long, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe, "Training Language Models to Follow Instructions with Human Feedback," in "Advances in Neural Information Processing Systems," Vol. 35 2022, pp. 27730–27744.

Ouyang, Shumiao, Hayong Yun, and Xingjian Zheng, "AI as Decision-Maker: Ethics and Risk Preferences of LLMs," 2024. Working Paper.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," *arXiv preprint arXiv:2302.06590*, 2023.

Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever, "Improving Language Understanding by Generative Pre-Training," 2018. OpenAI Technical Report.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal, "AI Models Collapse When Trained on Recursively Generated Data," *Nature*, 2024, *631*, 755–759.

Sims, Christopher A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

_ , "Rational Inattention and Monetary Economics," in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, Vol. 3, Elsevier, 2010, chapter 4, pp. 155–181.

Sommer, Will, "Navigating the Commoditization Trap as Token Costs Fall by Over 90% Through 2030," 2026.

Üstünel, Ali Süleyman, "Entropy, Invertibility and Variational Calculus of the Adapted Shifts on Wiener Space," *Journal of Functional Analysis*, 2009, *257*, 3655–3689.

Yang, Kai Hao, "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," *American Economic Review*, 2022, *112* (4), 1364–1393.

Zhong, Hongda, "Optimal Integration: Human, Machine, and Generative AI," *Management Science*, 2025.

Zhong, Weijie, "Optimal Dynamic Information Acquisition," *Econometrica*, 2022, *90*, 1537–1582.

Zhu, Kai, "Adverse Selection in the AI Data Commons," 2026. Working Paper.

Ziegler, Daniel M., Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, and Paul Christiano, "Fine-Tuning Language Models from Human Preferences," in "arXiv preprint arXiv:1909.08593" 2019.

Appendix: Proof of Propositions

Proof of Proposition 1

Based on the measurability constraints I impose that learning about the fundamentals is independent from about the task state, I decompose the Radon-Nikdoymn derviatve, $\frac{dv^i}{d\mu}$, as $\frac{dv^i}{d\mu} = \rho\left(-\delta\phi^i\right)\rho\left(-\delta\theta^i\right)$. which again $\rho\left(-\delta\theta^i\right)$ is a Girsanov exponential, which assuming Novikov's or Kamizaki's condition are honest martingales that will have unitary expectation. Let λ_i be the Lagrange multiplier on the rational inattention constraint and η_f^i and η_s^i the Lagrange multipliers on the constraints that the changes of measure have unitary expectation, respectively. I first express decision problem (28) as the Lagrangian under a change in measure

$$\begin{aligned}\mathcal{L}_i = \sup_{\phi^i,\theta^i\in H}\sup_{K^i\in\mathbb{R}}\inf_{\lambda_i,\eta_f^i,\eta_s^i\in\mathbb{R}}\mathbb{E}_\mu\left[L\left(\phi^i\right)L\left(\theta^i\right)U^i-\lambda_i\left(L\left(\phi^i\right)L\left(\theta^i\right)\left(\frac{1}{2}|\phi^i|_H^2+\frac{1}{2}|\theta^i|_H^2\right)-\kappa\right)\ |\mathcal{W}_{\tau_i}\right]\\ -\eta_f^i\left(\mathbb{E}_\mu\left[L\left(\phi^i\right)\ |\mathcal{W}_{\tau_i}\right]-1\right)-\eta_s^i\left(\mathbb{E}_\mu\left[L\left(\theta^i\right)\ |\mathcal{W}_{\tau_i}\right]-1\right).\end{aligned} \tag{A.1}$$

Notice that the expectation is taken with respect to μ, under which $\boldsymbol{f}$ and $\boldsymbol{a}(s)$ are independent because the marginals μ_f and $\mathbf{\Pi}_0$ are independent. In addition, $K^i \in \mathcal{W}_{\tau_i+1}$. I apply the Donsker-Varadhan lemma in which I impose that the admissible set of changes of measure are elements of the Cameron-Martin space, i.e., satisfy Girsanov's Theorem.

Step 1: Optimal Effort K^i:

As is standard, the optimal choice of effort at time $\tau_i + 1$ from decision problem (A.1) is

$$K^i = \mathbb{E}_\nu\left[\boldsymbol{a}(s_i)'\boldsymbol{f}\mid\mathcal{W}_{\tau_i+1}\right], \tag{A.2}$$

which by the imposed posterior independence imposes

$$K^i = \mathbb{E}_{V^i\mu_2}\left[\boldsymbol{a}(s_i)\ \mid\mathcal{W}_{\tau_i+1}\right]'\mathbb{E}_{Q^i\mu_1}\left[\boldsymbol{f}\mid\mathcal{W}_{\tau_i+1}\right]. \tag{A.3}$$

The optimal choice of effort is the agent's best estimates of the task state, s_i, and the fundamentals, $\boldsymbol{f}$, which are independent of each other.

Substituting the optimal choice of effort from Equation (A.3) into agent i's objective, I find

$$U^i = -\frac{1}{2}\left(\boldsymbol{a}\left(s_i\right)'\boldsymbol{f} - \mathbb{E}_{V^i\mu_2}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right]'\mathbb{E}_{Q^i\mu_1}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i+1}\right]\right)^2, \tag{A.4}$$

Step 2: Optimal Posterior v^i:

The first-order condition for the optimal change of measure ϕ^i for a variation $\hat{\phi})^i_{\xi_q} = \phi^i + \xi_q\epsilon^i_{1t}$ from Lagrangian (A.1) is given by the Gateaux derivative

$$\left(L\left(\theta^i\right)U^i - \lambda_i L\left(\theta^i\right)\log L\left(\phi^i\right) - \lambda_i - \eta^i_f\right)L\left(\phi^i\right)\epsilon^i_{1t} = 0, \tag{A.5}$$

from which follows, taking expectations with respect to μ_2 that

$$L\left(\phi^i\right) = e^{\frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[L\left(\theta^i\right)U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - 1 - \frac{1}{\lambda_i}\eta^i_f} = e^{\frac{1}{\lambda_i}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - 1 - \frac{1}{\lambda_i}\eta^i_f}. \tag{A.6}$$

Notice that I do not optimize over K^i by the Envelope Theorem. Imposing that $\mathbb{E}_{\mu_1}\left[L\left(\phi^i\right) \mid \mathcal{W}_{\tau_i}\right] = 1$,Equation (A.6) implies

$$e^{1+\frac{1}{\lambda_i}\eta^i_f} = \mathbb{E}_{\mu_1}\left[e^{\frac{1}{\lambda_i}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]} \mid \mathcal{W}_{\tau_i}\right], \tag{A.7}$$

from which follows by substituting Equation (A.7) into Equation (A.6) that

$$L\left(\phi^i\right) = \rho\left(-\delta\phi^i\right) = \frac{e^{\frac{1}{\lambda_i}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]}}{\mathbb{E}_{\mu_1}\left[e^{\frac{1}{\lambda_i}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]} \mid \mathcal{W}_{\tau_i}\right]}, \tag{A.8}$$

and the optimal change of measure θ^i satisfies

$$L\left(\theta^i\right) = \rho\left(-\delta\theta^i\right) = \frac{e^{\frac{1}{\lambda_i}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right]}}{\mathbb{E}_{\mu_2}\left[e^{\frac{1}{\lambda_i}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right]} \mid \mathcal{W}_{\tau_i}\right]}. \tag{A.9}$$

Notice that $e^{\frac{1}{\lambda_i}\mathbb{E}_{Q^i\mu_1}\left[U^i\right]}$ and $e^{\frac{1}{\lambda_i}\mathbb{E}_{V^i\mu_2}\left[U^i\right]}$ are independent because the former is a function of W_1 and the latter of W_2, and I have imposed independence in posteriors. Consequently, the changes of measure under the optimal information acquisition program are well-defined.

Step 3: Optimal Transport Processes:

Consider now variations of q^i and v^i, i.e., $\tilde{q}^i = q^i + \xi_q \dot{\epsilon}_{1t}$ and $\tilde{v}^i = v^i + \xi_v \dot{\epsilon}_{2t}$ for scalars ξ_q and ξ_v and arbitrary perturbations $\epsilon_1, \epsilon_2 \in H$. Importantly, q^i and v^i will be adapted to the larger base filtration, $\mathcal{B}_t^i$, with respect to which s_i and $\boldsymbol{f}$ are also measurable.

Building on Üstünel (2009), I recognize because U^i is Frechet-differentiable that for arbitrarily small perturbations (parameterized by $\xi_q$and ξ_v)

$$\begin{aligned}\mathbb{E}_\mu\left[L\left(\tilde{\phi}^i\right)L\left(\tilde{\theta}^i\right)U^i\,|\mathcal{W}_{\tau_i}\right] &= \mathbb{E}_\mu\left[U^i\circ Q^i\circ V^i\,|\mathcal{W}_{\tau_i}\right]+\mathbb{E}_\mu\left[\xi_f\nabla_1 U^i\circ Q^i\circ V^i+\xi_s\nabla_2 U^i\circ Q^i\circ V^i\,|\mathcal{W}_{\tau_i}\right]\\ &\quad +O\left(\xi_q^2,\xi_v^2\right).\end{aligned} \tag{A.10}$$

I can also express the mutual information between μ and ν under these perturbations of q^i and v^i

$$\begin{aligned}\mathbb{E}_\mu\left[\frac{1}{2}L\left(\phi_i\right)L\left(\theta_i\right)\int_{\tau_i}^{\tau_i+1}\left(|\dot{\phi}_t^i|^2+|\dot{\theta}_t^i|^2\right)dt\,|\mathcal{W}_{\tau_i}\right] &= \mathbb{E}_\mu\left[\frac{1}{2}\left(|\phi^i\circ Q^i|_H^2+|\theta^i\circ V^i|_H^2\right)\,|\mathcal{W}_{\tau_i}\right]\\ &\quad +\mathbb{E}_\mu\left[\xi_q\langle\phi^i\circ Q^i,\epsilon_1\rangle_H+\xi_v\langle\theta^i\circ V^i,\epsilon_2\rangle_H\,|\mathcal{W}_{\tau_i}\right].\end{aligned} \tag{A.11}$$

Substituting with Equations (A.10) and (A.11), I can express the problem of the optimal change of measure as

$$\mathcal{L}_i = \sup_{K^i\in\mathbb{R}}\sup_{\phi^i,\theta^i\in H}\mathbb{E}_\mu\left[\xi_q\langle\nabla_1 U^i\circ Q^i\circ V^i+\lambda_i\phi^i\circ Q^i,\epsilon_1\rangle_H+\xi_v\langle\nabla_2 U^i\circ Q^i\circ V^i+\lambda_i\theta^i\circ V^i,\epsilon_2\rangle_H\,|\mathcal{W}_{\tau_i}\right] \tag{A.12}$$

$$\begin{aligned}&+\mathbb{E}_\mu\left[\frac{1}{2}\left(|\phi^i\circ Q^i|_H^2+|\theta^i\circ V^i|_H^2\right)dt\,|\mathcal{W}_{\tau_i}\right]-\eta_f^i\left(\mathbb{E}_\mu\left[\rho\left(-\delta\tilde{\phi}^i\right)\,|\mathcal{W}_{\tau_i}\right]-1\right)\\ &-\eta_s^i\left(\mathbb{E}_\mu\left[\rho\left(-\delta\tilde{\theta}^i\right)\,|\mathcal{W}_{\tau_i}\right]-1\right)+\mathbb{E}_\mu\left[U^i\circ Q^i\circ V^i\,|\mathcal{W}_{\tau_i}\right]+\lambda_i\kappa.\end{aligned}$$

From Equation (A.12), the optimal change of measure q^i satisfies

$$\mathbb{E}_\mu\left[\langle p\left(\nabla_1 U^i\circ Q^i\circ V^i+\lambda_i\phi^i\circ Q^i\right),\epsilon_1\rangle_H\,|\mathcal{W}_{\tau_i}\right]-\eta_f^i\mathbb{E}_\mu\left[\rho\left(-\delta\phi^i\right)\langle W^1+q^i,\epsilon_1\rangle\,|\mathcal{W}_{\tau_i}\right]=0, \tag{A.13}$$

where $p\left(\cdot\right)$ is the dual predictable projection and I recall that $\phi^i = -q^i\circ\Phi^i$. However, it is

immediate that

$$\mathbb{E}_{\mu}\left[\rho\left(-\delta\phi^{i}\right)\left\langle W_{1}+q_{t}^{i},\epsilon_{1}\right\rangle \mid\mathcal{W}_{\tau_{i}}\right]=\mathbb{E}_{Q^{i}\mu_{1}}\left[\left\langle B_{1}^{i},\epsilon_{1}\right\rangle \mid\mathcal{W}_{\tau_{i}}\right]=\mathbb{E}_{Q^{i}\mu_{1}}\left[\delta\left(\epsilon_{1}\right)\mid\mathcal{W}_{\tau_{i}}\right]=0. \tag{A.14}$$

Consequently, substituting Equation (A.14) into (A.13), and recognizing that $\phi^{i}\in\mathcal{B}_{t}^{i}$, the first-order condition reduces under the law of iterated expectations to

$$\mathbb{E}_{\mu}\left[\mathbb{E}_{\mu}\left[\left\langle p\left(\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}+\lambda_{i}\phi^{i}\circ Q^{i}\right),\epsilon_{1}\right\rangle_{H}\mid\mathcal{B}_{\tau_{i}}^{i}\right]\mid\mathcal{W}_{\tau_{i}}\right]=0. \tag{A.15}$$

Importantly, compared to the standard Föllmer drift that operates on a fixed Brownian functional, the Föllmer drift here must account for how perturbing the path of the Brownian motion also changes the functional itself through how it changes the filter via $L\left(\phi^{i}\right)$. In the case of the quadratic objective, it enters through the Kallianpur-Striebel formula for $\hat{h}_{t}^{i}=\frac{\mathbb{E}_{\mu_{1}}\left[L\left(\phi^{i}\right)h\mid\mathcal{W}_{t}\right]}{\mathbb{E}_{\mu_{1}}\left[L(\phi^{i})\mid\mathcal{W}_{t}\right]}$. This is similar to the correction that arises in the correction to the Clark-Ocone-Hausmann martingale representation theorem under a Girsanov change of measure (e.g., Ocone and Karatzas, 1991).

To take this into account, I recognize that $\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}$ will have two terms

$$\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}=\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}|_{L(\phi^{i})}+\nabla_{\phi^{i}}U^{i}\circ Q^{i}\circ V^{i}|_{W_{1}} \tag{A.16}$$

Assuming the Gateaux derivatives, $\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}$ and $\nabla_{2}U^{i}\circ Q^{i}\circ V^{i}$, known as Gross-Sobolev derivatives admit Frechet representations that are absolutely continuous with respect to the Lesbegue measure, I can express Equation (A.10) as Malliavin derivatives. By the Riesz representation $\left\langle\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}|_{L(\phi^{i})},\epsilon_{1}\right\rangle_{H}=\int_{\tau_{i}}^{\tau_{i}+1}\mathcal{D}_{1t}U^{i}\circ Q^{i}\circ V^{i}\dot{\epsilon}_{1t}dt$ and D_{1} corresponds to W_{1t}. The second term that captures the impact on the filter is given by applying Malliavin integration-by-parts

$$\begin{aligned}
\mathbb{E}_{\mu_{1}}\left[\left\langle\nabla_{1}U^{i}\circ Q^{i}\circ V^{i}|_{W_{1}},\epsilon_{1}\right\rangle_{H}\mid\mathcal{W}_{\tau_{i}}\right] &= -\mathbb{E}_{\mu_{1}}\left[L\left(\phi^{i}\right)\frac{\partial U^{i}}{\partial L\left(\phi^{i}\right)}\circ Q^{i}\circ V^{i}\int_{\tau_{i}}^{\tau_{i}+1}\dot{\epsilon}_{1t}\left(dW_{1t}+\dot{\phi}_{t}^{i}dt\right)\mid\mathcal{W}_{\tau_{i}}\right]\\
&= -\mathbb{E}_{\nu_{1}}\left[\frac{\partial U^{i}}{\partial L\left(\phi^{i}\right)}\circ Q^{i}\circ V^{i}\int_{\tau_{i}}^{\tau_{i}+1}\dot{\epsilon}_{1t}dB_{1t}^{i}\mid\mathcal{W}_{\tau_{i}}\right]\\
&= -\mathbb{E}_{\nu_{1}}\left[\int_{\tau_{i}}^{\tau_{i}+1}\mathcal{D}_{1t}^{B}\left(\frac{\partial U^{i}}{\partial L\left(\phi^{i}\right)}\right)\circ Q^{i}\circ V^{i}\dot{\epsilon}_{1t}dt\mid\mathcal{W}_{\tau_{i}}\right],
\end{aligned} \tag{A.17}$$

where $\mathcal{D}^B_{1t}$ is the Malliavin derivative with respect to B^i_{1t}.

Because relation (A.15) must hold for any arbitrary variation ϵ_{1t}, this imposes point-wise

$$\dot{\phi}^i_t \circ Q^i = -\frac{1}{\lambda_i}\mathbb{E}_{\mu_1}\left[\left(\mathcal{D}_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - L\left(\phi^i\right)\mathcal{D}^B_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]\right) \circ Q^i \mid \mathcal{B}^i_t\right], \quad \text{(A.18)}$$

$dt \times d\mu$ a.s. since Q^i and $\mathcal{D}_{1t}$ are independent of μ_2 and the map V^i.

Recognizing that $q^i = -\phi^i \circ Q^i$, Equation (A.18) becomes

$$\dot{q}^i_t = \frac{1}{\lambda_i}\mathbb{E}_{\mu_1}\left[\left(\mathcal{D}_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - L\left(\phi^i\right)\mathcal{D}^B_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]\right) \circ Q^i \mid \mathcal{B}^i_t\right], \quad \text{(A.19)}$$

and I have from $\phi = -q^i \circ \Phi^i$ that

$$\dot{\phi}^i_t = -\frac{1}{\lambda_i}\mathbb{E}_{\mu_1}\left[\left(\mathcal{D}_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right] - L\left(\phi^i\right)\mathcal{D}^B_{1t}\mathbb{E}_{V^i\mu_2}\left[U^i \mid \mathcal{W}_{\tau_i}, \boldsymbol{f}\right]\right) \circ Q^i \mid \mathcal{B}^i_t\right] \circ \Phi^i. \quad \text{(A.20)}$$

Similarly, the optimal change of measure v^i and θ^i satisfy

$$\dot{v}^i_t = \frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[\left(\mathcal{D}_{2t}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right] - L\left(\theta^i\right)\mathcal{D}^B_{2t}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right]\right) \circ V^i \mid \mathcal{B}^i_t\right], \quad \text{(A.21)}$$

$$\dot{\theta}^i_t = -\frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[\left(\mathcal{D}_{2t}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right] - L\left(\theta^i\right)\mathcal{D}^B_{2t}\mathbb{E}_{Q^i\mu_1}\left[U^i \mid \mathcal{W}_{\tau_i}, s_i\right]\right) \circ V^i \mid \mathcal{B}^i_t\right] \circ \Theta^i. \quad \text{(A.22)}$$

Notice on the larger filtration $\mathcal{B}^i_t$, on which $\boldsymbol{f}$ and s_i are measurable, that $Q^i \circ \Phi^i = I_W$ and $V^i \circ \Theta^i = I_W$ trivially, i.e., the adapted perturbation of identities are invertible *a.s.*. When I next solve for representations of the optimal filters for Q^i and V^i, I will have to verify that the maps are invertible on the coarser observation filtration, i.e., that $\mathcal{Q}^i = \sigma\left(\{\hat{B}^i_{1t}\}_{t\in[\tau_i,\tau_i+1]}\right)$ satisfies $\mathcal{Q}^i = \mathcal{W}_1$ for $t \in [\tau_i, \tau_i + 1]$, and similarly for $\mathcal{V}^i = \sigma\left(\{\hat{B}^i_{2t}\}_{t\in[\tau_i,\tau_i+1]}\right)$ and $\mathcal{W}_2$.

Step 4: Conjecture a Linear Solution for q^i

Recall $\bar{\boldsymbol{f}} = \mathbb{E}_{\mu_1}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i}\right]$ and define $\bar{\boldsymbol{a}}^i = \mathbb{E}_{\mu_2}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i}\right] = \mathbb{E}_{V^i\mu_2}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i}\right]$ and $h^i = \bar{\boldsymbol{a}}^{i\prime}\boldsymbol{f}$. Writing agent i's objective U^i as

$$U^i = -\frac{1}{2}\left(h^i - \hat{h}^i_{\tau_i+1} + \left(\boldsymbol{a}\left(s_i\right) - \bar{\boldsymbol{a}}^i\right)'\boldsymbol{f}\right)^2, \quad \text{(A.23)}$$

I can then recover q^i from Equation (A.19) according to the Malliavin chain rule

$$\dot{q}_t^i = \frac{1}{\lambda_i}\mathbb{E}_{\mu_1}\left[\left(h^i - \hat{h}^i_{\tau_i+1}\circ Q^i\right)\left(\left(\mathcal{D}_{1t}\hat{h}^i_{\tau_i+1}\right) - \mathbb{E}_\nu\left[\mathcal{D}^B_{1t}\left(h^i - \hat{h}^i_{\tau_i+1}\right) \mid \mathcal{W}_t\right]\right)\circ Q^i \mid \mathcal{B}^i_t\right]. \tag{A.24}$$

Since $\mathbb{E}_\nu\left[e^i_t \mid \mathcal{W}_t\right] = 0$, where $e^i_{\tau_i+1} = h^i - \hat{h}^i_{\tau_i+1}$, it is a martingale difference sequence (has constant expectation 0) and by the Clark-Ocone-Hausmann martingale representation theorem,

$$\begin{aligned} e^i_{\tau_i+1} &= \mathbb{E}_\nu\left[e^i_{\tau_i+1} \mid \mathcal{W}_{\tau_i}\right] + \int_{\tau_i}^{\tau_i+1}\mathbb{E}_\nu\left[\mathcal{D}^B_{1t}\left(h^i - \hat{h}^i_{\tau_i+1}\right) \mid \mathcal{W}_t\right]dB^i_t \\ &= \int_{\tau_i}^{\tau_i+1}\mathbb{E}_\nu\left[\mathcal{D}^B_{1t}\left(h^i - \hat{h}^i_{\tau_i+1}\right) \mid \mathcal{W}_t\right]dB^i_t. \end{aligned} \tag{A.25}$$

Since $\mathbb{E}_\nu\left[e^i_{\tau_i+1} \mid \mathcal{W}_{\tau_i+1}\right] = 0$, it must be the case that $\mathbb{E}_\nu\left[\mathcal{D}^B_{1t}\left(h^i - \hat{h}^i_{\tau_i+1}\right) \mid \mathcal{W}_t\right] \equiv 0$. Consequently, the second term in Equation (A.24) vanishes.

Notice that, without loss of generality, I can express the conditional expectation $\hat{h}^i_t$, which will be a martingale under the optimal filter, as

$$d\hat{h}^i_t = \beta^i_t\left(dW_{1t} - \mathbb{E}_\nu\left[\dot{q}^i_t \mid \mathcal{W}_t\right]dt\right). \tag{A.26}$$

I now conjecture that $\dot{q}^i_t = \alpha^i_t\left(h^i - \hat{h}^i_t\circ Q^i\right)$, where α^i_t is deterministic, from which follows that $\mathbb{E}_\nu\left[\dot{q}^i_t \mid \mathcal{W}_t\right] = 0$ and from Equation (A.29) that

$$\hat{h}^i_{\tau_i+1} = \bar{h}^i + \int_{\tau_i}^{\tau_{i+1}}\beta^i_t dW_{1t}. \tag{A.27}$$

Conjecturing that β^i_t is also deterministic, it follows from Equation (A.50)

$$\mathcal{D}_{1t}\hat{h}^i_{\tau_i+1}\circ Q^i = \beta^i_t. \tag{A.28}$$

By contrast, I recognize that

$$d\hat{h}^i_t\circ Q^i = \beta^i_t\left(dW_{1t} + \dot{q}^i_t dt\right) = \beta^i_t\left(dW_{1t} + \alpha^i_t\left(h^i - \hat{h}^i_t\circ Q^i\right)dt\right), \tag{A.29}$$

which follows from Equation (A.29), since $d\hat{h}_t^i = d\left(\hat{h}_t^i - \bar{h}^i\right)$,

$$\hat{h}_t^i \circ Q^i = h^i + e^{-\int_{\tau_i}^t \beta_s^i \alpha_s^i ds}\left(\bar{h}^i - h^i\right) + \int_{\tau_i}^t e^{-\int_s^t \beta_s^i \alpha_s^i ds} \beta_s^i dW_{1s}, \tag{A.30}$$

and consequently

$$\mathbb{E}_{\mu_1}\left[\hat{h}_{\tau_i+1}^i \circ Q^i \mid \mathcal{B}_t^i\right] = h^i + e^{-\int_t^{\tau_i+1} \beta_s^i \alpha_s^i ds}\left(\hat{h}_t^i \circ Q^i - h^i\right). \tag{A.31}$$

Substituting Equations (A.51) and (A.31) into (A.24), I have that

$$\dot{q}_t^i = \frac{e^{\int_t^{\tau_i+1} \beta_s^i \alpha_s^i ds} \beta_t^i}{\lambda_i} \mathbb{E}_{\mu_1}\left[h^i - \hat{h}_{\tau_i+1}^i \circ Q^i \mid \mathcal{B}_t^i\right] = \frac{e^{-\int_t^{\tau_i+1} \beta_s^i \alpha_s^i ds} \beta_t^i}{\lambda_i}\left(h^i - \hat{h}_t^i \circ Q^i\right), \tag{A.32}$$

because $h^i \in \mathcal{B}_t^i$. Intuitively, the change of measure converts from μ, under which the conditional expectation of h^i is $\bar{h}^i \circ Q^i$, to ν, under which the conditional expectation of h^i is $\hat{h}_t^i$.

Matching Equation (A.32) with the conjecture

$$\dot{q}_t^i = \alpha_t^i\left(h^i - \hat{h}_t^i \circ Q^i\right) = \frac{e^{-\int_t^{\tau_i+1} \beta_s^i \alpha_s^i ds} \beta_t^i}{\lambda_i}\left(h^i - \hat{h}_t^i \circ Q^i\right), \tag{A.33}$$

from which follows that

$$\alpha_t^i = \frac{e^{-\int_t^{\tau_i+1} \beta_s^i \alpha_s^i ds} \beta_t^i}{\lambda_i}. \tag{A.34}$$

This functional form for $\dot{q}_t^i$ implies a signal structure $h^i - \hat{h}_t^i \circ Q^i$

$$dW_{1t} = \dot{q}_t^i dt + dB_{1t}^i = \alpha_t^i\left(h^i - \hat{h}_t^i \circ Q^i\right) dt + dB_{1t}^i. \tag{A.35}$$

Using the Kallianpur-Striebel formula (e.g., Liptser and Shiryaev, 2001), the conditional distribution of h^i admits a likelihood ratio representation, from which the law of motion for the conditional mean, $d\hat{h}_t^i$ follows

$$d\hat{h}_t^i = Cov_\nu\left[dW_{1t}, h^i \mid \mathcal{W}_t\right] dW_{1t} = \alpha_t^i \Sigma_t^i dW_{1t}, \tag{A.36}$$

where again Σ_t^i is the conditional variance of the posterior for h^i. It is then immediate from

comparing Equation (A.36) to (A.29) that

$$\beta_t^i = \alpha_t^i \Sigma_t^i, \tag{A.37}$$

and substituting for β_t^i in Equation (A.34) that

$$1 = \frac{e^{-\int_t^{\tau_{i+1}} \left(\alpha_s^i\right)^2 \Sigma_s^i ds} \Sigma_t^i}{\lambda_i}. \tag{A.38}$$

Evaluating Equation (A.38) at time $\tau_i + 1$ reveals that

$$\Sigma_{\tau_{i+1}}^i = \lambda_i, \tag{A.39}$$

and substituting Equation (A.39) back into Equation (A.38) that

$$\int_t^{\tau_{i+1}} \left(\alpha_s^i\right)^2 \Sigma_s^i ds = \log\left(\frac{\Sigma_t^i}{\Sigma_{\tau_{i+1}}^i}\right). \tag{A.40}$$

Notice that the law of motion for Σ_t^i is given by

$$d\Sigma_t^i = -Cov_\nu\left[dW_{1t}, h^i \mid \mathcal{W}_t\right]^2 dt - Cov_\nu\left[\left(h^i - \hat{h}_t^i\right)^2, dW_{1t} \mid \mathcal{W}_t\right] dW_{1t}. \tag{A.41}$$

Conjecturing a Gaussian posterior, in which higher-order moments than the second moment are zero, I arrive at the standard law of motion of Σ_t^i for the Kalman-Bucy filter

$$d\Sigma_t^i = -\left(\alpha_t^i \Sigma_t^i\right)^2 dt \;\rightarrow\; \Sigma_t^i = \frac{\Sigma_{\tau_i}^i}{1 + \Sigma_{\tau_i}^i \int_{\tau_i}^t \left(\alpha_s^i\right)^2 ds}. \tag{A.42}$$

As such, Σ_t^i falls monotonically over time. Equation (A.42) implies

$$\int_t^{\tau_{i+1}} \left(\alpha_s^i\right)^2 \Sigma_s^i ds = \frac{\int_t^{\tau_{i+1}} \Sigma_{\tau_i}^i \left(\alpha_s^i\right)^2 ds}{1 + \Sigma_{\tau_i}^i \int_{\tau_i}^s \left(\alpha_r^i\right)^2 dr} = \log\left(\frac{1 + \Sigma_{\tau_i}^i \int_{\tau_i}^{\tau_i+1} \left(\alpha_s^i\right)^2 ds}{1 + \Sigma_{\tau_i}^i \int_{\tau_i}^t \left(\alpha_s^i\right)^2 ds}\right) = \log\left(\frac{\Sigma_t^i}{\Sigma_{\tau_i+1}^i}\right), \tag{A.43}$$

which satisfies Equation (A.40) for any sequence $\{\alpha_t^i\}_{\tau_i}^{\tau_i+1}$. This confirms the affine conjecture for $\dot{q}_t^i$ and establishes that $\lambda_i = \Sigma_T^i$.

Since the sequence $\{\alpha_t^i\}_{\tau_i}^{\tau_i+1}$ is indeterminate, without loss I can specify $\alpha_t^i = \alpha^i$. Intuitively, only the total amount of learning about h^i is determined between time τ_i and $\tau_i + 1$. Evaluating Σ_t^i from Equation (A.42) at time $\tau_i + 1$ and recognizing $\Sigma_{\tau_i+1}^i = \lambda_i$, α^i is given by

$$\alpha^i = \sqrt{\frac{1}{\lambda} - \frac{1}{\Sigma_{\tau_i}^i}} = \sqrt{\frac{1}{\lambda_i} - \frac{1}{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}}, \tag{A.44}$$

because $\Sigma_{\tau_i}^i = \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$.

To verify that the map Q^i is indeed invertible, such that $Q^i \circ \Phi^i = I_W$, I recognize that $\phi^i = -q^i \circ \Phi^i$ and

$$\mathbb{E}_\nu\left[|\phi^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] = \mathbb{E}_\nu\left[\int_{\tau_i}^{\tau_i+1}\left(\alpha^i\right)^2\left(h^i - \hat{h}_t^i\right)^2 dt \mid \mathcal{W}_{\tau_i}\right] = \mathbb{E}_\nu\left[\int_{\tau_i}^{\tau_i+1}\left(\alpha^i\right)^2\Sigma_t^i dt \mid \mathcal{W}_{\tau_i}\right], \tag{A.45}$$

which substituting with Equation (A.42) implies the kinetic energy of Q^i satisfies

$$\frac{1}{2}\mathbb{E}_{\nu_1}\left[|\phi^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] = -\frac{1}{2}\mathbb{E}_{\nu_1}\left[\int_{\tau_i}^{\tau_i+1}\frac{d\Sigma_t^i}{\Sigma_t^i} \mid \mathcal{W}_{\tau_i}\right] = \frac{1}{2}\log\left(\frac{\Sigma_{\tau_i}^i}{\Sigma_{\tau_i+1}^i}\right) = \mathbb{D}\left(\nu_1||\mu_1\right), \tag{A.46}$$

is equal to the relative entropy. By Theorem 7 of Üstünel (2009), an adapted perturbation of identity is invertible if and only if its kinetic energy equals the relative entropy of the measure it induces. Since this equality holds, Q^i is *a.s.* invertible. By Theorem 9 of Üstünel (2009), invertibility is equivalent to the innovation conjecture: the innovations $\hat{B}_1^i$ generate the same filtration as the observations, i.e., $\mathcal{Q}^i = \mathcal{W}_1$. Consequently, the optimal filter derived on $\mathcal{B}_1^i$ can be projected onto $\mathcal{W}_1$ without any loss of information, i.e., the kinetic energy exceeds the relative entropy.

Consequently, with a Gaussian prior and a linear signal structure, given by Equation (A.35), the optimal posterior about h^i at time $\tau_i + 1$ is Gaussian $h^i \mid \mathcal{W}_t \sim \mathcal{N}\left(\hat{h}_t^i, \Sigma_t^i\right)$, where $\hat{h}_t^i$ and Σ_t^i are the first two conditional moments of h^i under measure ν^i.

Step 5: Conjecture a Linear Solution for θ^i

Recall $\bar{\boldsymbol{a}} = \mathbb{E}_{\mu_2}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i}\right]$ and define $\bar{\boldsymbol{f}} = \mathbb{E}_{\mu_1}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i}\right] = \mathbb{E}_{Q^i\mu_1}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i}\right]$ and $g^i = \boldsymbol{a}\left(s_i\right)'\bar{\boldsymbol{f}}$. Writing agent i's objective U^i as

$$U^i = -\frac{1}{2}\left(g^i - \hat{g}_{\tau_i+1}^i + \boldsymbol{a}\left(s_i\right)'\left(\boldsymbol{f} - \bar{\boldsymbol{f}}\right)\right)^2, \tag{A.47}$$

I can then recover θ^i from Equation (A.21) according to the Malliavin chain rule

$$\dot{v}_t^i = \frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[\left(g^i - \hat{g}^i_{\tau_i+1}\circ V^i\right)\left(\left(\mathcal{D}_{2t}\hat{g}^i_{\tau_i+1}\right) - \mathbb{E}_{\nu_2}\left[\mathcal{D}^B_{2t}\left(g^i - \hat{g}^i_{\tau_i+1}\right) \mid \mathcal{W}_t\right]\right)\circ V^i \mid \mathcal{B}^i_t\right], \quad \text{(A.48)}$$

where $\mathcal{D}^B_{2t}$ is the Malliavin derivative with respect to B^i_{2t}. By similar arguments to those in Step 4, I conclude the second term in Equation (A.48) vanishes.

Notice that, without loss of generality, I can express the conditional expectation $\hat{h}^i_t$, which will be a martingale under the optimal filter, as

$$d\hat{g}^i_t = \gamma^i_t\left(dW_{2t} - \mathbb{E}_{\nu_2}\left[\dot{v}^i_t \mid \mathcal{W}_t\right]dt\right). \quad \text{(A.49)}$$

I now conjecture that $\dot{v}^i_t = \delta^i_t\left(g^i - \hat{g}^i_t\circ V^i\right)$, where δ^i_t is deterministic, from which follows that $\mathbb{E}_{\nu_2}\left[\dot{v}^i_t \mid \mathcal{W}_t\right] = 0$ and from Equation (A.49) that

$$\hat{g}^i_{\tau_i+1} = \bar{g}^i + \int_{\tau_i}^{\tau_{i+1}}\gamma^i_t dW_{2t}. \quad \text{(A.50)}$$

It follows from Equation (A.50) that

$$\mathcal{D}_{2t}\hat{g}^i_{\tau_i+1} = \gamma^i_t. \quad \text{(A.51)}$$

Using the Kallianpur-Striebel formula (e.g., Liptser and Shiryaev, 2001), the conditional distribution of the task-state probabilities $\boldsymbol{\pi}$ admits a likelihood ratio representation, from which the law of motion for the conditional mean of the task-state probability vector, $\hat{\boldsymbol{\pi}}$ follows

$$d\hat{\boldsymbol{\pi}}^i_t = \delta^i_t\left(G^i - \hat{g}^i_t I_{|\mathcal{S}_i|}\right)\hat{\boldsymbol{\pi}}^i_t dW_{2t}, \quad \text{(A.52)}$$

where G^i is the $|\mathcal{S}_i|\times|\mathcal{S}_i|$ diagonal matrix with diagonal entries $\boldsymbol{a}(s_i)'\bar{\boldsymbol{f}}$ and $I_{|\mathcal{S}_i|}$ is $|\mathcal{S}_i|\times|\mathcal{S}_i|$ identity matrix. Equation (A.52) implies a law of motion for the conditional mean of g^i, $d\hat{g}^i_t$

$$d\hat{g}^i_t = \delta^i_t\sum_{s\in\mathcal{S}_i}\hat{\pi}^i_t(s_i)\left(g^i(s_i) - \hat{g}^i_t\right)^2 dW_{2t} = \delta^i_t Var_{\nu_2}\left(g^i \mid \mathcal{W}_t\right)dW_{2t}. \quad \text{(A.53)}$$

Matching Equation (A.53) with (A.49) imposes that

$$\gamma_t^i = \delta_t^i Var_{\nu_2}\left(g^i \mid \mathcal{W}_t\right). \tag{A.54}$$

Next, I recognize that $\hat{g}^i_{\tau_i+1} \circ V^i$ has a law of motion

$$d\hat{g}_t^i \circ V^i = \gamma_t^i \left(dW_{2t} + \delta_t^i \left(g^i - \hat{g}_t^i \circ V^i\right) dt\right), \tag{A.55}$$

which implies that

$$\hat{g}_t^i \circ V^i = g^i + e^{-\int_{\tau_i}^t \delta_s^i \gamma_s^i ds}\left(\bar{g}^i - g^i\right) + \int_0^t e^{-\int_s^t \delta_u^i \gamma_u^i du} \gamma_s^i dW_{2s}, \tag{A.56}$$

and consequently that

$$\mathbb{E}_{\mu_2}\left[\hat{g}^i_{\tau_i+1} \circ V^i \mid \mathcal{B}_t^i\right] - g^i = \mathbb{E}_{\mu_2}\left[e^{-\int_t^{\tau_i+1} \delta_s^i \gamma_s^i ds} \mid \mathcal{B}_t^i\right]\left(\hat{g}_t^i \circ V^i - g^i\right). \tag{A.57}$$

Substituting Equation (A.57) into (A.48), and recognizing $\gamma_t^i \circ V^i \in \mathcal{B}_t^i$, I arrive at

$$\dot{v}_t^i = \frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[e^{-\int_t^{\tau_i+1} \delta_s^i \gamma_s^i \circ V^i ds} \mid \mathcal{B}_t^i\right]\left(g^i - \hat{g}_t^i \circ V^i\right)\gamma_t^i \circ V^i, \tag{A.58}$$

Comparing to my conjecture, Equation (A.58) imposes that

$$\delta_t^i = \frac{1}{\lambda_i}\mathbb{E}_{\mu_2}\left[e^{-\int_t^{\tau_i+1} \delta_s^i \gamma_s^i \circ V^i ds} \mid \mathcal{B}_t^i\right]\gamma_t^i \circ V^i, \tag{A.59}$$

which confirms the conjecture. Substituting Equation (A.54) into (A.59), I arrive at

$$\mathbb{E}_{\mu_2}\left[e^{-\int_t^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(g^i \mid \mathcal{W}_s\right) \circ V^i ds} \mid \mathcal{B}_t^i\right] Var_{\nu_2}\left(g^i \mid \mathcal{W}_t\right) \circ V^i = \lambda_i. \tag{A.60}$$

Evaluating Equation (A.60) at $t = \tau_i + 1$ imposes

$$\lambda_i = Var_{\nu_2}\left(g^i \mid \mathcal{W}_{\tau_i+1}\right) \circ V^i, \tag{A.61}$$

from which follows by substituting this into Equation (A.60) and with $g^i = \bar{f}' \boldsymbol{a}(s_i)$

$$\mathbb{E}_{\mu_2}\left[e^{-\int_t^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i)\mid \mathcal{W}_s\right)\circ V^i ds} \mid \mathcal{B}_t^i\right] = \frac{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)\circ V^i}{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_t\right)\circ V^i}. \tag{A.62}$$

Finally, I can compute the Wonham filter's kinetic energy by recognizing $\theta^i = -v^i \circ \Theta^i$ and

$$\frac{1}{2}\mathbb{E}_{\nu_2}\left[|\theta^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] = \frac{1}{2}\mathbb{E}_{\nu_2}\left[\int_{\tau_i}^{\tau_i+1}\left(\delta_t^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_t\right)dt \mid \mathcal{W}_{\tau_i}\right]. \tag{A.63}$$

Notice that if the sequence $\{\delta_t^i\}_{\tau_i}^{\tau_i+1}$ is not too large, i.e., small gain, then I can approximate

$$\begin{aligned}\mathbb{E}_{\mu_2}\left[e^{-\int_{\tau_i}^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i)\mid \mathcal{W}_s\right)\circ V^i ds} \mid \mathcal{B}_t^i\right] &\approx 1 - \mathbb{E}_{\mu_2}\left[\int_{\tau_i}^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_s\right)\circ V^i ds \mid \mathcal{B}_t^i\right] \\ &= 1 - \mathbb{E}_{\nu_2}\left[\int_{\tau_i}^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_s\right)ds \mid \mathcal{B}_t^i\right]\end{aligned} \tag{A.64}$$

and along with Equation (A.62) evaluated at $t = \tau_i$ implies the approximate relation by the law of iterated expectations

$$\mathbb{E}_{\nu_2}\left[\int_{\tau_i}^{\tau_i+1}\left(\delta_s^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_s\right)ds \mid \mathcal{W}_{\tau_i}\right] \approx \frac{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right) - Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)}{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right)}, \tag{A.65}$$

since $Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right)$ and $Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)\circ V^i = \lambda_i$ are constants.

Consequently, under the small gain approximation, the kinetic energy from Equation (A.63) pulling back to ν_2 becomes

$$\begin{aligned}\frac{1}{2}\mathbb{E}_{\nu_2}\left[|\theta^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] &= \frac{1}{2}\mathbb{E}_{\nu_2}\left[\int_{\tau_i}^{\tau_i+1}\left(\delta_t^i\right)^2 Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_t\right)dt \mid \mathcal{W}_{\tau_i}\right] \\ &\approx \frac{1}{2}\frac{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right) - Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)}{Var_{\nu_2}\left(\bar{f}'\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right)}.\end{aligned} \tag{A.66}$$

The kinetic energy is equal to the relative entropy when the adapted perturbations, V^i and Θ^i, are invertible and satisfy finite-energy conditions. In this case, this expression also approximates the relative entropy of the measure change from μ_2 to ν_2, $\mathbb{D}(\nu_2||\mu_2)$.

Recognizing $Var_{\nu_2}\left(g^i \mid \mathcal{W}_t\right) = \bar{f}' Var_{\nu_2}\left(\boldsymbol{a}(s_i) \mid \mathcal{W}_t\right)\bar{f}$ with $Var_{\nu_2}\left(\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i}\right) = \Sigma_a$, i.e., the

prior variance of task weights, I can express Equation (A.66) as

$$\frac{1}{2}\mathbb{E}_{\nu_2}\left[|\theta^i|_H^2 \mid \mathcal{W}_{\tau_i}\right] \approx \frac{1}{2}\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} - \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}, \tag{A.67}$$

where I recall that $\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = \lambda_i$.

To confirm that the innovation filtration, $\mathcal{V}^i$, is equal to the observation filtration, $\mathcal{W}_2$, I appeal to Allinger and Mitter (1981) who show that innovation conjecture is satisfied when the signal and noise are independent and the signal has finite energy. In the present setting, the signal is the task state s_i, which is independent of W_2 under μ_2, and the finite energy condition $\mathbb{E}_{\mu_2}\left[\int_{\tau_i}^{\tau_i+1}|\theta^i|_H^2 dt\right] < \infty$ is satisfied because attention capacity is bounded by the rational inattention constraint at κ.

Step 6: Rational Inattention Constraint

Note that it must be optimal for the rational inattention constraint to bind; otherwise the agent could acquire more information and improve her decision-making. Consequently, λ_i must satisfy

$$\frac{1}{2}\mathbb{E}_{\nu}\left[|\phi^i|_H^2 + |\theta^i|_H^2 \,|\mathcal{W}_{\tau_i}\right] = \kappa, \tag{A.68}$$

provided the perturbation maps, Q^i and V^i, are invertible such that the relative entropy is equal to the kinetic energy.

Substituting with Equations (A.46) with $\Sigma^i_{\tau_i} = \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$ and (A.67) under the small gain approximation, I arrive at

$$\frac{1}{2}\log\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}{\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i}\right) + \frac{1}{2}\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} - \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = \kappa, \tag{A.69}$$

which I can rewrite as

$$\log\left(\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i\right) + \frac{\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = 1 + \log\left(\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i\right) - 2\kappa. \tag{A.70}$$

Since when there is learning about fundamentals, $\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i = \lambda_i$, while when there is learning about the task, $\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = \lambda_i$, this implies a water-filling strategy for optimal information acquisition.

If the agent learns about both the fundamentals and his task, then $\bar{\boldsymbol{a}}^{i\prime}\Omega^{i}_{\tau_i+1}\bar{\boldsymbol{a}}^{i} = \bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = \lambda_i$, and Equation (A.70) implies that λ_i solves

$$\log\left(\lambda_i\right) + \frac{\lambda_i}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = 1 + \log\left(\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}\right) - 2\kappa, \tag{A.71}$$

Expontentiating both sides of Equation (A.71), I arrive at

$$\frac{\lambda_i}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}e^{\frac{\lambda_i}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}} = \frac{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}e^{1-2\kappa}, \tag{A.72}$$

which has a known unique solution

$$\lambda_i = \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}\, W\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}e^{1-2\kappa}\right), \tag{A.73}$$

where $W\left(\cdot\right)$ is the Lambert W function.

If instead $\bar{\boldsymbol{a}}^{i\prime}\Omega^{i}_{\tau_i+1}\bar{\boldsymbol{a}}^{i} = \lambda_i > \bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i}\right)\bar{\boldsymbol{f}} = \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$, then agent devotes all his attention to learning about the fundamentals. In this case, λ_i from Equation (A.70) satisfies

$$\lambda_i = e^{-2\kappa}\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}. \tag{A.74}$$

This is optimal if and only if $\lambda_i > \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$, or $\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} < e^{-2\kappa}\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}$

By contrast, if $\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = \lambda_i > \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}$, then agent devotes all his attention to learning about his task. In this case, λ_i from Equation (A.70) satisfies

$$\lambda_i = \left(1-2\kappa\right)\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}. \tag{A.75}$$

This is optimal if and only if $\lambda_i > \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i}$, or $\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i} < \left(1-2\kappa\right)\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$. This can only be a solution if $\kappa < \frac{1}{2}$ since $\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^{i} > 0$.

Since only one of these three conditions can hold, the information acquisition policy of agent i is unique.

This complete my characterization of the agent's optimal information acquisition strategy.

Step 7: Expected Utility

Recognizing that

$$\begin{aligned}\mathbb{E}\left[\left(\boldsymbol{a}\left(s_i\right)-\bar{\boldsymbol{a}}^i\right)'\Omega_{\tau_i+1}\left(\boldsymbol{a}\left(s_i\right)-\bar{\boldsymbol{a}}^i\right)\right] &= tr\left(\Omega^i_{\tau_i+1}\mathbb{E}_{\nu_2}\left[\left(\boldsymbol{a}\left(s_i\right)-\bar{\boldsymbol{a}}^i\right)\left(\boldsymbol{a}\left(s_i\right)-\bar{\boldsymbol{a}}^i\right)'\right]\right)\\ &= tr\left(\Omega^i_{\tau_i+1}\Sigma_a\right), \end{aligned} \tag{A.76}$$

with some manipulation that Equation (A.4) can be written as

$$\begin{aligned}\mathbb{E}_{\nu^i}\left[U^i\right] &= -\frac{1}{2}\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i-\frac{1}{2}\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\mid\mathcal{W}_{\tau_{i+1}}\right)\bar{\boldsymbol{f}}-\frac{1}{2}Tr\left(\Sigma_a\Omega^i_{\tau_i+1}\right)\\ &\quad -\frac{1}{2}Tr\left(\left(\Omega_0-\Omega^i_{\tau_i+1}\right)Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\mid\mathcal{W}_{\tau_{i+1}}\right)\right). \end{aligned} \tag{A.77}$$

Proof of Proposition 2

Recall that each agent j that lives on $t\in[\tau_j,\tau_{j+1}]$ generates an optimal change of measure ν_1^j over the fundamentals, $\boldsymbol{f}$, parameterized by

$$\frac{d\nu^j}{d\mu_1}=L\left(\varphi^j\right)=\rho\left(-\delta\varphi^j\right), \tag{A.78}$$

where $\varphi^j\in L^2(\mu,H)$ is supported on $[\tau_j,\tau_{j+1}]$ and $\rho\left(-\delta\varphi^j\right)=\exp\left(-\delta\left(\varphi^j\right)-\frac{1}{2}|\varphi^j|^2_H\right)$ is the Girsanov exponential. Under ν_1^j, the observation process admits the signal-plus-noise decomposition

$$dW_{1t}=\dot{q}^j_t dt+dB^j_{1t},\qquad t\in\left[\tau_j,\tau_{j+1}\right], \tag{A.79}$$

where $\dot{q}^j_t=\alpha^j\left(h^j-\hat{h}^j_t\circ Q^j\right)$ is adapted to $\mathcal{B}^j_t=\mathcal{W}_t\vee\sigma\left(\boldsymbol{f},s_j\right)$ and $\varphi^j=-q^j\circ\Phi^j$.

I now extend each $\dot{\varphi}^j$ to all of $t\geq 0$ by setting $\dot{\varphi}^j_t=0$ for $t\notin\left[\tau_j,\tau_{j+1}\right]$.

Step 1: Representing Past Agent (Private) Information

Consider the optimal changes of measure ν^j parameterized by $\frac{d\nu^j}{d\mu_1}=L(\varphi^j)$ for $j=0,1,\ldots,i-1$. Because the agents live on disjoint time intervals $[\tau_j,\tau_j+1)$, the Cameron-Martin elements $\varphi^j\in H$ have disjoint temporal supports. Define

$$\varphi^a=\sum_{j=0}^{i-1}\varphi^j,\qquad \text{so that } \dot{\varphi}^a_t=\dot{\varphi}^j_t \text{ for } t\in[\tau_j,\tau_j+1). \tag{A.80}$$

The Brownian increments $\{W_{1t} - W_{1s} : s, t \in [\tau_j, \tau_{j+1})\}$ are independent across generations under μ_1 because they correspond to disjoint time intervals of the canonical Brownian motion. Consequently, the Skorokhod integrals decomposes

$$\delta(\varphi^a) = \sum_{j=0}^{i-1} \delta\left(\varphi^j\right), \qquad |\varphi^a|_H^2 = \sum_{j=0}^{i-1} |\varphi^j|_H^2, \tag{A.81}$$

where the cross-terms $\int_0^{\tau_i} \dot{\varphi}_t^j \dot{\varphi}_t^k \, dt = 0$ for $j \neq k$ by disjointness of supports. It follows that the Girsanov exponential factorizes according to

$$L(\varphi^a) = \exp\left(-\delta(\varphi^a) - \frac{1}{2}|\varphi^a|_H^2\right) = \prod_{j=0}^{i-1} \exp\left(-\delta\left(\varphi^j\right) - \frac{1}{2}|\varphi^j|_H^2\right) = \prod_{j=0}^{i-1} L\left(\varphi^j\right). \tag{A.82}$$

Note that I would arrive at the same result if agents learned from the same realizations of the Brownian motions through sequential updating. In this case, agent j updates from ν^{j-1} (rather than from μ_1), and

$$L\left(\varphi^{j-1}\right) \cdot \exp\left(-\int_{\tau_j}^{\tau_{j+1}} \dot{\varphi}_t^j \, dW_{1t} - \frac{1}{2}\int_{\tau_j}^{\tau_{j+1}} |\dot{\varphi}_t^j|^2 \, dt\right) = L\left(\varphi^{j-1} + \varphi^j\right), \tag{A.83}$$

because the cross-term $\int_0^{\tau_{j+1}} \dot{\varphi}_t^j \dot{\varphi}_t^{j-1} \mathbf{1}_{[\tau_{j-1}, \tau_j)} \mathbf{1}_{[\tau_j, \tau_{j+1})} \, dt = 0$ vanishes by disjointness of temporal supports. Iterating on Equation (A.83) yields $L(\varphi^a) = L\left(\sum_{j=0}^{i-1} \varphi^j\right)$.

This further identifies an adapted perturbation of identity, Φ^a, for which I can construct an inverse, $Q^a = I_W + q^a$, where $\dot{q}_t^a = \dot{q}_t^j$ for $t \in [\tau_j, \tau_j + 1)$. This induces a signal

$$dW_{1t} = \dot{q}_t^a dt + dB_{1t}^a. \tag{A.84}$$

A subtle issue is that it need not be the case that the AI can generically recover the each agent's innovation $\hat{B}_{1t}^i$ from simply observing W_{1t}, i.e., that $\mathcal{Q}^a = \mathcal{W}_1$. However, from Proposition 1, I have shown that the innovation conjecture holds for each agent i. Since the agents' filters are disjoint over time, this consequently holds for the AI as well. If the innovation conjecture failed, which it could for a generic nonlinear filter, then there would be information loss from the AI recovering instead $\Phi^a \circ Q^a = Z^a$, where Z^a is a μ_1-Brownian motion with respect to $\mathcal{Q}_t^a \subset \mathcal{W}_t$.

Consequently, without any alteration to the data that the artificial intelligence receives, an

agent i that uses artificial intelligence would start from an updated prior μ_1^a where $\frac{d\mu_1^a}{d\mu_1} = L(\phi^a)$.[30]

An important distinction from agent learning is that the AI is trying to recover a posterior for the vector of fundamentals, $\boldsymbol{f}$, rather than some linear combination of them, i.e., $\bar{\boldsymbol{a}}^{i\prime}\boldsymbol{f}$. Consequently, its prior is $\boldsymbol{f} \sim \mathcal{N}\left(\bar{f}, \Omega_0\right)$.

Step 2: Blackwell Garbling of Past Agents' Information

Suppose that the AI observes only the observations on a coarser filtration $\mathcal{Z}_t \subset \mathcal{W}_t$. This induces a coarser measure, $\hat{\nu}^a$, adapted to $\mathcal{Z}_t$, and I impose that $\mathcal{Z}_t$ is generated by a diffusion process so that it induces a continuous measure when ν^a is restricted to $\mathcal{Z}_t$.[31] Let $\mathcal{E}_t = \mathcal{Z}_t \vee \{\boldsymbol{f}\}$. It is immediate that

$$\begin{aligned}
\mathbb{E}_{\nu_1^a}\left[1_{\{x\in A\}} \mid \mathcal{E}_0\right] &= \mathbb{E}_{\mu_1}\left[\frac{d\mu_1^a}{d\mu_1} 1_{\{x\in A\}} \mid \mathcal{E}_0\right] = \mathbb{E}_{\mu}\left[\mathbb{E}_{\mu_1}\left[\frac{d\mu_1^a}{d\mu_1} 1_{\{x\in A\}} \mid \mathcal{B}_0^i\right] \mid \mathcal{E}_0\right] \\
&= \mathbb{E}_{\mu_1}\left[\mathbb{E}_{\mu_1}\left[L(\phi) 1_{\{x\in A\}} \mid \mathcal{B}_0^i\right] \mid \mathcal{E}_0\right] = \mathbb{E}_{\mu_1}\left[L(\phi) 1_{\{x\in A\}} \mid \mathcal{E}_0\right] \\
&= \mathbb{E}_{\mu_1}\left[\mathbb{E}_{\mu_1}\left[L(\phi) \mid \mathcal{E}_0\right] 1_{\{x\in A\}} \mid \mathcal{E}_0\right] \qquad \text{(A.85)}
\end{aligned}$$

and because $\frac{d\mu_1^a}{d\mu_1} \in \mathcal{B}_t^i$ for $\mathcal{E}_t \subset \mathcal{B}_t^i$, I conclude $dt \times d\mu_1$ *a.s.*

$$\frac{d\hat{\mu}_1^a}{d\mu_1} = \mathbb{E}_{\mu_1}\left[\frac{d\mu_1^a}{d\mu_1} \mid \mathcal{E}_0\right] = \mathbb{E}_{\mu_1}\left[L(\phi^a) \mid \mathcal{E}_0\right] = L\left(\hat{\phi}^a\right). \qquad \text{(A.86)}$$

Recognizing that

$$\begin{aligned}
d\mathbb{E}_{\mu_1}\left[L(\phi^a) \mid \mathcal{E}_0\right] &= -\mathbb{E}_{\mu_1}\left[L(\phi^a)\dot{\phi}\left(dW_{1t} + \hat{\dot{\phi}}_t^a dt - \hat{\dot{\phi}}_t^a dt\right) \mid \mathcal{E}_t\right] \\
&= -\mathbb{E}_{Q^a\mu_1}\left[\dot{\phi}_t^a d\hat{B}_{1t}^a \mid \mathcal{E}_t\right] + \mathbb{E}_{Q^a\mu_1}\left[\dot{\phi}_t^a \hat{\dot{\phi}}_t^a dt \mid \mathcal{E}_t\right], \qquad \text{(A.87)}
\end{aligned}$$

[30] A subtle observation is that the generational structure of the agents ensures that their signals, $\hat{B}_{1t}^i$ do not overlap in time. Otherwise, I would have to concatenate overlapping posteriors for overlapping agent i according to $\phi^i \circ \Phi^a$ rather than ϕ^i. This is because the aggregate inverse Φ^a is no longer the identity on agent i's interval $[\tau_i, \tau_i + 1)$ when other agents simultaneously contribute signals.

[31] Otherwise, Duffie and Lando (2001) show that if the coarser filtration involves discretely sampling diffusion processes, then the restriction of ν^a to $\mathcal{E}_t$ induces a jump rather than a diffusion measure. This would instead require the jump information structure described in Footnote 4.

I have, since $\mathbb{E}_{\mu_1}\left[L\left(\phi^a\right) \mid \mathcal{E}_0\right] = L\left(\hat{\phi}^a\right)$, that

$$dL\left(\hat{\phi}^a\right) = -\frac{\mathbb{E}_{\mu_1}\left[L\left(\phi^a\right)\dot{\phi}_t^a \mid \mathcal{E}_t\right]}{\mathbb{E}_{\mu_1}\left[L\left(\phi^a\right) \mid \mathcal{E}_t\right]} L\left(\hat{\phi}^a\right)\left(d\hat{B}_{1t}^a - \hat{\phi}_t^a dt\right) = -\mathbb{E}_{\mu_1^a}\left[\dot{\phi}_t^a \mid \mathcal{E}_t\right] L\left(\hat{\phi}^a\right) dZ_{1t}^a, \tag{A.88}$$

from which follows, since $dL\left(\hat{\phi}^a\right) = -L\left(\hat{\phi}^a\right)\dot{\hat{\phi}}_t^a dZ_{1t}^a$, that

$$\dot{\hat{\phi}}_t^a = \mathbb{E}_{Q^a\mu_1}\left[\dot{\phi}_t^a \mid \mathcal{E}_t\right]. \tag{A.89}$$

Consequently, a garbling of the signals that the AI uses for its training data is equivalent to adapting ϕ^a to the coarser filtration $\mathcal{E}_t$ under the pushed forward measure $Q^a\mu_1$, i.e., the image of μ_1 under Q^a.

Step 3: Post-processing of Garbled Posterior with RLHF

Suppose the designers of the artificial intelligence model employ a post-processing algorithm to the AI after gathering and curating its training data, which from Step 2 induced a posterior indexed $\hat{\phi}^a$ adapted to $\mathcal{E}_t$. This post-processing training is akin to the Reinforcement Learning and Human Feedback (RLHF) step used in many text-prediction Large Language Models (LLMs).

Specifically, the designers impose a tilt for the posterior, parameterized by $L\left(\varphi^a\right)$, that solves the static variational problem

$$\mathcal{L}_A = \sup_{\frac{d\mu_1^A}{d\hat{\mu}_1^a} \in L_+^1\left(\hat{\mu}_1^a, \mathcal{E}_0\right)} \mathbb{E}_{\hat{\mu}_1^a}\left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} R\left(\boldsymbol{f}\right) - \lambda_T \mathbb{D}\left(\mu_1^A || \hat{\mu}_1^a\right) \mid \mathcal{E}_0\right] - \eta_a^i\left(\mathbb{E}_{\hat{\mu}_1^a}\left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} \mid \mathcal{E}_0\right] - 1\right), \tag{A.90}$$

where $R\left(\boldsymbol{f}\right)$ is the reward function and λ_T corresponds to the temperature of the AI. As is standard, the solution to this problem can be recovered by applying the Donsker-Varadhan Lemma

$$\lambda_T \log \mathbb{E}_{\hat{\mu}_1^a}\left[e^{R(\boldsymbol{f})/\lambda_T} \mid \mathcal{E}_0\right] = \sup_{\frac{d\mu_1^A}{d\hat{\mu}_1^a} \in L_+^1\left(\hat{\mu}_1^a, \mathcal{E}_0\right)} \left\{\mathbb{E}_{\hat{\nu}^a}\left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} R\left(\boldsymbol{f}\right) - \lambda_T \mathbb{D}\left(\mu_1^A || \hat{\mu}_1^a\right) \mid \mathcal{E}_0\right] - \eta_a^i\left(\mathbb{E}_{\hat{\nu}^a}\left[\frac{d\mu_1^A}{d\hat{\mu}_1^a} \mid \mathcal{E}_0\right] - 1\right)\right\}, \tag{A.91}$$

where

$$\frac{d\mu_1^A}{d\hat{\mu}_1^a} = \frac{e^{R(\boldsymbol{f})/\lambda_T}}{\mathbb{E}_{\mu_1^A}\left[e^{R(\boldsymbol{f})/\lambda_T} \mid \mathcal{E}_0\right]}, \tag{A.92}$$

and I assume that $\mathbb{E}_{\hat{\mu}_1^a}\left[e^{R(f)/\lambda_T} \mid \mathcal{E}_0\right] < \infty$.

In what follows, I select a quadratic reward objective, $R(\boldsymbol{f})$, for the AI

$$R(\boldsymbol{f}) = -\frac{1}{2}(\boldsymbol{f} - \boldsymbol{f}^*)' \Gamma (\boldsymbol{f} - \boldsymbol{f}^*), \tag{A.93}$$

where the $N \times 1$ vector $\boldsymbol{f}^*$, and the $N \times N$ matrix Γ are reward function parameters. As such, Equation (A.94) becomes

$$\frac{d\mu_1^A}{d\hat{\mu}_1^a} = \frac{e^{-\frac{1}{2\lambda_T}(\boldsymbol{f}-\boldsymbol{f}^*)'\Gamma(\boldsymbol{f}-\boldsymbol{f}^*)}}{\mathbb{E}_{\hat{\mu}_1^a}\left[e^{-\frac{1}{2\lambda_T}(\boldsymbol{f}-\boldsymbol{f}^*)'\Gamma(\boldsymbol{f}-\boldsymbol{f}^*)} \mid \mathcal{E}_0\right]}, \tag{A.94}$$

Step 4: Gaussian Learning over Fundamentals

Recall that, in equilibrium, the optimal information acquisition of agent i is $\hat{q}_t^i = \alpha^i\left(h^i - \hat{h}_t^i\right) = \alpha^i \bar{\boldsymbol{a}}^{i\prime}\left(\boldsymbol{f} - \hat{\boldsymbol{f}}_t^i\right)$. Concatenating these Cameron-Martin shifts across agents into the AI, I find that the AI effectively observes a degenerate signal process about $\boldsymbol{f}$

$$d\hat{B}_{1t}^i = \alpha^i \bar{\boldsymbol{a}}^{i\prime}\left(\boldsymbol{f} - \hat{\boldsymbol{f}}_t^i\right)dt + dW_{1t}. \tag{A.95}$$

With a Gaussian prior at $t = 0$, $\boldsymbol{f} \sim \mathcal{N}\left(\bar{\boldsymbol{f}}, \Omega_0\right)$, it follows that the optimal filter will be the Kalman-Bucy filter with a time-varying Kalman-Bucy gain.

To Blackwell garble the signal process in a manner that maintains its continuity, I inject noise into $d\hat{B}_{1t}^i$ by adding a second Brownian motion, W_{3t}, to it with diffusion $\sqrt{\frac{1}{\beta^2} - 1}$ for $\beta \leq 1$ such that the AI instead uses the less informative signal

$$d\tilde{B}_{1t}^i = \alpha^i \bar{\boldsymbol{a}}^{i\prime}\left(\boldsymbol{f} - \hat{\boldsymbol{f}}_t^a\right)dt + dW_{1t} + \sqrt{\frac{1}{\beta^2} - 1}W_{3t} = \alpha^i \bar{\boldsymbol{a}}^{i\prime}\left(\boldsymbol{f} - \hat{\boldsymbol{f}}_t^a\right)dt + \frac{1}{\beta}d\tilde{W}_{1t}. \tag{A.96}$$

A lower β corresponds to more noise injected into the signal.

It is immediate that based on the properties of the Kalman-Bucy filter, this gives rise to a posterior at time t for the AI, $\hat{\mu}_1^a$, is given by $\boldsymbol{f} \mid \mathcal{E}_t \sim \mathcal{N}\left(\hat{\boldsymbol{f}}_t^P, \Omega_t^a\right)$ where

$$d\hat{\boldsymbol{f}}_t^P = \beta\alpha^i \Omega_t^P \bar{\boldsymbol{a}}^i d\tilde{B}_{1t}^i, \tag{A.97}$$

and

$$d\Omega_t^P = -\left(\beta\alpha^i\right)^2 \Omega_t^P \bar{\boldsymbol{a}}^i \bar{\boldsymbol{a}}^{i\prime} \Omega_t^P dt. \tag{A.98}$$

Recognizing that $d\left(\Omega_t^P\right)^{-1} = -\left(\Omega_t^P\right)^{-1} d\Omega_t^P \left(\Omega_t^P\right)^{-1}$, I can rewrite the Ricatti Equation (A.98) as

$$\left(\Omega_{\tau_i}^P\right)^{-1} = \Omega_0^{-1} + \sum_{j=0}^{i-1} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}. \tag{A.99}$$

Finally, applying the Gibbs tilt from Equation (A.94) to $\hat{\mu}_1^a$, I conclude that the AI's posterior after RLHF at time τ_i is given by the Gaussian measure $\mu_1^{AI} \sim \mathcal{N}\left(\hat{\boldsymbol{f}}_{\tau_i}^A, \Omega_{\tau_i}^A\right)$, where

$$\hat{\boldsymbol{f}}_{\tau_i}^A = \Omega_{\tau_i}^A \left(\left(\Omega_{\tau_i}^P\right)^{-1} \hat{\boldsymbol{f}}_{\tau_i}^P + \frac{1}{\lambda_T}\Gamma \boldsymbol{f}^*\right) = \Omega_{\tau_i}^A \left(\left(\Omega_0^{-1} + \sum_{j=0}^{i-1} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}\right) \hat{\boldsymbol{f}}_{\tau_i}^P + \frac{1}{\lambda_T}\Gamma \boldsymbol{f}^*\right), \tag{A.100}$$

and

$$\left(\Omega_{\tau_i}^A\right)^{-1} = \frac{1}{\lambda_T}\Gamma + \left(\Omega_{\tau_i}^P\right)^{-1} = \Omega_0^{-1} + \sum_{j=0}^{i-1} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime} + \frac{1}{\lambda_T}\Gamma. \tag{A.101}$$

In my model, I choose $f^* = \bar{f}$ both to avoid introducing a bias that agents must rationally anticipate (or form expectations about) and to give an interpretation of the RLHF tilt as acting as a form of Bayesian shrinkage toward the prior mean.

Step 5: Non-invertibility of the AI Posterior

Suppose agent i observes the AI's posterior when composed with her prompt, $\bar{\boldsymbol{a}}^i, \mathcal{N}\left(\bar{\boldsymbol{a}}^{i\prime}\hat{\boldsymbol{f}}_{\tau_i}^A, \bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A\bar{\boldsymbol{a}}^i\right)$, where $\hat{\boldsymbol{f}}_{\tau_i}^A$ is given by Equation (A.100) and $\Omega_{\tau_i}^A$ is given by Equation (A.101). It is immediate that agent i cannot unravel the tilt that the AI developers impose through the RLHF post-processing to recover $\hat{\boldsymbol{f}}_{\tau_i}^P$ and $\Omega_{\tau_i}^P$.

Proof of Proposition 3

The case in which agent i chooses not to acquire information is already characterized in Proposition 1. I first examine the case in which the agent employs the AI from τ_i to $\tau_i + 1$. I then consider her incentives to cancel her subscription halfway through her tenure. In the case where agent i uses AI, her informative measure is $\nu^i = Q^i \mu_1^{AI}$

Step 1: Characterization if AI is used from τ_i to τ_{i+1}

The key difference in this case is the prior that agent i uses when forming expectations about the fundamentals $\boldsymbol{f}$. Modifying the rational inattention constraint Equation (A.69) to reflect the use of AI, in which the agent anticipates the higher precision in her initial prior over $\boldsymbol{f}$ from the AI's posterior, $\Omega^A_{\tau_i}$,

$$\frac{1}{2}\log\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i}{\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i}\right)+\frac{1}{2}\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}-\bar{\boldsymbol{f}}'Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\mid\mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}=\kappa. \tag{A.102}$$

Conditional on this updated prior, the optimal choice of effort, K^i, still satisfies Equation (A.3). Given the disjoint time intervals between the AI's learning (up to time τ_i) and agent i's learning (from time τ_i to τ_i+1, I do not need to modify the map Φ^i based on the AI's transport map, Φ^A, i.e., on $[\tau_i,\tau_i+1]$, $Q^i\circ Q^A=Q^i$.

Since Γ is positive semi-definite, $\Omega_0-\Omega^A_{\tau_i}$ is also positive semi-definite. If this were not the case, then the AI would have a more diffuse posterior than the prior and agent i would never rationally pay to adopt it.

Step 2: Incentives to Use AI at τ_i

Agent i will use the AI at time τ_i if the differential expected utility from doing so exceeds the total subscription cost from τ_i to τ_i+1, c. The expected utility from not using the AI is given by Equation (A.77) in the proof of Proposition 1. It is straightforward to compute the expected utility from using the AI, and having the updated (but potentially biased) prior over the fundamentals, $\mathcal{N}\left(\hat{\boldsymbol{f}}^A_{\tau_i},\Omega^A_{\tau_i}\right)$

$$\begin{aligned}\mathbb{E}_{\nu^i}\left[U^i\mid\mathcal{W}_{\tau_i}\right] &= -\frac{1}{2}\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i-\frac{1}{2}\mathbb{E}_{\nu^i}\left[\hat{\boldsymbol{f}}^{A\prime}_{\tau_i}Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\mid\mathcal{W}_{\tau_i+1}\right)\hat{\boldsymbol{f}}^A_{\tau_i}\mid\mathcal{W}_{\tau_i}\right]\\ &\quad -\frac{1}{2}Tr\left(\left(\Omega^A_{\tau_i}-\Omega^i_{\tau_i+1}\right)Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\mid\mathcal{W}_{\tau_i+1}\right)\right)-\frac{1}{2}Tr\left(\Sigma_a\Omega^i_{\tau_i+1}\right),\end{aligned} \tag{A.103}$$

which telescopes back to Equation (A.77) but generically with different optimal choices of $\Omega^i_{\tau_i+1}$ and $Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\right)$.

Consequently, the difference in expected utility from using the AI, $\Delta\mathbb{E}_{\nu^i}\left[U^i\right]$, is given by

$$\begin{aligned}\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] &= -\frac{1}{2}\bar{\boldsymbol{a}}^{i\prime}\Delta\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i - \frac{1}{2}\bar{\boldsymbol{f}}'\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\bar{\boldsymbol{f}} - \frac{1}{2}Tr\left(\Omega_0\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right) \\ &\quad -\frac{1}{2}Tr\left(\Sigma_a\Delta\Omega^i_{\tau_i+1}\right) + \frac{1}{2}\Delta Tr\left(\Omega^i_{\tau_i+1}Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right). \qquad \text{(A.104)}\end{aligned}$$

The first two terms are weakly positive because they are quadratic forms of $\Delta\Omega^i_{\tau_i+1}$ and $\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)$, respectively, which are (weakly) negative semi-definite matrices because agent i can learn more about both fundamentals and tasks when using the AI's updated posterior. The third and fourth term are weakly positive because the trace of a negative semi-definite matrix and a positive semi-definite matrix is non-positive. However, the fifth term is (weakly) negative by the product rule and again recognizing the trace of a (weakly) negative semi-definite matrix multiplied by a positive semi-definite matrix is non-positive.

Notice that the overall difference is positive because

$$\begin{aligned}\Delta Tr\left(\Omega^i_{\tau_i+1}Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right) &= Tr\left(\Delta\Omega^i_{\tau_i+1}Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)|_{AI}\right) \\ &\quad + Tr\left(\Omega^i_{\tau_i+1}|_{No\ AI}\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right), \qquad \text{(A.105)}\end{aligned}$$

and consequently

$$\begin{aligned}&\frac{1}{2}\Delta Tr\left(\Omega^i_{\tau_i+1}Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right) - \frac{1}{2}Tr\left(\Sigma_a\Delta\Omega^i_{\tau_i+1}\right) - \frac{1}{2}Tr\left(\Omega_0\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right)\right)\right) \\ = \ &-\frac{1}{2}Tr\left(\Delta\Omega^i_{\tau_i+1}\left(\Sigma_a - Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)|_{AI}\right)\right) - \frac{1}{2}Tr\left(\left(\Omega_0 - \Omega^i_{\tau_i+1}|_{No\ AI}\right)\Delta Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right)\right) \\ > \ &0, \qquad \text{(A.106)}\end{aligned}$$

because the trace of the product of a positive and a negative semi-definite matrix is nonpositive.

As such, I conclude from Equations (A.104) and (A.106) that

$$\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq 0. \qquad \text{(A.107)}$$

Agent i therefore adopts AI from time τ_i to $\tau_i + 1$ if $\Delta\mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq c$ and does not otherwise.

Step 3: Incentives to Discontinue AI Use at $\tau_i + \frac{1}{2}$

Suppose at time τ_i, agent i who adopts AI decides that she will cancel her subscription at time $\tau_i + \frac{1}{2}$. In this case, since she chose to adopt AI from time τ_i to $\tau_i + \frac{1}{2}$, it must be the case that $\Delta \mathbb{E}_{\nu^i}\left[U^i \mid \mathcal{W}_{\tau_i}\right] \geq \frac{c}{2}$.

However, notice that if she cancels her subscription, then she loses all of the AI's information when making her final effort decision. This tightens her rational inattention constraint and forces her to choose how to continue learning to satisfy it. However, by revealed preference, this altered information acquisition plan must be (weakly) suboptimal in expectation to the plan she would have made if she choose how to acquire information without using AI from time τ_i to $\tau_i + 1$. Her expected utility at time τ_i of her utility at time $\tau_i + \frac{1}{2}$ is therefore bounded from above by her expected utility from not using the AI at all.

Not buying an AI subscription consequently strictly dominates buying a subscription at time τ_i and discontinuing it at $\tau_i + \frac{1}{2}$, both because of the (weakly) higher expected continuation utility at time $\tau_i + \frac{1}{2}$ from adopting and discontinuing, and the saved cost of the initial subscription, $\frac{c}{2}$.

I therefore conclude that agent i either buys an AI subscription from time τ_i to $\tau_i + 1$ or does not buy one at all.

Step 4: Substitution Toward Learning about Tasks

Suppose agent i adopts AI at time τ_i. Then it is immediate from comparing Equations (A.69) and (A.102) that agent i devotes (weakly) more attention to learning about her task state, s_i. To see this, I examine the three cases of information acquisition.

First, suppose all information acquisition is devoted to learning about the fundamentals. Then, from the proof of Proposition 1, $\lambda_i = e^{-2\kappa} \bar{\boldsymbol{a}}^{i\prime} \Omega^A_{\tau_i} \bar{\boldsymbol{a}}^i < e^{-2\kappa} \bar{\boldsymbol{a}}^{i\prime} \Omega_0 \bar{\boldsymbol{a}}^i$. This tightens the constraint because it is now more likely that $\bar{\boldsymbol{f}}' \Sigma_a \bar{\boldsymbol{f}} > \lambda_i$, violating the condition for agent i to learn only about the fundamentals.

Second, suppose all information acquisition is devoted to learning about the task state. Then, from the proof of Proposition 1, $\lambda_i = \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right) \bar{\boldsymbol{f}} = \left(1 - 2\kappa\right) \bar{\boldsymbol{f}}' \Sigma_a \bar{\boldsymbol{f}}$ and it is more likely that $\lambda_i > \bar{\boldsymbol{a}}^{i\prime} \Omega^A_{\tau_i} \bar{\boldsymbol{a}}^i$, the required condition for agent i to learn only about his task state.

Finally, if agent i acquires information about both the fundamentals and his task state. Then, from the proof of Proposition 1, $\lambda_i = \bar{\boldsymbol{a}}^{i\prime} \Omega^i_{\tau_i+1} \bar{\boldsymbol{a}}^i = \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right) \bar{\boldsymbol{f}}$ and λ_i now solves

from Equation (A.102)

$$\log(\lambda_i) + \frac{\lambda_i}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = \log\left(\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i\right) + 1 - 2\kappa. \tag{A.108}$$

Since $\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i \leq \bar{\boldsymbol{a}}^{i\prime}\Omega^0_{\tau_i}\bar{\boldsymbol{a}}^i$ if the agent uses AI, λ_i is lower when the agent uses the AI. Applying the Implicit Function Theorem to Equation (A.108), I find

$$\frac{\partial \log(\lambda_i)}{\partial \log\left(\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i\right)} = \frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}{\lambda_i + \bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} < 1, \tag{A.109}$$

and consequently $\log\lambda_i$ increases less than one-for-one with $\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i$ so that attention shifts more toward learning about her task. As such, κ_s (weakly) increases with the use of AI.

Step 5: Value of a Subscription

Notice because $\hat{\boldsymbol{f}}^A_{\tau_i}$ is an unbiased conditional expectation for $\boldsymbol{f}$, agent i cannot forecast it when choosing whether to buy an AI subscription. Consequently, from Equation (A.102), the subscription's ex ante value is increasing in the extent to which $\bar{\boldsymbol{a}}^{i\prime}\Omega^A_{\tau_i}\bar{\boldsymbol{a}}^i$ is smaller than $\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$. Consequently, the vector $\bar{\boldsymbol{a}}^i$ that maximizes the subscription's ex ante value solves the following program (since the expected task weights $\bar{\boldsymbol{a}}^i$ sum to 1)

$$\inf_{\boldsymbol{x}} \frac{1}{2}\boldsymbol{x}'\left(\Omega^A_{\tau_i} - \Omega_0\right)\boldsymbol{x}, \tag{A.110}$$

$$s.t. \quad : \quad \boldsymbol{x}'\iota_N = 1,$$

$$s.t. \quad : \quad \boldsymbol{x} \geq 0 \qquad (element - by - element), \tag{A.111}$$

where $\Omega^A_{\tau_i} - \Omega_0$ is a negative semi-definite matrix.

Recognizing that this is a minimization over a concave objective, the answer is a corner solution on the N-dimensional simplex, i.e., $\boldsymbol{x} = \boldsymbol{e}_n$, where $\boldsymbol{e}_n$ is the Euclidian $N \times 1$ basis vector with the n^{th} entry equal to 1 and the rest zeroes, and n is chosen as the vertex such that $n = \arg\min_n\{\Omega^A_{\tau_i} - \Omega_0\}_{nn}$, i.e., the minimum diagonal entry of the matrix $\Omega^A_{\tau_i} - \Omega_0$.

By contrast, the AI adds the least ex ante value for the vector $\bar{\boldsymbol{a}}^i$ that solves the following

program

$$\sup_{\boldsymbol{x}} \frac{1}{2}\boldsymbol{x}' \left(\Omega^A_{\tau_i} - \Omega_0\right) \boldsymbol{x}, \tag{A.112}$$

$$s.t. \quad : \quad \boldsymbol{x}'\boldsymbol{\iota}_N = 1,$$

$$s.t. \quad : \quad \boldsymbol{x} \geq 0 \qquad (element - by - element). \tag{A.113}$$

If none of the nonnegativity constraints bind, and $\Omega_0 - \Omega^A_{\tau_i}$ is nonsingular, then this problem has the straightforward solution

$$\boldsymbol{x} = \frac{1}{\boldsymbol{\iota}'_N \left(\Omega_0 - \Omega^A_{\tau_i}\right)^{-1} \boldsymbol{\iota}_N} \left(\Omega_0 - \Omega^A_{\tau_i}\right)^{-1} \boldsymbol{\iota}_N, \tag{A.114}$$

where $\boldsymbol{\iota}_N$ is the $N \times 1$ vector of ones.

Proof of Proposition 4

Step 1: Innovation to Invention with AI

Suppose agent i adopts AI. If there is no change in the task space, agent $i+1$ will also purchase an AI subscription because her expected task is the same as agent i's, i.e., $\bar{\boldsymbol{a}}^{i+1} = \bar{\boldsymbol{a}}^i$, and the AI has incrementally learned more about the fundamentals associated with the average task from agent i's training data. Then, from Proposition 3, κ^i_s is higher because of AI and therefore so is the Poisson compensator, $\zeta\left(\kappa^i_s\right) dt$, for a jump in the task space. Consequently, AI usage increases the likelihood of the invention of new tasks.

Step 2: Invention to Innovation with AI

Suppose that there is a Poisson jump at time $t \in [\tau_i, \tau_i + 1]$ and the task space has changed, i.e., some previous tasks are now antiquated and some new tasks are added to $\mathcal{S}_t$. Notice that the AI's precision, given by Equation (A.101), beyond the time-invariant RLHF tilt, $\frac{1}{\lambda_T}\Gamma$, is based on the history of expected task weights of previous agents, $\sum_{j=0}^{i} \left(\beta\alpha^j\right)^2 \bar{\boldsymbol{a}}^j \bar{\boldsymbol{a}}^{j\prime}$. Consequently, the expertise of the AI is backward-looking and concentrated on the fundamentals on which past agents expected their tasks to load the most.

Next, recall from the proof of Proposition 3 that agents for whom $\bar{\boldsymbol{a}}'\left(\Omega_{\tau_i}^A - \Omega_0\right)\bar{\boldsymbol{a}}$ is least negative benefit the least from their subscription when acquiring information. Suppose that $\Omega_{\tau_i}^A - \Omega_0$ is a rank deficient negative semi-definite matrix. Then, this is maximized when

$$\left(\Omega_{\tau_i}^A - \Omega_0\right)\bar{\boldsymbol{a}} = \mathbf{0}, \tag{A.115}$$

i.e., $\bar{\boldsymbol{a}}$ is in the null space of $\Omega_{\tau_i}^A - \Omega_0$.

If this is the case, then the AI has no impact on the optimal information acquisition policy value for the first agent after the jump. As such, from (A.104), $\Delta\mathbb{E}_{\nu^{i+1}}\left[U^{i+1}\right] = 0$, and that agent does not purchase an AI subscription.[32]

If agent $i+1$ does not adopt AI, then she (weakly) devotes more attention to learning about fundamentals since she then uses the prior about fundamentals, $\mathcal{N}\left(\bar{\boldsymbol{f}}, \Omega_0\right)$. As such, κ_s^{i+1} (weakly) decreases, and the likelihood of a jump to a new task span consequently (weakly) falls.

However, if agent $i+1$ does not use AI, then there exists a future agent $j > i+1$ that adopts AI in the absence of another jump because the AI uses the training data from all agents $i+1$ to j (not just subscribers) to improve the AI's precision about the new span of task states. Agent j must exist because non-subscribing agents' fundamental learning updates the AI in the new task directions, so the AI's precision about the relevant fundamentals increases until the benefit of subscribing to AI exceeds c, provided that c is not too large.

Consequently, after an invention of new task states, i.e., a jump, AI adoption can then begin to recover, and similarly with the rate of invention, $\zeta\left(\kappa_s^i\right)$.

Step 3: Invention and Innovation Cycles

From the logic of Steps 1 and 2, the use of AI increases the likelihood of invention (a jump that changes the task span), after which there is a finite number of generations that do not use AI until the training data has sufficiently updated in the absence of a subsequent jump. In this case, the pace of invention falls as agents devote more attention to learning about fundamentals and eschew AI subscriptions.

[32] Alternatively, when the matrix is full rank, the quadratic form is minimized when the new expected task weights satisfy Equation (A.114). However, in this case, generically $\Delta\mathbb{E}_{\nu^{i+1}}\left[U^{i+1}\right] \neq 0$.

Consequently, the economy cycles between periods of AI adoption and task innovation (in which the Poisson rate is high) and periods of AI disruption and fundamental re-learning (in which the Poisson rate falls). The cycle is self-sustaining because each phase creates the conditions for the next.

Proof of Proposition 5

Recall from Equation (27) that Ω_t^A is decreasing in a positive semi-definite sense in the RLHF tilt, $\frac{1}{\lambda_T}\Gamma$ such that if $\Gamma_1 \succeq \Gamma_2$, then $\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A(\Gamma_1)\,\bar{\boldsymbol{a}}^i \leq \bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A(\Gamma_2)\,\bar{\boldsymbol{a}}^i$.

Since from Equation (A.109), λ_i is increasing in $\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A\bar{\boldsymbol{a}}^i$, and from Equation (A.109) falls less than one-for-one with $\bar{\boldsymbol{a}}^{i\prime}\Omega_{\tau_i}^A\bar{\boldsymbol{a}}^i$, it follows that agent i acquires less information about the fundamentals under RLHF tilt Γ_1 than under RLHF tilt Γ_2.

By similar logic, I recognize that decreasing the temperature, λ_T, is equivalent to dilating all elements of Γ such that $\frac{1}{\lambda_{T1}}\Gamma \succeq \frac{1}{\lambda_{T2}}\Gamma$. I therefore conclude agent i acquires less information about the fundamentals if the temperature λ_T falls.

Proof of Proposition 6

I begin with the optimal pricing, c_{21}, and temperature, λ_{T1}, policies of the developer, taking β and Γ, which are fixed at time 0, as given.

Step 1: Optimal Pricing c_1

Because the developer is a large player, it is immediate that he will engage in discriminatory pricing with each agent if he wants the agent to subscribe to the AI. As a result, he will set the subscription cost c_1 to extract all the agent's surplus from using the AI according to

$$c_1 = \Delta\mathbb{E}_{\nu^1}\left[U^1 \mid \mathcal{W}_1\right]. \tag{A.116}$$

It is consequently optimal for agent 1 to purchase a subscription, $m_1^1 = 1$, if

$$c_1 \geq \chi_1, \tag{A.117}$$

Step 2: Optimal Temperature λ_{T1}

Notice that the temperature, λ_{T1}, impacts $\Delta\mathbb{E}_{\nu^2}\left[U^1 \mid \mathcal{W}_{\tau_1}\right]$ only through the first term, which depends on the AI's posterior at $t=1$, μ_1^A, and that by the Envelope Condition applied to Equation (A.103)

$$\frac{d\Delta\mathbb{E}_{\nu^1}\left[U^1 \mid \mathcal{W}_{\tau_1}\right]}{d\lambda_{T1}} = -\frac{1}{2}Tr\left(\frac{d\Omega_1^A}{d\lambda_{T1}}Var_{\nu_1}\left(\boldsymbol{a}\left(s_1\right) \mid \mathcal{W}_2\right)\right) - \frac{\lambda_1}{2\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}\frac{d\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}{d\lambda_{T1}} < 0, \tag{A.118}$$

since $\frac{d\Omega_1^A}{d\lambda_{T1}} = \frac{1}{\lambda_{T1}^2}\Omega_1^A\Gamma\Omega_1^A$ is a positive semi-definite matrix, as is $Var_{\nu_1}\left(\boldsymbol{a}\left(s_1\right) \mid \mathcal{W}_{\tau_2}\right)$, and the trace of the product of two positive semi-definite matrices is nonnegative. Furthermore, $\frac{d\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}{d\lambda_{T1}} > 0$ from Proposition 5.

Therefore, from Step 1,

$$\frac{dc_1}{d\lambda_{T1}} < 0, \tag{A.119}$$

and therefore the developer chooses $\lambda_{T1} = \underline{\lambda}_T$.

This complete the characterization of the developer's static problem at $t=1$.

Proof of Proposition 7

Step 1: Optimal Pricing c_0

Similar to the case of $t=1$, the developer will engage in discriminatory pricing with agent 0 if he wants the agent to subscribe to the AI. As a result, he will set the subscription cost c_0 to extract all the agent's surplus from using the AI according to

$$c_0 = \Delta\mathbb{E}_{\nu^0}\left[U^0 \mid \mathcal{W}_0\right]. \tag{A.120}$$

Since the subscription price does not impact agent 0's information acquisition decision, the optimal subscription pricing is still determined by static incentives.

Step 2: Optimal Data Degradation β

Notice that the degree of data degradation β only affects the developer's problem at $t=1$ through the precision of the AI's posterior and the cost of running the AI, χ_1. The first-order

necessary condition, applying the Envelope Condition applied to Equation (A.103), is

$$\begin{aligned}\frac{d\left(\mathbb{E}_{\mu^D}\left[c_1-\chi_1 \mid \mathcal{W}_0\right]\right)}{d\beta} &= -\frac{1}{2}\mathbb{E}_{\mu^D}\left[Tr\left(\frac{d\Omega_1^A}{d\beta}Var_{\nu_1}\left(\boldsymbol{a}\left(s_1\right) \mid \mathcal{W}_2\right)\right)+\frac{\lambda_1}{\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}\frac{d\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}{d\beta} \mid \mathcal{W}_0\right] \\ &\quad -C\beta\left(\alpha^0\right)^2 \geq 0, \end{aligned} \tag{A.121}$$

with equality if $\beta < 1$. Substituting $\frac{d\Omega_1^A}{d\beta} = -2\beta\left(\alpha^0\right)^2\Omega_1^A\bar{\boldsymbol{a}}^0\bar{\boldsymbol{a}}^{0\prime}\Omega_1^A$ and $\frac{d\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1}{d\beta}$ and employing the property of the trace operator $Tr\left(ABC\right) = Tr\left(CAB\right)$, Equation (A.121) becomes

$$\mathbb{E}_{\mu^D}\left[\bar{\boldsymbol{a}}^{0\prime}\Omega_1^A Var_{\nu_1}\left(\boldsymbol{a}\left(s_1\right) \mid \mathcal{W}_2\right)\Omega_1^A\bar{\boldsymbol{a}}^0+\lambda_1\frac{\left(\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^0\right)^2}{\bar{\boldsymbol{a}}^{1\prime}\Omega_1^A\bar{\boldsymbol{a}}^1} \mid \mathcal{W}_0\right] \geq C. \tag{A.122}$$

The expectation is taken over the likelihood of a jump in the task span at $t = 1$.

For C in an intermediate range, β will have an interior solution.

Step 3: Optimal Temperature λ_{T0}

Building on the derivation for the optimal temperature at $t = 1$ from Proposition 6, the first-order condition is

$$\begin{aligned}-\frac{1}{2}Tr\left(\frac{d\Omega_0^A}{d\lambda_{T0}}Var_{\nu_0}\left(\boldsymbol{a}\left(s_0\right) \mid \mathcal{W}_1\right)\right)-\frac{\lambda_0}{2\bar{\boldsymbol{a}}^{0\prime}\Omega_0^A\bar{\boldsymbol{a}}^0}\frac{d\bar{\boldsymbol{a}}^{0\prime}\Omega_0^A\bar{\boldsymbol{a}}^0}{d\lambda_{T0}}+\frac{d\left(\mathbb{E}_{\mu^D}\left[c_1-\chi_1 \mid \mathcal{W}_0\right]\right)}{d\lambda_{T0}} \\ +\frac{d\kappa_s^0}{d\lambda_{T0}}\zeta'\left(\kappa_s^0\right)\Delta\mathbb{E}_{\mu^D}\left[c_1 \mid \mathcal{W}_0\right] \leq 0, \end{aligned} \tag{A.123}$$

with equality if $\lambda_{T0} > \underline{\lambda}_T$. Notice that $\frac{d\left(\mathbb{E}_{\mu^D}\left[c_1-\chi_1 \mid \mathcal{W}_0\right]\right)}{d\lambda_{T0}} \propto \frac{d\left(\mathbb{E}_{\mu^D}\left[c_1-\chi_1 \mid \mathcal{W}_0\right]\right)}{d\beta}$, i.e., the impact of a change in temperature on future data quality is proportional to a change in the degradation of the data because temperature impacts future data quality through agent 0's choice of Kalman-Bucy gain, α^0. Both derivatives balance the marginal benefit of more informative training data against the marginal cost of a higher cost of operating the AI for the developer. Therefore, the condition for the optimal β, Equation (A.122), has an interior solution, $\frac{d\left(\mathbb{E}_{\mu^D}\left[c_1-\chi_1 \mid \mathcal{W}_0\right]\right)}{d\lambda_{T0}} = 0$.

Recognizing this, $\kappa_s^0 = \kappa - \kappa_f^0 = \kappa - \log\left(\frac{\bar{\boldsymbol{a}}^{0\prime}\Omega_0^A\bar{\boldsymbol{a}}^0}{\bar{\boldsymbol{a}}^{0\prime}\Omega_1^1\bar{\boldsymbol{a}}^0}\right)$, and substituting $\frac{d\Omega_0^A}{d\lambda_{T0}} = \frac{1}{\lambda_{T0}^2}\Omega_0^A\Gamma\Omega_0^A$,

Equation (A.123) becomes

$$-Tr\left(\Omega_0^A\Gamma\Omega_0^A Var_{\nu_0}\left(\boldsymbol{a}\left(s_0\right)\mid\mathcal{W}_1\right)\right)-\frac{\bar{\boldsymbol{a}}^{0\prime}\Omega_0^A\Gamma\Omega_0^A\bar{\boldsymbol{a}}^0}{\bar{\boldsymbol{a}}^{0\prime}\Omega_0^A\bar{\boldsymbol{a}}^0}\left(\lambda_0+\zeta'\left(\kappa_s^0\right)\Delta\mathbb{E}_{\mu^D}\left[c_1\mid\mathcal{W}_0\right]\right)\leq 0, \quad \text{(A.124)}$$

As in the case of λ_{T1}, The first two terms are negative (lower temperature inflates the perceived value of the AI). However, the third term related to the change in the subscription fee at $t=1$, c_1, can be positive if the higher probability of a change in task span reduces the value of the AI.

Step 4: Optimal RLHF Tilt Γ

Consider a variation $\Gamma^\epsilon=\Gamma+\epsilon G$ for a positive semi-definite (p.s.d.) matrix G, and I recognize that $\boldsymbol{x}'A\boldsymbol{x}=Tr\left(A\boldsymbol{x}\boldsymbol{x}'\right)$. Then, at a stationary point, the developer's optimal choice of Γ satisfies

$$Tr\left(\left[\sum_{t=0}^{1}\frac{\rho^t}{2\lambda_{Tt}}\Omega_t^A\mathbb{E}_{\mu^D}\left[Var_{\nu_t}\left(\boldsymbol{a}\left(s_t\right)\mid\mathcal{W}_{t+1}\right)+\frac{\bar{\boldsymbol{a}}^t\bar{\boldsymbol{a}}^{t\prime}}{\bar{\boldsymbol{a}}^{t\prime}\Omega_t^A\bar{\boldsymbol{a}}^t}\left(\lambda_t+\zeta'\left(\kappa_s^0\right)\Delta\mathbb{E}_{\mu^D}\left[c_1\mid\mathcal{W}_0\right]\mathbf{1}_{\{t=0\}}\right)\mid\mathcal{W}_0\right]\Omega_t^A-\vartheta I_N+S\right]G\right)=0, \quad \text{(A.125)}$$

where ϑ is the Lagrange multiplier on the $Tr\left(\Gamma\right)=1$ constraint, S is the multiplier on the constraint that Γ is p.s.d., and I recognize from the previous steps that if β has an interior solution, then the terms related to the marginal benefit and cost of improving the training data at $t=1$ vanish. I search for a Γ matrix that is p.s.d.

Define

$$F_0=\sum_{t=0}^{1}\frac{\rho^t}{2\lambda_{Tt}}\Omega_t^A\mathbb{E}_{\mu^D}\left[Var_{\nu_t}\left(\boldsymbol{a}\left(s_t\right)\mid\mathcal{W}_{t+1}\right)+\frac{\bar{\boldsymbol{a}}^t\bar{\boldsymbol{a}}^{t\prime}}{\bar{\boldsymbol{a}}^{t\prime}\Omega_t^A\bar{\boldsymbol{a}}^t}\left(\lambda_t+\zeta'\left(\kappa_s^0\right)\Delta\mathbb{E}_{\mu^D}\left[c_1\mid\mathcal{W}_0\right]\mathbf{1}_{\{t=0\}}\right)\mid\mathcal{W}_0\right]\Omega_t^A, \quad \text{(A.126)}$$

and notice that F_0 is symmetric.

Equation (A.125) becomes

$$Tr\left(\left(F_0-\vartheta I_N+S\right)G\right)=0, \quad \text{(A.127)}$$

which implies since G is an arbitrary p.s.d. matrix that $F_0-\vartheta I_N$ is negative semi-definite. Consequently, ϑI_N-F_0 is p.s.d. Choosing $G=\Gamma$, Equation (A.127) becomes

$$Tr\left(\left(\vartheta I_N-F_0\right)\Gamma\right)=0, \quad \text{(A.128)}$$

Recognizing the product of two p.s.d. matrices has trace zero only if the two matrices are

orthogonal to each other, Equation (A.128) implies

$$F_0\Gamma = \vartheta\Gamma. \tag{A.129}$$

Therefore, ϑ is an eigenvalue of F_0 and Γ (a linear combination of) its eigenspace such that, imposing the trace constraint

$$\vartheta = Tr\left(F_0\Gamma\right), \tag{A.130}$$

and Equation (A.129) becomes

$$F_0\Gamma = Tr\left(F_0\Gamma\right)\Gamma. \tag{A.131}$$

Step 4: AI Subscription at $t = 0$

Since the developer is risk-neutral, he sells a subscription to agent 0 at $t = 0$ if it (weakly) raises his overall expected profit

$$c_0 - \chi_0 \geq -\rho\Delta\mathbb{E}_{\mu^D}\left[c_1 - \chi_1 \mid \mathcal{W}_0\right]. \tag{A.132}$$

Proof of Proposition 8

Step 1: Optimal Effort

Following the proof of Proposition 1, agent i's optimal level of effort at time $\tau_i + 1$ is now

$$\begin{aligned} K^i &= \mathbb{E}_{\nu_2^i}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right]' \mathbb{E}_{\nu_1^i}\left[\left(1-w\right)\boldsymbol{f} + w\mathbb{E}_{\mu_1^A}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_i+1}\right] \mid \mathcal{W}_{\tau_{i+1}}\right] \\ &= \mathbb{E}_{\nu_2^i}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right]' \mathbb{E}_{\nu_1^i}\left[\left(1-w\right)\boldsymbol{f} + w\hat{\boldsymbol{f}}_{\tau_i}^A \mid \mathcal{W}_{\tau_{i+1}}\right]. \end{aligned} \tag{A.133}$$

Because of the independence of posteriors, her optimal effort is her best estimate of her task state multiplied by a weighted average between her best estimate of the hidden fundamentals and her best estimate of the AI's best estimate of the fundamentals.

Regardless of whether agent i purchases an AI subscription, her optimal effort collapses to

$$K^i = \mathbb{E}_{\nu_2^i}\left[\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_{i+1}}\right]' \mathbb{E}_{\nu_1^i}\left[\boldsymbol{f} \mid \mathcal{W}_{\tau_{i+1}}\right]. \tag{A.134}$$

This is because either agent i adopts AI, and consequently shares the same posterior as the AI, or she does not adopt it, and then her best forecast of the AI's (unbiased) conditional expectation is her own conditional expectation.

Step 2: The Optimal Filters

It is straightforward to see that the altered objective does not change how the agent learns about the fundamentals. The agent continues to use the Kalman-Bucy filter to learn about the fundamentals. This is because the agent does not internalize that altering the path of her own information acquisition shapes the AI's information acquisition because it is her training data that the AI uses to update its posterior at time $\tau_i + 1$.

Similarly, she continues to learn about her task state through $g^i = \bar{\boldsymbol{f}}' \boldsymbol{a}(s_i)$ when she does not purchase an AI subscription, and through $g^i = \hat{\boldsymbol{f}}^{A\prime}_{\tau_i} \boldsymbol{a}(s_i)$ when she does, using the controlled Wonham filter as characterized in the proof of Proposition 1.

Step 3: Optimal Information Acquisition Policy

The key difference from the baseline model is agent i's information acquisition policy. This is because when learning from the Kalman-Bucy filter, agent i learns about $h^i = \bar{\boldsymbol{a}}^{i\prime}\left((1-w)\boldsymbol{f} + w\hat{\boldsymbol{f}}^A_{\tau_i+1}\right)$ regardless of whether she buys an AI subscription.

Following the proof of Proposition 1, and maintaining the small gain approximation for the Wonham filter, her rational inattention constraint reduces in the case that she does not buy an AI subscription reduces to

$$\frac{1}{2}\log\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\left(\Omega_0 - w(2-w)\Omega^A_{\tau_i+1}\right)\bar{\boldsymbol{a}}^i}{\bar{\boldsymbol{a}}^{i\prime}\left(\Omega^i_{\tau_i+1} - w^2\Omega^A_{\tau_i+1}\right)\bar{\boldsymbol{a}}^i}\right) + \frac{1}{2}\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} - \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = \kappa, \tag{A.135}$$

where I recognize that $Cov_{\nu^i}\left(\boldsymbol{f}, \hat{\boldsymbol{f}}^A_{\tau_i+1} \mid \mathcal{W}_{\tau_i+1}\right) = \Omega^i_{\tau_i+1}$ and $Var_{\nu^i}\left(\hat{\boldsymbol{f}}^A_{\tau_i+1} \mid \mathcal{W}_{\tau_i+1}\right) = \Omega^i_{\tau_i+1} - \Omega^A_{\tau_i+1}$ for noise $\epsilon^A_{\tau_i}$ in the AI's initial estimate of $\boldsymbol{f}$, $\hat{\boldsymbol{f}}^A_{\tau_i}$.

If she adopts AI, the rational inattention constraint is instead

$$\frac{1}{2}\log\left(\frac{\bar{\boldsymbol{a}}^{i\prime}\left(\Omega^A_{\tau_i} - w^2\Omega^A_{\tau_i+1}\right)\bar{\boldsymbol{a}}^i}{(1-w)^2\,\bar{\boldsymbol{a}}^{i\prime}\Omega^i_{\tau_i+1}\bar{\boldsymbol{a}}^i}\right) + \frac{1}{2}\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} - \bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}(s_i) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}}}{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}} = \kappa, \tag{A.136}$$

because the AI and agent share the same posterior at time $\tau_i + 1$ so that $\hat{\boldsymbol{f}}^A_{\tau_i+1} \in \mathcal{W}_{\tau_i+1}$ under ν^i.

In both cases, it is clear that having agent i's utility depend on the AI's validation reduces the unconditional variance of h^i below that in the case where $w = 0$. This reduces the attention that agent i devotes to learning about the fundamentals through two forces. First, reducing the numerator of the first term in both Equation (A.135) and (A.136), i.e., there is less initial uncertainty about h^i because part of it is now driven by $\hat{\boldsymbol{f}}^A_{\tau_i+1}$, which has lower (un)conditional variance than $\boldsymbol{f}$. Second, it reduces the denominators of the first terms in both expressions, so that a higher $\Omega^i_{\tau_i+1}$ achieves the same denominator as in the case in which $w = 0$.

Consequently, I conclude that $w > 0$ reduces the incentives of an agent to acquire information about the fundamentals regardless of whether she adopts AI.

Step 4: Knowledge Collapse

Suppose $w = 1$ so that agent i's entire utility is derived from the AI's evaluation, and suppose by way of contraction that agent i adopts AI. If she adopts the AI, then from Equation (A.136), the only feasible solution is $\Omega^i_{\tau_i+1} = \Omega^A_{\tau_i+1} = \Omega^A_{\tau_i}$, i.e., the agent relies on the AI and devotes zero attention to acquiring information about the fundamentals. In this case, knowledge collapses because the AI does not receive any new unique training data from which to update its posterior. As this is true at $t = 0$, it follows that $\Omega^A_{\tau_i} = \Omega^A_0 = \Omega_0$ and $\hat{\boldsymbol{f}}^A_{\tau_i} = \bar{\boldsymbol{f}}$. However, since agent i starts with the same prior, and the AI subscription costs $c > 0$, she will never pay for it, a contradiction.

If she does not adopt AI, then an equilibrium in which $\Omega^A_{\tau_i+1} = \Omega_0$ is self-fulfilling. Since the AI adds zero value in terms of higher precision about fundamentals, $\Delta\mathbb{E}\left[U^i \mid \mathcal{W}_{\tau_i}\right] = 0$ and agent i will never adopt since $c > 0$. Consequently, knowledge collapse is the unique equilibrium.

Suppose now $w \in (0, 1)$. From Proposition 1, if agent i devotes all her attention to learning about her task, then $\bar{\boldsymbol{f}}' Var_{\nu_2}\left(\boldsymbol{a}\left(s_i\right) \mid \mathcal{W}_{\tau_i+1}\right)\bar{\boldsymbol{f}} = (1-2\kappa)\,\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}$. From Equation (A.136), agent i now does not acquire any information about the fundamentals if $(1-2\kappa)\,\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} \geq \left(1-w^2\right)\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$. Therefore, if $w \geq \sqrt{1-(1-2\kappa)\,\frac{\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}}}{\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i}}$ for all i, then knowledge collapse is the unique equilibrium. Given that $\bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$ is bounded, it follows there exists a minimum w^* such that for $w \geq w^*$, knowledge collapse is the unique equilibrium. If $(1-2\kappa)\,\bar{\boldsymbol{f}}'\Sigma_a\bar{\boldsymbol{f}} > \bar{\boldsymbol{a}}^{i\prime}\Omega_0\bar{\boldsymbol{a}}^i$, then trivially $w^* = 0$.