The AI Skill Trap

Zhifeng Cai*

May 5, 2026

Abstract

This paper studies aggregate implications of artificial intelligence (AI), focusing on long-run skill formation. Workers accumulate expertise through learning-by-doing and AI learns from an aggregate dataset of workers' past actions. Using AI raises workers' short-run productivity but may slow skill growth. A data quality externality—workers do not bear the cost their AI use imposes on aggregate data quality—leads to over-adoption of AI. This generates an *AI skill trap* with multiple steady states. In the calibrated economy, the skill distribution is bimodal, with a high-skill group coexisting with a fully AI-dependent low-skill "underclass". A planner internalizes this externality, but does not eliminate the trap. Improving AI's data-processing efficiency does not weaken the case for intervention; the welfare gains rise with it.

JEL Classification: E24, J24, O33, D62

Keywords: artificial intelligence, human capital, skill formation, data quality, externality, polarization

*First draft: March, 2026. Department of Economics, Rutgers University. Email: zhifeng.cai@rutgers.edu. I am grateful to Pascual Restrepo, Mathieu Taschereau-Dumouchel, and Laura Veldkamp for comments. While AI provided substantial assistance in computation, simulation, and draft writing, all errors remain my own.

1 Introduction

As AI transforms knowledge work, a growing concern is that individuals who delegate tasks to AI may fail to accumulate the expertise they would otherwise develop. A growing body of evidence supports this concern: access to AI impairs learning in educational settings (Bastani et al., 2025; Jošt et al., 2024; Barcaui, 2025), slows skill acquisition among software developers (Shen and Tamkin, 2026), and reduces diagnostic performance among physicians who rely on AI assistance (Budzyń et al., 2025). AI developers have voiced a similar concern: that less-skilled workers could permanently dependent on AI, forming a kind of "underclass" (Amodei, 2024, 2026).

Why does this matter? If AI can perform a task, why must workers retain the underlying skill? After all, past waves of automation have rendered certain human skills obsolete without obvious social cost. The situation with modern AI is different: its continued improvement depends on a steady supply of high-quality human-generated data (Sutskever, 2024; Villalobos et al., 2024), and when workers delegate to AI their actions increasingly reflect what AI already knows rather than new human ideas. A computer-science literature documents the consequence: when AI systems are recursively trained on their own outputs, successive generations degrade in quality (Alemohammad et al., 2023; Briesch et al., 2023; Shumailov et al., 2024; Gerstgrasser et al., 2024).[1] The loss of human expertise is therefore not merely a private cost but a social one.

Individual workers, however, do not internalize this cost. The private benefits of AI use are immediate and concentrated—AI improves current task performance, particularly for less-experienced workers (Brynjolfsson et al., 2025; Cui et al., 2025b; Galdin and Silbert, 2025)—while the resulting decline in data quality is gradual and external to any individual's payoff. This divergence between private and social returns gives rise to a *data quality externality* (Cai, 2024): from a social perspective, AI is underpriced, and its decentralized adoption is inefficiently high. These observations raise two questions this paper intends to address: how does AI shape skill formation and the long-run distribution of expertise, and what role, if any, is there for policy intervention?

We develop a model in which long-lived workers choose each period how to allocate tasks between their own labor and AI, allowing for any degree of human-AI collaboration. The model features two key ingredients. The first is *skill accumulation through learning-by-doing*. A worker's

[1] Whether reinforcement learning and synthetic data can substitute for fresh human signals—and where they can't—is a question we return to below.

private expertise evolves through a Bayesian learning process whose speed depends on her intensity of AI use: performing tasks through her own labor generates rich feedback that builds expertise rapidly, while delegating to AI yields only coarse signals and accumulates skill at a much lower rate. Workers optimally choose a mix of labor and AI that trades off current task performance against future skill growth, and this mix shifts endogenously over a worker's career as her expertise evolves.

The second ingredient is the data quality externality. Workers who perform tasks with their own labor contribute informative signals that reflect their expertise; workers who delegate to AI contribute actions that largely replicate information already in the dataset. Cai (2024) studies this mechanism in a setting with short-lived agents, where each generation is born fresh and the externality bears only on the quality of AI at a point in time. In our long-lived setting, the same externality additionally distorts workers' lifetime skill accumulation: a worker who over-adopts AI today both contributes less informative data and fails to build the expertise that would have made her future actions more valuable sources of data. This dynamic interactions between individual skill dynamics and the aggregate dataset is what generates the qualitatively new phenomenon at the center of the paper—the skill trap—which is absent in the short-lived framework.

The interaction between these two ingredients produces S-shaped individual skill dynamics and, in the aggregate, an AI skill trap. Because AI is underpriced, workers on the margin are pushed toward higher AI intensity than is socially optimal. Their skills then depreciate: after several periods of heavy AI use, their expertise has deteriorated to the point where switching back to labor is no longer worthwhile—the investment required to rebuild skill exceeds what a finite career can support. A critical skill threshold emerges: workers who start above it invest in labor, accumulate expertise, and converge to a high steady state, while workers who start below it rely on AI, experience persistent skill depreciation, and are permanently absorbed into a low steady state. The long-run skill distribution therefore depends on how new workers enter the economy. Under bottom entry—where new cohorts arrive with low initial skill, as is natural for recent graduates—the trap is particularly severe, since entrants are immediately drawn into AI dependence.

The mechanism does not require AI use to impair learning at every intensity. What it requires is only that workers who delegate *heavily* to AI slow their own skill accumulation—an empirical

regularity with relatively strong support in the empirical and computer-science evidence cited above. In a later section we relax the benchmark model to allow moderate AI use to be genuinely *complementary* to learning, so that workers who blend modest AI assistance with their own effort accumulate skill *faster* than under pure labor. The trap persists: as long as full delegation continues to slow learning, the bimodal long-run distribution and the welfare case for intervention survive.

We calibrate the model by targeting moments of the pre-AI economy—including wage growth over the life cycle (Lagakos et al., 2018) and worker turnover rates—and match key features of early AI adoption patterns and learning outcomes under AI use. The post-AI distribution depends on how new workers enter, so we report results under two entry distributions: a *uniform* entry, in which initial skill is drawn uniformly across the pre-AI support, representing an economy with diverse educational backgrounds and lateral entry; and a *bottom* entry, in which every entrant arrives at the lowest pre-AI skill level, representing the natural case for recent graduates in knowledge-intensive occupations. The multiple steady states yield a *bimodal* aggregate skill distribution under both entries: a low-skill mass at the AI trap coexists with a high-skill mass at the expert level. Under uniform entry the two masses are roughly comparable, because entrants split across the tipping point. Under bottom entry the low-skill mass dominates, since every entrant starts below it. In both cases a substantial mass is absorbed into a *trap*—a low-skill region from which skill cannot recover within a finite career, because the marginal return to labor is too small to overcome past AI-induced skill loss. Average skill and aggregate welfare fall substantially below what is attainable once the data externality is corrected, driven by a sharp deterioration of aggregate data quality.

We analyze welfare through a constrained planner who directly chooses each worker's AI intensity to maximize aggregate welfare, internalizing how these choices shape the evolution of aggregate data quality and the entire skill distribution. The data quality externality is the key source of inefficiency, but correcting it does not abolish the trap. Quantitatively, the planner cuts mean AI intensity from 0.73 to 0.40 under uniform entry and from 0.88 to 0.39 under bottom entry; aggregate data quality recovers by 27% under uniform entry and by 51% under bottom entry, and mean skill rises from 20.3 to 23.8 and from 16.8 to 21.7, respectively. The long-run welfare gains are large: 7.2% in lifetime terms and 5.9% in flow terms under uniform entry, 11.8% and 12.3% respectively under bottom entry, where the decentralized distortion is most severe. These gains

remain economically meaningful across sensitivity exercises varying the AI learning rate, the AI signal precision, and the turnover rate.[2]

What the planner does *not* do is eliminate the AI-dependent low-skill group. Under both entry distributions, the planner's stationary skill distribution remains visibly bimodal: a residual mass of workers stays in full AI use near the bottom of the support. The optimal correction is therefore best read not as a uniform reduction in AI use but as a redistribution of mass: for the lowest-skill cohort the labor-investment return cannot overcome the deskilling cost of past AI reliance, so the planner concedes this region and concentrates the correction on the medium range where the marginal return to building skill is largest and the data-supply benefit is strongest. Bimodality, in this reading, is a robust feature of an AI economy with a data externality and learning differences under AI: the externality governs its severity but not its existence. The trap has two distinct sources. The first is a market failure that the planner can correct: each worker fails to internalize her impact on the shared dataset. The second is a technological constraint that the planner cannot: in the trap region, learning by doing is too slow to overcome past skill loss within a finite career, so even a well-designed policy must concede this region to AI dependence.

A further finding addresses the natural conjecture that the role for intervention may vanish as AI technology improves. We show that this conjecture is not valid: even a twenty-fold improvement in how efficiently AI processes data barely moves the competitive equilibrium. When nearly all workers delegate to AI, their recorded actions in the dataset are mostly echoes of the existing dataset, and processing efficiency cannot extract valuable information that was never generated. The market fails to provide adequate incentives for novel data generation, so a role for policy intervention persists regardless of how advanced AI's data-processing technology becomes. To preserve the supply of novel human-generated data, the planner keeps aggregate AI intensity substantially below the competitive level—rising only from 0.39 to 0.61 as γ_D grows twenty-fold—while the competitive equilibrium drifts toward full AI ($\bar{\alpha}$ from 0.88 to 0.99). The flow welfare gain from intervention grows correspondingly, from $+12.3\%$ at the baseline to $+36.6\%$ at a twenty-fold improvement.

[2]The welfare gains are smaller once we account for the transition path—about $+1.5\%$ under uniform entry and $+0.3\%$ under bottom entry in discounted lifetime terms—because the planner must accept a temporary reduction in flow output while workers rebuild skill under labor, and because the planner's discount factor $\beta_s = 0.96$ attenuates the long-run benefits of higher data quality. The transition gains are nonetheless positive in both cases, indicating that the planner's reallocation is worthwhile even after netting out the short-run adjustment cost.

An important assumption for the argument above is that human-generated data is essential to AI improvement. The natural objection is that reinforcement learning (Christiano et al., 2017; Ouyang et al., 2022; Bai et al., 2022) and synthetic data (Gunasekar et al., 2023; Dubey et al., 2024) can substitute for fresh human signals, bootstrapping further progress without continued human input. They can—but only in narrow regimes, where reinforcement learning is most effective as the reward signal is both *well-defined* (the system can verify whether an action moves closer to the target) and *frequent* (so improvement compounds at scale). Closed-form games, formal mathematics, and protein-structure prediction satisfy both conditions, and self-play has extended the frontier in those domains. Most knowledge work satisfies neither: the reward for a clinical assessment, a legal analysis, or a strategic recommendation is delayed, multidimensional, often subjective, and observable only after a human takes the action and lives with the consequences (so that the reward function is sparse). In these settings, reinforcement learning still depends on human preferences and ratings or on real-world outcomes that only humans can produce, and synthetic-data pipelines remain bounded by their human-trained base model. Our results therefore apply most strongly in tasks where the worker's action is itself the primary record of her expertise—tacit professional judgment, complex case reasoning, strategic decisions—and where the quality of that action becomes apparent only through delayed, noisy, often subjective feedback. The mechanism is weaker in tasks where outcomes are objectively verifiable regardless of who acted, because the dataset can be enriched through those outcomes even under heavy AI delegation.

Related literature. This paper is most closely related to the emerging literature on AI, data, and knowledge accumulation. Cai (2024) models the feedback loop between AI adoption and data quality with short-lived agents, showing that AI-generated data can corrupt future datasets. The present paper extends that framework to long-lived agents, adding skill formation and the resulting polarization; the model-collapse phenomenon (Shumailov et al., 2024) provides the empirical foundation for the data-quality channel. Several contemporaneous papers study related aspects of AI and collective learning, but none closes the feedback loop from data back to AI. Acemoglu et al. (2026a) model agentic AI substituting for human learning effort with exogenous AI precision; Acemoglu et al. (2026b) study a DeGroot model where AI is a mechanical aggregator with exogenous weights; in both, AI quality does not respond to workers' choices. Ide (2025) studies how AI-driven automation of entry-level tasks disrupts intergenerational tacit-knowledge trans-

mission through a mentor-composition channel that is complementary to ours. Empirical work corroborates the data-supply concern: del Rio-Chanona et al. (2024) document a 25% decline in weekly Stack Overflow contributions following ChatGPT's release, and Doshi and Hauser (2024) show generative AI raises individual creativity while reducing collective novelty.

More broadly, this paper sits in the literature on automation, AI, and human capital. The task-based framework (Acemoglu and Restrepo, 2018, 2019, 2020, 2022; Acemoglu, 2025) studies how automation reshapes the task content of jobs and wage inequality; complementary work analyzes distributional consequences (Korinek and Stiglitz, 2019; Korinek and Suh, 2024; Moll et al., 2022; Guerreiro et al., 2022; Hémous and Olsen, 2022) or frames AI as a prediction technology that changes returns to human judgment (Agrawal et al., 2019; Autor, 2024). The classical human-capital framework (Becker, 1964; Ben-Porath, 1967) studies life-cycle skill investment but does not feature the option to substitute AI for skill. A common feature of this literature is that human capital is treated as exogenous; our contribution is to endogenize it. The resulting polarization is also distinct from the classical *job-polarization* literature (Autor et al., 2003; Goos and Manning, 2007; Goos et al., 2014; Jaimovich and Siu, 2020), where computerization disproportionately displaces routine occupations. In our model, polarization occurs *within* the same job: workers performing identical tasks diverge in skill because of differential investment driven by the AI externality. Skill polarization arises not from the nature of tasks being automated but from the endogenous feedback between AI adoption and human capital accumulation.

Finally, the paper contributes to the literature on endogenous information quality and social learning. Grossman and Stiglitz (1980) show that the informativeness of market prices depends on how many traders acquire costly information, creating a tension between private acquisition and public revelation; our setting has an analogous structure for the AI dataset. Veldkamp (2011) develops a broader framework for information choice in macroeconomics, and Farboodi and Veldkamp (2025) treat data as an endogenous economic resource. Puri and Veldkamp (2025) show that AI-driven personalization can amplify cognitive inequality through information choices interacting with heterogeneous priors; the polarization in our paper is complementary but arises from an endogenous feedback between each worker's AI intensity and the quality of the shared dataset. The classical social-learning literature (Banerjee, 1992; Bikhchandani et al., 1992; Vives, 1993) analyzes how aggregation breaks down when private incentives to act informatively are weak; Fajgel-

baum et al. (2017) show such information externalities can generate uncertainty traps. Our skill trap is the human-capital counterpart.

Paper outline. Section 2 presents the model. Section 3 characterizes the competitive equilibrium and derives conditions for polarization. Section 4 describes the calibration. Section 5 presents the quantitative results. Section 6 formulates the constrained planner's problem and analyzes optimal policy. Section 7 concludes.

2 Model

2.1 Primitives, Technology, and Friction

Time is discrete and runs forever, $t = 0, 1, 2, \ldots$. The economy is populated by a total of $\bar{N}$ workers. Each worker has a finite career: in every period, each worker exits with probability $\delta \in (0, 1)$ and is replaced by a newborn entrant. A surviving worker discounts the future at rate $\beta \in (0, 1)$, so that the effective discount factor is $\beta(1 - \delta)$.

Workers produce a good of quality A_t^i that depends on how well their actions match an underlying economic fundamental θ_t. Specifically, the quality of worker i's output is

$$A_t^i = \bar{A} - (\theta_t - a_t^i)^2, \tag{1}$$

where a_t^i is the worker's action and $\bar{A}$ is the maximum quality, achieved when the action perfectly matches θ_t.[3] The variable θ_t represents the underlying state of the economy—it can be interpreted as reflecting prevailing market conditions, optimal business strategy, or the best approach to a professional task. It evolves according to

$$\theta_t = \rho\, \theta_{t-1} + \eta_t, \qquad \eta_t \sim N(0, 1/\gamma_\eta), \tag{2}$$

where $\rho \in (0, 1)$ governs persistence and γ_η is the precision of the innovation.

The only friction in the model is informational: θ_t is not directly observable to any participant. Absent this friction, workers would always set $a_t^i = \theta_t$ and achieve the maximum quality $\bar{A}$. Because θ_t is time-varying, even access to an unlimited amount of historical data cannot fully resolve the uncertainty—old observations become less relevant as θ_t drifts, which constrains any

[3] The constant $\bar{A}$ plays no role in any of the results—it drops out of all optimization problems. We retain it only so that output quality has a natural interpretation as a level.

forecasting technology, including AI.

2.2 Pre-AI Economy: Information and Skill Accumulation

We begin with the economy before the arrival of AI. In this benchmark, all workers perform tasks through their own labor, and the only way to learn about θ_t is through direct experience. The cost of labor is normalized to zero; what matters for the AI adoption decision is the relative cost of AI versus labor, which is captured by the AI price p_A introduced below.

Worker i enters period t with *private expertise* $\gamma_t^{H,i} > 0$, which represents the accumulated precision of her beliefs about the fundamental from all past experience. By performing her task, the worker generates a private signal about θ_t:

$$s_t^{L,i} = \theta_t + \varepsilon_t^{L,i}, \qquad \varepsilon_t^{L,i} \sim N\big(0,\, 1/\gamma^l(\gamma_t^{H,i})\big), \tag{3}$$

where $\gamma^l(.) > 0$ is the precision of the labor signal, written as a function of the worker's expertise to reflect that more experienced workers extract richer information from each task. The functional form of $\gamma^l(\cdot)$ is disciplined by the empirical wage–experience profile and specified in Section 4. The worker then chooses her action a_t^i to maximize expected output quality $E[A_t^i \mid \gamma_t^{H,i}, s_t^{L,i}]$. Given the quadratic structure of (1) and the normality of signals, the optimal action is the posterior mean of θ_t given all available information. Let $\mathcal{H}_t^i = \{s_\tau^{L,i}\}_{\tau<t}$ denote the worker's prior information set—the collection of all labor signals she has observed throughout her career. This history determines her private expertise $\gamma_t^{H,i}$, which summarizes the precision of her accumulated beliefs about the fundamental. The optimal action is then

$$a_t^i = E(\theta_t \mid \mathcal{H}_t^i, s_t^{L,i}). \tag{4}$$

The worker's total forecasting precision is $P_t^i = \gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})$, combining what she already knows from past experience with what she learns from today's labor signal. Substituting the optimal action into (1) and taking expectations yields per-period expected output

$$E[A_t^i \mid \mathcal{H}_t^i] = \bar{A} - \frac{1}{P_t^i} = \bar{A} - \frac{1}{\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})}, \tag{5}$$

increasing in expertise.

For the purposes of calibration we interpret the forecasting precision P_t^i as the worker's *labor productivity*—the quality-adjusted output she generates per period—and identify it empirically with labor earnings. The per-period payoff is observationally equivalent to a CRRA utility over

earnings with relative risk aversion equal to two: $u(C) = \bar{A} + C^{(1-\sigma)}/(1-\sigma), \sigma = 2$. A more productive worker has higher precision, hence higher earnings, and hence higher utility. The worker takes no strategic action beyond Bayesian updating, and all the relevant economics is encoded in the law of motion for γ^H that we derive next.

After acting, the worker's expertise updates to reflect the information gained during the period. Because the fundamental is persistent but not permanent (equation 2), the worker's accumulated knowledge partially depreciates as θ_t drifts. The law of motion for private expertise follows from Bayesian updating. By standard results for normal signals, the updating rule is:

$$\frac{1}{\gamma_{t+1}^{H,i}} = \frac{\rho^2}{\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})} + \frac{1}{\gamma_\eta}. \tag{6}$$

The first term on the right-hand side is the variance inherited from the current period's posterior, scaled up by ρ^2 to account for the drift in the fundamental. A higher persistence ρ means past information remains more relevant, so less precision is lost between periods. The second term adds the variance of the new innovation η_{t+1}, which the worker cannot resolve until she observes next period's signals.

Under this law of motion, expertise converges monotonically to a unique steady state $\bar{\kappa}$ that solves

$$\frac{1}{\bar{\kappa}} = \frac{\rho^2}{\bar{\kappa} + \gamma^l(\bar{\kappa})} + \frac{1}{\gamma_\eta}.$$

Existence and uniqueness of $\bar{\kappa}$ are immediate when $\gamma^l(\cdot)$ is weakly increasing (Assumption 1 below), since both sides of the equation are then monotone in $\bar{\kappa}$ with opposite slopes near a fixed point. We refer to $\bar{\kappa}$ as the *full-labor steady state*. A worker who enters at $\gamma^H \geq \gamma_0^H$ converges to $\bar{\kappa}$ over time, so in the pre-AI economy there is a single long-run skill level—every worker, no matter how unskilled at entry, converges to $\bar{\kappa}$. We let $\gamma_0^H \equiv \gamma_\eta(1-\rho^2)$ denote the *no-learning lower bound*—the steady state expertise of a worker who never learns ($\gamma^l \equiv 0$). γ_0^H is the lowest expertise level that arises in equilibrium: it is the entry-level skill of new workers and the lower bound of the support throughout the paper.

In every period, a fraction δ of workers exit and are replaced by newborn entrants whose initial expertise is drawn from a distribution F_0 over the positive reals—this is the distribution of entry-level precision, reflecting the heterogeneity in training, education, or innate ability among new workers entering the labor force. The cross-sectional distribution of expertise, G_t, evolves as

surviving workers accumulate skill and new workers enter. In the pre-AI economy, cross-sectional variation in expertise arises from two sources: differences in career length (time since entry) and differences in initial skill across cohorts.

2.3 Introduction of AI

We now suppose that, at some point in time, an AI technology becomes available to all workers at a per-unit cost $p_A > 0$.[4] Each period, worker i chooses an *AI intensity* $\alpha_t^i \in [0,1]$, representing the share of her task delegated to AI, and pays $\alpha_t^i \cdot p_A$. The remaining share $1 - \alpha_t^i$ is performed through her own labor. When $\alpha_t^i = 0$, the worker operates exactly as in the pre-AI economy; when $\alpha_t^i = 1$, she fully delegates to AI; interior values represent human-AI collaboration, in which the worker combines her own judgment with AI assistance.

We treat p_A as an exogenous model parameter and interpret the AI sector as competitive: p_A is the marginal cost of supplying AI services—for example, the per-token cost of inference—taken as given by both workers and AI providers. Cai (2024) studies an alternative environment with monopolistic AI supply, where the AI provider chooses a price to maximize profit. Here we abstract from the AI provider's pricing problem to focus on the worker-side externality through the data-quality channel.

The AI has access to two sources of information for predicting θ_t. The first is a historical dataset Ω_{t-1}, compiled from all workers' past actions. The contents and construction of this dataset are described in detail in Section 2.5 below. Let Q_t denote the precision of the conditional distribution of θ_t given Ω_{t-1}:

$$Q_t = \frac{1}{\text{Var}(\theta_t \mid \Omega_{t-1})}. \tag{7}$$

This is the *aggregate data quality*: it summarizes how informative the historical record is for predicting the current fundamental. Unlike private expertise $\gamma_t^{H,i}$, which is specific to each worker, Q_t is common to all AI users—any worker who delegates to AI draws on the same collective dataset.

The use of large historical datasets to predict outcomes is central to modern AI applications. Retail firms use past sales data to forecast demand and set prices (Adams et al., 2026), hospitals

[4]Since the cost of labor is normalized to zero, p_A represents the relative cost of AI versus labor. In practice, this could be mapped to the subscription cost of AI tools relative to the opportunity cost of performing a task manually. At the benchmark calibration ($p_A = 0.007$), the cost of full AI delegation equals roughly 0.7% of the expected output loss from forecasting with the prior alone—a modest cost relative to the informational gain AI provides to low-skill workers.

use patient records to aid diagnosis, and financial firms use market histories to guide trading strategies. Because θ_t is time-varying, old data points become less relevant as the fundamental drifts: even with unlimited data, AI cannot perfectly predict θ_t.

The second source of information is AI's own privately generated signal:

$$s_t^{A,i} = \theta_t + \varepsilon_t^{A,i}, \qquad \varepsilon_t^{A,i} \sim N(0, 1/\gamma^a), \tag{8}$$

with precision $\gamma^a > 0$. This captures AI's ability to produce novel information beyond what is already in the historical dataset—for example, through reinforcement learning, exploration of new strategies, or pattern recognition on proprietary inputs. The parameter γ^a governs the quality of this novel information. If $\gamma^a \to 0$, AI converges to a pure prediction tool that relies entirely on past data. If $\gamma^a \to \infty$, AI's action would depend only on its newly generated information and not on historical data. As discussed in the introduction, AI extends its frontier autonomously only in narrow regimes—closed-form games, formal mathematics, protein-structure prediction—where the reward signal is well-defined and frequent enough to support self-play. In most knowledge-work tasks, AI's novel signal is bounded: it can recombine existing patterns from its training data but rarely produces information that is not already implicit in the dataset. We therefore view AI's competitive advantage in our benchmark as coming overwhelmingly from the aggregate dataset rather than from independent information generation.

When a worker uses a mix of labor and AI, her implemented action combines both forecasts. Let $f_t^{L,i}$ denote the worker's own forecast of θ_t based on her private expertise and labor signal, and let $f_t^{A,i}$ denote the AI's forecast based on the historical dataset Ω_{t-1} and the fresh AI signal $s_t^{A,i}$. The worker's implemented action is

$$a_t^i = (1 - \alpha_t^i)\, f_t^{L,i} + \alpha_t^i\, f_t^{A,i}, \tag{9}$$

where

$$f_t^{L,i} = E(\theta_t \mid \mathcal{H}_t^i, s_t^{L,i}), \qquad f_t^{A,i} = E(\theta_t \mid \mathcal{H}_t^i, s_t^{A,i}, \Omega_{t-1}).$$

Here $\mathcal{H}_t^i$ denotes worker i's prior information set (which determines her private expertise $\gamma_t^{H,i}$). The labor forecast $f_t^{L,i}$ combines the worker's accumulated experience with the current labor signal; the AI forecast $f_t^{A,i}$ additionally draws on the historical dataset. At the extremes, $\alpha_t^i = 0$ recovers pure labor and $\alpha_t^i = 1$ recovers full AI delegation.

Notice that the AI forecast $f_t^{A,i}$ conditions on the worker's prior information $\mathcal{H}_t^i$—the AI is per-

sonalized. We adopt this specification as our benchmark because it strengthens AI's competitive advantage (the AI can leverage both historical data *and* the worker's private knowledge), making the skill trap result stronger: even under this best-case scenario for AI, over-adoption and skill depreciation arise.[5]

Private expertise continues to evolve through Bayesian learning, but the rate of learning now depends on how the worker allocates her task. When a worker performs her task through labor, she processes the labor signal $s_t^{L,i}$ directly and accumulates expertise rapidly, as in the pre-AI economy. When she delegates to AI, the AI processes the information on her behalf. Although the AI's forecast may be highly accurate—because it draws on the rich historical dataset Q_t—the worker herself does not directly engage with the underlying signals, and her private expertise grows more slowly.

The law of motion for private expertise generalizes (6):

$$\frac{1}{\gamma_{t+1}^{H,i}} = \frac{\rho^2}{\gamma_t^{H,i} + (1-\alpha_t^i)\,\gamma^l(\gamma_t^{H,i}) + \alpha_t^i\,\tilde{\gamma}^a(\gamma_t^{H,i})} + \frac{1}{\gamma_\eta}, \tag{10}$$

where $\tilde{\gamma}^a(\gamma^H)$ is the *effective AI learning rate*—the precision of the feedback that AI-assisted work contributes to the worker's own skill accumulation. We specify the AI rate as a fixed proportion of the labor rate at every level of expertise:

$$\tilde{\gamma}^a(\gamma^H) = c \cdot \gamma^l(\gamma^H), \qquad c \in (0,1). \tag{11}$$

This proportional specification preserves the relative deskilling penalty across the skill distribution: AI delivers a fraction c of the learning that pure labor would at the same expertise level. We discuss robustness to this functional form in Appendix D.5, where we add an interior learning-synergy term that allows a moderate mix of labor and AI to build skill faster than either pure mode; the headline results survive.

Assumption 1 (Learning Technology)**.** *For all γ^H:*

(i) *$\gamma^l(\gamma^H)$ is weakly increasing in γ^H (more experienced workers extract richer information from each labor draw); and*

(ii) *$\tilde{\gamma}^a(\gamma^H) < \gamma^l(\gamma^H)$ (labor builds expertise faster than AI-assisted work). Under (11) this reduces to*

[5]Under impersonal AI—where the AI produces the same forecast $E(\theta_t \mid s_t^{A,i}, \Omega_{t-1})$ regardless of who uses it—the skill trap is at least as severe, because AI's informational advantage over low-skill workers is smaller and the data contribution from AI actions is even more redundant.

$c < 1$.

Both parts are supported empirically. We calibrate part (ii) using the experimental literature on AI-assisted learning, and part (i) using the concavity of the life-cycle wage profile; we defer detailed discussion to Section 4.

It is important to distinguish between the information used for *current performance* and the information that contributes to *skill accumulation*. AI's forecast draws on the aggregate dataset Q_t, which delivers predictions substantially more accurate than what the worker could produce alone. But the worker's own expertise grows only from the effective learning input in (10), which does not include Q_t: the collective knowledge that improves AI's forecast does not transfer to the worker's private understanding.[6] AI therefore boosts today's performance but can slow tomorrow's skill growth.

Under full labor ($\alpha_t^i = 0$), expertise converges to $\bar{\kappa}$ as in the pre-AI economy. Under full AI ($\alpha_t^i = 1$), it converges to a lower steady state $\underline{\kappa}$ that solves

$$\frac{1}{\underline{\kappa}} = \frac{\rho^2}{\underline{\kappa} + \tilde{\gamma}^a(\underline{\kappa})} + \frac{1}{\gamma_\eta}.$$

Since $\tilde{\gamma}^a(\gamma^H) < \gamma^l(\gamma^H)$ pointwise (Assumption 1), we have $\underline{\kappa} < \bar{\kappa}$: workers who permanently delegate to AI converge to a lower level of expertise. The gap $\bar{\kappa} - \underline{\kappa}$ measures the long-run cost of permanent AI dependence. In the pre-AI economy, this gap does not exist: all workers converge to the same $\bar{\kappa}$. The introduction of AI creates the possibility of a second, lower steady state, which is the foundation for the skill trap.

2.4 Worker's Problem

Each period, worker i faces a dynamic decision that unfolds in three stages:

1. *Allocation.* The worker observes her current expertise $\gamma_t^{H,i}$ and the aggregate data quality Q_t, and chooses her AI intensity $\alpha_t^i \in [0, 1]$.

2. *Action.* She draws her labor and AI signals, forms the two forecasts $f_t^{L,i}$ and $f_t^{A,i}$ as defined in (9), and implements the combined action $a_t^i = (1 - \alpha_t^i) f_t^{L,i} + \alpha_t^i f_t^{A,i}$. Her current payoff is

[6] This reflects the limited interpretability of modern AI systems (Lipton, 2018; Rudin, 2019). Modern AI is largely a black box: a worker who delegates to AI observes its outputs but cannot directly inspect the reasoning that produced them. She therefore cannot recover the underlying knowledge from Q_t simply by watching the AI act, so the AI's reliance on Q_t does not show up in her own learning rate $\tilde{\gamma}^a$.

$\bar{A} - (\theta_t - a_t^i)^2 - \alpha_t^i \, p_A$.

3. *Learning*. Her private expertise updates according to the skill law of motion (10), with the effective learning input depending on her choice of α_t^i.

The worker's choice of α_t^i thus affects both her *current* output quality—a higher α_t^i improves the forecast precision whenever $Q_t + \gamma^a > \gamma^l(\gamma^H)$—and her *future* expertise, since delegating to AI replaces the rich feedback from labor (rate $\gamma^l(\gamma^H)$) with the weaker feedback from AI-assisted work (rate $\tilde{\gamma}^a(\gamma^H)$). Taking the path of aggregate data quality $\{Q_t\}$ as given, the worker solves:

$$V(\gamma^H; Q) = \max_{\alpha \in [0,1]} \left\{ E\left[\bar{A} - (\theta - a)^2\right] - \alpha \, p_A + \beta(1-\delta) \, V(\gamma^{H\prime}; Q') \right\}, \tag{12}$$

where $a = (1-\alpha) f^L + \alpha f^A$, and $\gamma^{H\prime}$ evolves according to (10). The worker takes Q_t as given because she is atomistic: her individual choice has a negligible effect on the aggregate dataset. We derive the explicit form of the flow payoff and characterize the optimal policy in Section 3.

2.5 Data Generation

When a worker takes an action, her behavior is recorded as a data point that enters the aggregate dataset. Specifically, the data point from worker i's action in period t is

$$D_t^i = a_t^i + \varepsilon_t^{D,i}, \qquad \varepsilon_t^{D,i} \sim N(0, \bar{N}/\gamma_D), \tag{13}$$

where $\gamma_D > 0$ is the data recording precision.[7] A key assumption is that only the worker's observable action enters the dataset—her private signals and private expertise are not directly recorded. This is consistent with real-world practice: a firm's business decisions are visible to outsiders, but the private analysis behind those decisions is not.

The aggregate dataset available to AI at the start of period t consists of the complete history of all past recorded actions:

$$\Omega_{t-1} = \{D_{t-j}^i\}_{j \geq 1}^{i=1,\ldots,\bar{N}}. \tag{14}$$

Note that not all data points are equally informative about θ_t—the informativeness of a data point depends on how much of the underlying action reflects genuinely new information versus information that the dataset has already captured from past observations. We characterize this

[7] The noise variance $\bar{N}/\gamma_D$ is scaled by the population size so that, in the $\bar{N} \to \infty$ limit, the law of large numbers delivers a well-defined aggregate precision Φ_t from the cross-section of data points. Each individual data point becomes infinitely noisy, but the aggregate of $\bar{N}$ such points retains finite precision.

precisely in Section 3, where we show that the composition of labor versus AI activities determines the aggregate flow of novel information Φ_t into the dataset, and hence the evolution of aggregate data quality Q_t.

2.6 Competitive Equilibrium

We focus on the limiting economy as $\bar{N} \to \infty$, in which idiosyncratic noise in data points washes out and the evolution of Q_t becomes deterministic given the cross-sectional distribution of choices.

Definition 1. *A* competitive equilibrium *consists of a value function $V(\gamma^H; Q)$, a policy function $\alpha^*(\gamma^H; Q)$, a path for the skill distribution $\{G_t\}_{t=0}^{\infty}$, and a path for aggregate data quality $\{Q_t\}_{t=0}^{\infty}$ such that:*

1. *V and α^* solve the worker's problem (12) given $\{Q_t\}$.*
2. *Aggregate data quality Q_t evolves according to (25), consistently with the cross-sectional distribution G_t and the policy α^*.*
3. *The skill distribution G_t evolves consistently with individual skill dynamics (10), the exit rate δ, and the entry distribution F_0:*

$$G_{t+1}(B) = (1-\delta) \int \mathbf{1}\Big\{ H(\gamma^H, \alpha^*(\gamma^H; Q_t)) \in B \Big\}\, dG_t(\gamma^H) + \delta\, F_0(B), \tag{15}$$

where H denotes the law of motion (10).

Each worker solves (12) taking Q_t as given, but Q_t is itself determined by the distribution of choices across all workers through (25). Because no individual worker internalizes her contribution to Q_t, the equilibrium features over-adoption of AI relative to the social optimum.

3 Characterization

We now characterize the equilibrium of the model. We proceed in three steps. First, we derive the optimal actions and the resulting flow payoffs for workers using labor, AI, and collaborative mixes. Second, we analyze the individual skill dynamics and establish the conditions under which the skill trap and polarization arise. Third, we characterize the information content of each type of action and derive the explicit law of motion for aggregate data quality Q_t.

3.1 Flow Payoffs and Optimal AI Intensity

We now characterize the flow payoffs and the worker's optimal choice of AI intensity α.

The quadratic structure of (1) implies that each optimal forecast is the posterior mean and the expected loss equals the inverse of the posterior precision. The labor forecast $f_t^{L,i}$ defined in (9) has precision $\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})$, combining the worker's accumulated expertise with her fresh labor signal. The AI forecast $f_t^{A,i}$ has precision $\gamma_t^{H,i} + Q_t + \gamma^a$, additionally drawing on the aggregate dataset. When the worker combines these two forecasts with weights $(1-\alpha, \alpha)$, the precision of the combined forecast is additive in its components. This follows because, conditional on the worker's prior information $\mathcal{H}_t^i$, the two fresh signals $s_t^{L,i}$ and $s_t^{A,i}$ are drawn independently—each provides information about θ_t that the other does not contain. Since the expected loss under quadratic payoffs equals the inverse of total precision, the expected flow payoff is:

$$u(\gamma^H, \alpha; Q) = \bar{A} - \frac{1}{\gamma^H + (1-\alpha)\gamma^l(\gamma^H) + \alpha(Q + \gamma^a)}. \tag{16}$$

The denominator is the worker's effective precision: her own expertise γ^H contributes regardless of how she allocates the task, while the remaining precision comes from labor (precision $\gamma^l(\gamma^H)$, weight $1-\alpha$) or AI (precision $Q + \gamma^a$, weight α).

Substituting (16) into the Bellman equation (12), the worker's problem in recursive form is:

$$V(\gamma^H; Q) = \max_{\alpha \in [0,1]} \left\{ \bar{A} - \frac{1}{\gamma^H + (1-\alpha)\gamma^l(\gamma^H) + \alpha(Q + \gamma^a)} - \alpha\, p_A + \beta(1-\delta)\, V(\gamma^{H\prime}; Q') \right\}, \tag{17}$$

where $\gamma^{H\prime}$ evolves according to (10).

In the full dynamic problem, the worker's choice of α balances three forces. The first-order condition for an interior $\alpha^* \in (0,1)$ is:

$$\underbrace{\frac{Q + \gamma^a - \gamma^l(\gamma^H)}{\left[\gamma^H + (1-\alpha)\gamma^l(\gamma^H) + \alpha(Q + \gamma^a)\right]^2}}_{\text{information gain from AI}} = \underbrace{p_A}_{\text{AI cost}} \underbrace{- \beta(1-\delta)\frac{\partial H(\gamma^H, \alpha)}{\partial \alpha} V'(\gamma^{H\prime}; Q')}_{\text{skill loss of AI}}, \tag{18}$$

where $\partial H / \partial \alpha < 0$ because a higher AI intensity substitutes the lower learning rate $\tilde{\gamma}^a(\gamma^H)$ for the labor rate $\gamma^l(\gamma^H)$ in the skill law of motion (10).[8]

The left-hand side is the information gain: the marginal improvement in current output from shifting one more unit of the task to AI. The right-hand side is the total cost, with two components. The first is the direct cost premium p_A. The second is the dynamic skill cost: each unit of α

[8] The value function V is increasing and concave in γ^H—higher expertise strictly improves both current and future payoffs, with diminishing returns—so $V' > 0$ and the skill-loss term is strictly positive whenever $\tilde{\gamma}^a(\gamma^H) < \gamma^l(\gamma^H)$.

substitutes $\tilde{\gamma}^a(\gamma^H)$ for $\gamma^l(\gamma^H)$ in the learning process, reducing the rate of skill accumulation.

The three-way decomposition in (18) explains why the optimal AI intensity $\alpha^*(\gamma^H)$ is decreasing in expertise. A low-skill worker has large forecast errors, so the information gain from AI is large and outweighs the costs. A high-skill worker already forecasts accurately; the marginal information from AI is negligible, and the cost premium and skill cost dominate. The key insight is that expertise and AI are substitutes for current performance but expertise and labor are complements for skill accumulation. This pattern is well documented empirically: Brynjolfsson et al. (2025) find that AI-based tools boost productivity primarily for novice workers; Cui et al. (2025b) report a similar gradient among software developers; and Galdin and Silbert (2025) and Cui et al. (2025a) show that AI writing assistance compresses the quality distribution, disproportionately benefiting low-skill applicants.

3.2 Individual Skill Dynamics and the Skill Trap

We now analyze how the optimal policy feeds back into long-run skill dynamics. The central object is the *skill transition map* H: given a worker's current expertise γ^H and the aggregate data quality Q, the map returns her next-period expertise under the optimal policy:

$$H(\gamma^H; Q) = \left(\frac{\rho^2}{\gamma^H + (1-\alpha^*)\gamma^l(\gamma^H) + \alpha^* \tilde{\gamma}^a(\gamma^H)} + \frac{1}{\gamma_\eta} \right)^{-1}, \tag{19}$$

where $\alpha^* = \alpha^*(\gamma^H; Q)$ is the optimal AI intensity. A crucial distinction is embedded in this equation: the worker's *policy* is shaped by AI's current-period precision $Q + \gamma^a$, which exceeds labor precision $\gamma^l(\gamma^H)$ when the aggregate dataset is rich enough. But the *learning input*—the denominator of the first term in (19)—depends on $\tilde{\gamma}^a(\gamma^H)$, which is strictly lower than $\gamma^l(\gamma^H)$ under Assumption 1. This wedge between what determines AI adoption today and what determines skill growth tomorrow is the source of the skill trap.

Fixed points of H—values γ^H satisfying $H(\gamma^H; Q) = \gamma^H$—are long-run skill levels to which workers converge. The shape of H determines whether the economy has a unique long-run skill level or whether multiple steady states coexist.

To build intuition, we begin with the myopic case $\beta = 0$, where the worker cares only about current output.

Lemma 1 (Optimal AI Intensity)**.** *Suppose $Q + \gamma^a > \gamma^l(\gamma^H)$ for all γ^H in the support. When $\beta =$*

0, *under Assumption 1 and for appropriate values of* p_A*, the optimal AI intensity* $\alpha^*(\gamma^H; Q)$ *is strictly decreasing in expertise* γ^H*, with three regimes:*

(i) Full AI *(*$\alpha^* = 1$*): workers with low expertise delegate entirely to AI.*

(ii) Collaboration *(*$\alpha^* \in (0, 1)$*): workers with intermediate expertise use a mix of AI and labor.*

(iii) Full labor *(*$\alpha^* = 0$*): workers with high expertise rely entirely on their own expertise.*

The proof and closed-form expressions for the policy and the zone boundaries are in Appendix A. The monotone relationship between expertise and AI intensity—less skilled workers rely more on AI, more skilled workers rely more on labor—creates a feedback loop: workers who lack expertise adopt AI, which slows their learning, which keeps their expertise low, which makes AI even more attractive. Whether this feedback is strong enough to generate a genuine trap—with multiple long-run steady states—is the question we turn to next.

Recall from Section 2.3 that permanent full-AI use ($\alpha = 1$) leads to the low steady state $\underline{\kappa}$, while permanent labor ($\alpha = 0$) leads to $\bar{\kappa} > \underline{\kappa}$. Whether the optimal policy produces a genuine trap—where workers at different initial skill levels permanently diverge—depends on whether the substitution feedback documented above is strong enough to make the derivative of the skill transition map exceed one.

The proposition requires $Q + \gamma^a > \gamma^l(\gamma^H)$ at every γ^H in the support: AI must be more productive than labor at every skill level. This condition is not restrictive. The aggregate dataset Q reflects the accumulated knowledge of the entire workforce—in any economy where AI draws on a large body of historical data, Q will be a large number relative to the precision of a single worker's labor signal Q ranges from 20 to 50 while $\gamma^l(\gamma^H)$ ranges from 1 to about 5.

Proposition 1 (Existence of the Skill Trap)**.** *Suppose* $\beta = 0$ *and Assumption 1 hold. Then there exist values of* $Q > 0$ *and* $p_A > 0$ *such that the derivative of the skill transition map satisfies* $H'(\gamma^H; Q) > 1$ *for some* γ^H*. The map* H *therefore has at least two stable fixed points: a low steady state near* $\underline{\kappa}$ *with high AI intensity, and a high steady state near* $\bar{\kappa}$ *with low AI intensity.*

The proof is in Appendix A.

In the pre-AI economy, the derivative of the skill transition map H' is everywhere below 1: every worker converges to the same $\bar{\kappa}$, and no trap can form. The skill trap requires $H' > 1$

somewhere. Decomposing H' reveals the economic force that makes this possible. Let

$$L(\gamma^H, \alpha) \equiv \gamma^H + (1-\alpha)\,\gamma^l(\gamma^H) + \alpha\,\tilde{\gamma}^a(\gamma^H)$$

denote the *learning input* that enters the skill law of motion (10)—both labor and AI learning rates are functions of the worker's own expertise. On the collaboration zone, where the worker uses a mix of AI and labor,

$$H'(\gamma^H) = \underbrace{\frac{\rho^2\gamma_\eta^2}{\left(\rho^2\gamma_\eta + L(\gamma^H, \alpha^*)\right)^2}}_{\text{Kalman gain } (<1)} \times \underbrace{L'(\gamma^H)}_{\text{amplification } (>1)}, \tag{20}$$

where $\alpha^* = \alpha^*(\gamma^H; Q)$ is the optimal AI intensity from (17) and the prime on L' denotes the total derivative along the policy. The first factor is the standard Kalman gain—the same attenuation present in the pre-AI economy. The second factor satisfies

$$L'(\gamma^H) \;\geq\; 1 + \frac{\gamma^l(\gamma^H) - \tilde{\gamma}^a(\gamma^H)}{Q + \gamma^a - \gamma^l(\gamma^H)} \;>\; 1, \tag{21}$$

with equality when $\gamma^l(\cdot)$ is constant. The substitution term captures *endogenous AI-to-labor substitution*: when a worker gains a marginal unit of expertise, she reduces her AI intensity (Lemma 1), replacing AI-assisted learning at rate $\tilde{\gamma}^a(\gamma^H)$ with the faster labor learning at rate $\gamma^l(\gamma^H)$. The gap $\gamma^l(\gamma^H) - \tilde{\gamma}^a(\gamma^H)$ in the numerator is the deskilling wedge; the denominator $Q + \gamma^a - \gamma^l(\gamma^H)$ is the informational advantage of AI over labor at expertise γ^H. Whenever Assumption 1 holds with strict inequality (i.e., the labor-learning profile rises with expertise), there is an additional non-negative *profile* contribution to L' from the rising $\gamma^l(\gamma^H)$, which strengthens the amplification further; the proof in Appendix A writes out the exact expression.

The product $H' > 1$ when the amplification overcomes the Kalman attenuation. This happens when AI is informationally superior to labor ($Q + \gamma^a > \gamma^l(\gamma^H)$) but not overwhelmingly so—when the denominator of L' is moderate, the amplification is large. For any Q in the range where AI is better than labor but the gap is not too wide, there exists a price interval for p_A such that the collaboration zone is non-empty and $H' > 1$ on part of it (Appendix A). The resulting map H has at least two stable fixed points: a low one in the full-AI zone and a high one where the worker uses less AI. The high steady state need not involve full labor—in the calibrated economy, even the expert fixed point has interior AI intensity, reflecting the fact that skilled workers still find some AI use beneficial when the aggregate dataset is rich enough.

Figure 1 illustrates Proposition 1. In the full-AI zone (low γ^H), the worker delegates entirely

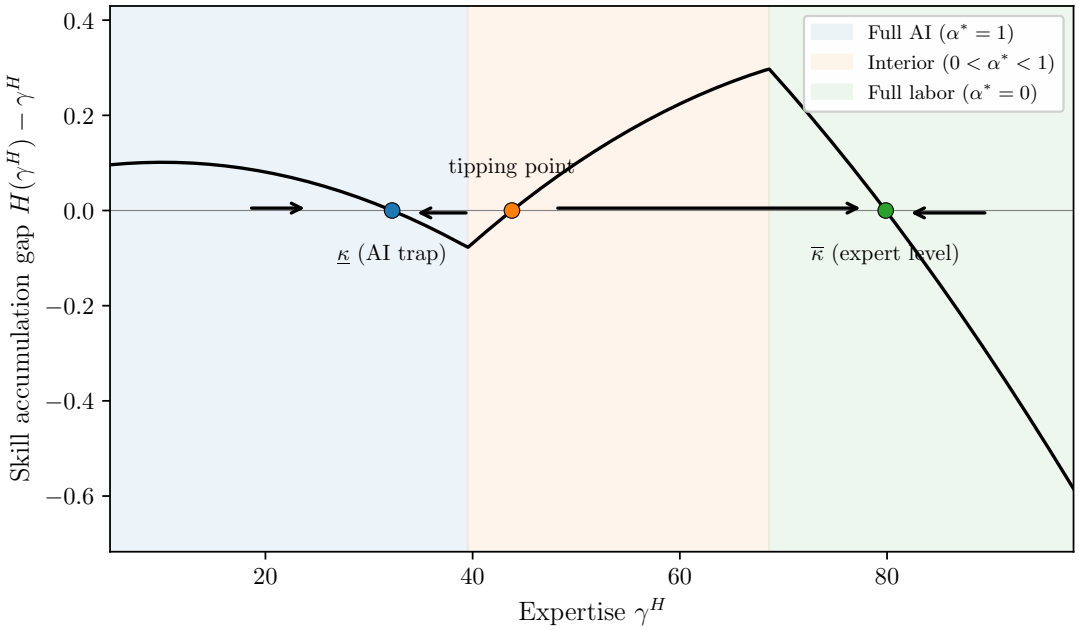

Figure 1: The skill trap under the myopic policy ($\beta = 0$), illustrative parameters. The plot shows the expected one-step skill change $H(\gamma^H) - \gamma^H$ as a function of current expertise γ^H, with the three regimes of the optimal policy (Lemma 1) shaded: full AI ($\alpha^* = 1$), interior collaboration ($0 < \alpha^* < 1$), and full labor ($\alpha^* = 0$). The zero crossings are the fixed points of H: a stable AI trap at $\underline{\kappa}$ (labeled κ_A in the figure), an unstable tipping point in the collaboration zone (labeled κ_M), and the full-labor steady state at $\bar{\kappa}$ (labeled κ_L). The arrows indicate the direction of skill dynamics: workers who start below the tipping point converge to the AI trap, while those above converge to the expert level.

to AI. Skill accumulates slowly, and above $\underline{\kappa}$ the depreciation from the fundamental's drift overwhelms the meager learning: the worker is trapped. In the collaboration zone (intermediate γ^H), the substitution from AI to labor has a compounding effect: each unit of labor that replaces AI raises the learning input by $\gamma^l(\gamma^H) - \tilde{\gamma}^a(\gamma^H) = (1 - c)\gamma^l(\gamma^H)$, which accelerates skill growth, which pushes expertise higher, which further reduces α^*. This positive feedback loop—driven by $H' > 1$—carries the worker through the unstable region and toward the high steady state, where the combination of high expertise and moderate AI use is self-sustaining.

3.3 Data Quality and the Externality from AI Adoption

The analysis so far has taken the aggregate data quality Q as given. But Q is itself determined by workers' collective actions: each period, the dataset is updated with the recorded actions of all workers, and the informational content of these recordings depends on how much each worker

relied on AI versus labor. We now derive how AI adoption shapes Φ_t—the aggregate flow of novel information into the dataset—and show that higher AI intensity systematically degrades future data quality, creating a negative externality that individual workers do not internalize.

To characterize the information content of workers' actions, we adopt the following maintained assumption:

Assumption 2 (Signal-Only Data Contribution)**.** *The dataset updates only from the current-period signal components of workers' actions.*

Under this assumption, a worker's prior $E(\theta_t \mid \mathcal{H}_t^i)$—which is a function of her entire history of past actions—is treated as already known to the dataset from previous recordings. This greatly simplifies the analysis by making the information content of each action transparent. Without the assumption, the dataset would partially re-learn from the prior component of each worker's action in every period, introducing a layer of repeated inference that complicates the aggregation but does not change the overall mechanism.

To see the content of the assumption, consider the decomposition of each forecast type. The labor forecast separates into a prior component and a fresh signal:

$$f_t^{L,i} = \underbrace{\frac{\gamma_t^{H,i}}{\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})} E(\theta_t \mid \mathcal{H}_t^i)}_{\text{prior (assumed known)}} + \underbrace{\frac{\gamma^l(\gamma_t^{H,i})}{\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})} s_t^{L,i}}_{\text{innovation (new)}}. \tag{22}$$

The AI forecast has an additional component drawn from the aggregate dataset Ω_{t-1}:

$$f_t^{A,i} = \underbrace{\frac{\gamma_t^{H,i}}{\gamma_t^{H,i} + Q_t + \gamma^a} E(\theta_t \mid \mathcal{H}_t^i)}_{\text{prior (assumed known)}} + \underbrace{\frac{Q_t}{\gamma_t^{H,i} + Q_t + \gamma^a} E(\theta_t \mid \Omega_{t-1})}_{\text{dataset echo (known)}} + \underbrace{\frac{\gamma^a}{\gamma_t^{H,i} + Q_t + \gamma^a} s_t^{A,i}}_{\text{innovation (new)}}. \tag{23}$$

When γ_D is finite, the dataset cannot perfectly reconstruct workers' past actions, so the prior component $E(\theta_t \mid \mathcal{H}_t^i)$ embedded in today's action contains some information the dataset has not fully absorbed. This gives value to repeated learning—extracting residual information from the prior-driven part of each action. As γ_D grows large, the dataset's recordings become increasingly precise, and the prior component becomes redundant: the dataset already knows what the worker knows from her history. Assumption 2 takes this limiting case as exact.

Two considerations support this simplification. The first is that the approximation becomes exact in the limit $\gamma_D \to \infty$, and the higher γ_D is, the more accurate it becomes. This is reassuring because Section 6.3 shows that the welfare gains from planner intervention are *larger*, not smaller,

at higher values of γ_D: the regime where Assumption 2 is tightest is also the regime where the data externality is most consequential, so the approximation is sharpest exactly where it matters most. The second is that we can relax Assumption 2 directly and compute the full Bayesian counterpart as a robustness check; Appendix E does this and shows that the skill trap, the planner's monotone-declining $\alpha^*(\gamma^H)$, and large welfare gains all survive under the exact treatment. At the baseline calibration the exact Bayesian planner gain is in fact slightly *larger* than the signal-only gain, because the exact Bayesian filter has a second information channel—joint filtering of cohort action dynamics—that is ignored in the competitive equilibrium. Internalizing the externality is therefore *more* valuable.

We maintain this assumption so that we can get the following intuitive expression for the aggregate data quality evolution:

Proposition 2 (Aggregate Data Quality)**.** *Under Assumption 2, in the limit $\bar{N} \to \infty$, the aggregate precision gain from period t's data is*

$$\Phi_t = \gamma_D \int \left[(1 - \alpha_t^i) \frac{\gamma^l(\gamma_t^{H,i})}{\gamma_t^{H,i} + \gamma^l(\gamma_t^{H,i})} + \alpha_t^i \frac{\gamma^a}{\gamma_t^{H,i} + Q_t + \gamma^a} \right]^2 dG_t, \tag{24}$$

and the law of motion for aggregate data quality is

$$\frac{1}{Q_{t+1}} = \frac{\rho^2}{Q_t + \Phi_t} + \frac{1}{\gamma_\eta}. \tag{25}$$

The law of motion (25) has the same Kalman filter structure as the individual skill recursion (10): the dataset starts the period with precision Q_t, gains precision Φ_t from the new data, and then loses some precision due to the drift in the underlying state θ_t. The key economic content is in the data flow Φ_t.

To see why AI adoption degrades data quality, consider the innovation weights in (24). A worker with expertise γ^H who uses labor contributes information with innovation weight $\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))$. A worker who uses AI contributes information with innovation weight $\gamma^a/(\gamma^H + Q + \gamma^a)$. The AI innovation weight is systematically smaller for two reasons. First, the AI's fresh signal precision γ^a is typically smaller than the labor signal precision $\gamma^l(\gamma^H)$ (the AI learns less from each interaction). Second, the AI's denominator includes the aggregate data quality Q, reflecting the fact that most of the AI's forecast comes from recycling what the dataset already knows—the "dataset echo" in equation (23)—rather than generating genuinely new information. When a

worker delegates to AI, her recorded action is mostly a noisy copy of the existing dataset, contributing little that the dataset did not already contain. The following result formalizes this.

Proposition 3 (Data Corruption from AI Adoption)**.** *Suppose that AI is more productive than labor ($Q + \gamma^a > \gamma^l(\gamma^H)$) but generates less novel information ($\gamma^a < \gamma^l(\gamma^H)$) at every γ^H in the support. Holding the cross-sectional distribution of expertise G_t fixed, the aggregate data flow Φ_t is strictly decreasing in each worker's AI intensity α_t^i. Higher AI intensity today degrades data quality Q_{t+1} tomorrow.*

The proof is in Appendix A. The key intuition is that a worker who delegates to AI contributes an action whose innovation weight—$\gamma^a/(\gamma^H + Q + \gamma^a)$—is smaller than the innovation weight under labor—$\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))$—because the AI's denominator includes the aggregate data quality Q. Most of the AI's forecast comes from recycling what the dataset already knows rather than generating genuinely new information. Switching from labor to AI therefore replaces a data-rich action with a data-poor one, reducing Φ_t and hence future Q.

Proposition 3 identifies the data externality at the heart of the model. Each worker takes Q as given when choosing her AI intensity, but her choice of α affects Φ_t, which affects Q_{t+1}, which affects every other worker's AI quality in the next period. The externality is negative because the AI innovation weight $\gamma^a/(\gamma^H + Q + \gamma^a)$ is smaller than the labor weight $\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))$. The proposition writes this out under the sufficient condition $\gamma^a < \gamma^l(\gamma^H)$, but the qualitative result is more robust than this condition suggests: the AI denominator carries Q, which is typically much larger than $\gamma^l(\gamma^H)$ in any economy where the AI draws on a substantial historical dataset. Even if AI's fresh signal precision matched or exceeded labor's, the dataset's existing precision would dilute it—an AI-assisted action remains mostly an echo of what the dataset already knows, and replacing labor with AI still records a less informative action. The externality therefore survives across a wide range of γ^a, including values that violate $\gamma^a < \gamma^l(\gamma^H)$ (Appendix D.4). Individual workers do not internalize this cost because Q is a public good—each worker's contribution to the dataset is infinitesimal, but the aggregate effect of widespread AI adoption can be large.

The preceding analysis has characterized the individual skill dynamics (Proposition 1) and the data quality externality (Propositions 2–3) as separate channels. In the competitive equilibrium, these channels interact: the cross-sectional distribution of AI intensity determines data quality through Φ_t, data quality shapes AI's informational advantage and hence the optimal pol-

icy through Q, and the optimal policy feeds back into skill dynamics through the transition map H. With forward-looking agents ($\beta > 0$) and endogenous data quality, closed-form characterization of the full equilibrium is no longer tractable. We therefore turn to a quantitative analysis. The next section calibrates the model to match pre-AI skill dynamics and the experimental evidence on AI-assisted learning, and the subsequent sections evaluate the model's implications for skill polarization, welfare, and optimal policy.

4 Calibration

We calibrate the model at an annual frequency. The strategy is to pin the skill dynamics from the pre-AI economy—where the life-cycle profile of expertise is well-documented—and then calibrate the AI-related parameters from the experimental literature on AI and learning.

4.1 Pre-AI skill dynamics

The labor signal precision is allowed to depend linearly on the worker's expertise, $\gamma^l(\gamma^H) = \gamma_0^l + \xi(\gamma^H - \gamma_0^H)$, with the entry-level precision normalized to $\gamma_0^l = 1$ and $\gamma_0^H = \gamma_\eta(1-\rho^2)$ the no-learning fixed point. The parameters (ρ, γ_η, ξ) jointly govern how fast the fundamental drifts, how rapidly workers accumulate expertise, and how the labor signal becomes more informative with experience. A worker's total forecasting precision under labor is $P = \gamma^H + \gamma^l(\gamma^H)$, and her expected output loss is $1/P$. Wage growth over the career thus tracks the growth of P.

We calibrate (ρ, γ_η, ξ) to the US experience-wage profile of Lagakos et al. (2018). The calibration yields $\rho = 0.946$, $\gamma_\eta = 128.21$, $\xi = 0.239$, and Figure 2 displays the fit. Implied entry expertise is $\gamma_0^H = 13.5$ and the full-labor steady state is $\bar{\kappa} = 29.6$.

Note that the calibrated $\gamma^l(\gamma^H)$ is increasing in γ^H, so assumption 1(i) holds in the data. This may seem at odds with the concavity of the wage profile—if more experienced workers extract richer information from each task, why doesn't wage growth accelerate at older ages? The model reconciles the two through Bayesian attenuation: as γ^H approaches the innovation-precision ceiling $\gamma_\eta = 128.2$, the marginal value of each additional signal shrinks regardless of its precision. Wages therefore rise rapidly in early career, when each labor signal moves expertise substantially against a low base, and plateau at older ages, when the worker is close to her information ceiling—

even though $\gamma^l(\gamma^H)$ itself is still growing. The empirical concavity of the wage–experience profile is thus consistent with rising γ^l, not evidence against it.

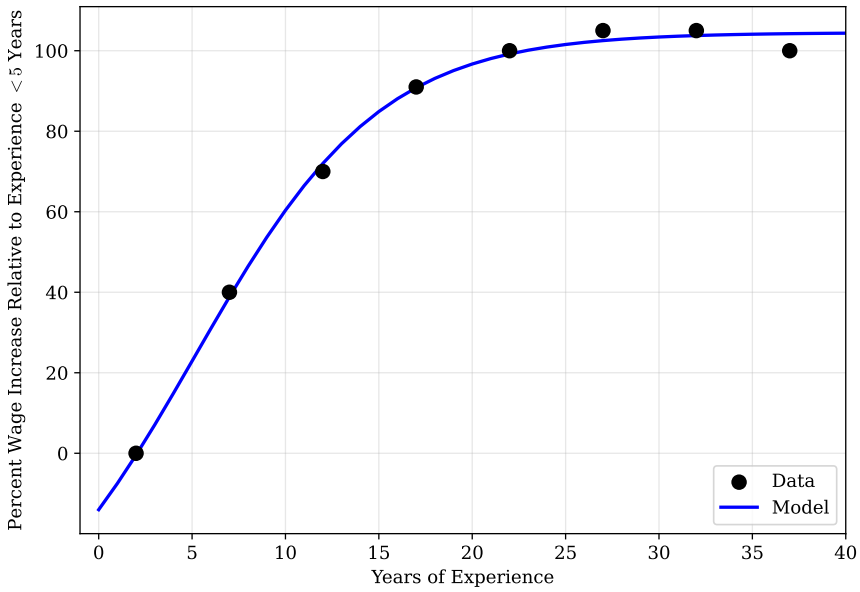

Figure 2: Pre-AI wage profile: data and model fit.

4.2 Demographics

$\beta = 0.96$ is the standard annual discount factor. We set $\delta = 1/30 \approx 0.033$, implying an expected career of $1/\delta = 30$ years in the expertise-building phase, consistent with the horizon over which the cognitive wage profile remains informative (Lagakos et al., 2018 Figure 10 reports profiles through $\tau = 37$). The exit rate δ should not be interpreted literally as retirement. Rather, it captures the effective rate at which a worker's accumulated occupation-specific human capital becomes less relevant—whether through career transitions, promotions into roles that require different skills, or technological changes that reshape the nature of work within an occupation. Of course, not all job changes reset expertise: workers who switch employers within the same occupation carry over much of their experience. But the fraction of human capital that is truly occupation-specific, and that must be rebuilt after a major career transition, is substantial (Kambourov and Manovskii, 2009).

4.3 AI-related parameters

We set $\gamma^a = 0.08$, implying that AI's own novel signal is small relative to an entry-level worker's labor signal: $\gamma^a/\gamma^l(\gamma_0^H) = 0.08$, where $\gamma^l(\gamma_0^H) = 1.0$ is the labor signal precision at entry. Because $\gamma^l(\gamma^H)$ rises with expertise (Assumption 1), the ratio $\gamma^a/\gamma^l(\gamma^H)$ is even smaller for more experienced workers; in steady state, AI's fresh signal is negligible compared to the labor signal of any worker who has accumulated substantial expertise. AI's advantage comes overwhelmingly from the aggregate dataset Q, not from independent information generation. The low value of γ^a is motivated by the model collapse literature: when AI systems are recursively trained on their own outputs, the quality of successive generations degrades (Shumailov et al., 2024), revealing that AI-generated content contains far less genuinely novel information than human-produced data. Cai (2024) formalizes this insight in an economic model with social learning, showing that the rate of data quality degradation under recursive AI training identifies the ratio of novel to recycled information in AI output—precisely γ^a in our model. We treat γ^a as a fixed property of the AI system, independent of the user's expertise.

We set the AI/labor learning ratio $c = \tilde{\gamma}^a(\gamma^H)/\gamma^l(\gamma^H) = 0.25$, a single scalar that disciplines AI's learning contribution at every level of expertise under (11). We calibrate this ratio from the experimental literature on AI-assisted learning. Several randomized studies measure how much workers learn when they delegate to AI versus performing tasks independently. We convert each into an annualized learning ratio $\tilde{\gamma}^a/\gamma^l$. Bastani et al. (2025) find that high school math students with unrestricted ChatGPT access learn 83% as much as a control group over one month of practice ($0.83^{12} \approx 0.11$ annualized). Shen and Tamkin (2026) find that software developers who fully delegated coding to AI learned roughly 49% as much as the control over a one-hour session—the AI group encountered a median of one error compared to three for the control, receiving far less of the hands-on diagnostic practice that builds expertise—which annualizes to near zero. Barcaui (2025) find that undergraduate students using ChatGPT retained 84% of the learning of a control group over 4.5 months including a retention test ($0.84^{2.7} \approx 0.63$ annualized). The annualized estimates range from near zero to 0.63, with the lowest values consistently associated with full delegation rather than guided use.

Evidence from clinical medicine corroborates these findings. Budzyń et al. (2025) find that

endoscopists whose colonoscopies were randomly assigned to use AI assistance roughly half the time saw their *unassisted* detection rate fall from 28.4% to 22.4%—a 21% decline—after just three months. Interpreting detection rates as a proxy for expertise, we can map this finding to the model. Under half AI use ($\alpha = 0.5$), the model's long-run expertise decline is 16% at $\tilde{\gamma}^a/\gamma^l = 0.25$ and 20% at $\tilde{\gamma}^a/\gamma^l = 0.10$. The observed 21% decline after a single quarter already exceeds the long-run prediction at $\tilde{\gamma}^a/\gamma^l = 0.25$, pointing to a much lower value.

We adopt $c = 0.25$ as our benchmark, so $\tilde{\gamma}^a(\gamma^H) = 0.25\,\gamma^l(\gamma^H)$ at every γ^H. The direct experimental estimates from education and software development cluster between 0 and 0.25 when workers fully delegate, and the medical evidence is consistent with a value in this range.

The proportional specification $\tilde{\gamma}^a(\gamma^H) = c\,\gamma^l(\gamma^H)$ also rules out an interior learning peak—a regime in which a moderate mix of labor and AI builds skill faster than either pure mode. The empirical case for an interior peak is weak (the studies cited above find AI access uniformly impairs independent learning), but a smaller "scaffolded use" literature suggests guided AI could complement learning. In Bastani et al. (2025), adding guardrails that prevented ChatGPT from giving direct answers eliminated the learning impairment entirely. The effective $\tilde{\gamma}^a$ thus depends on AI design and intensity of use.

The data recording precision γ_D governs how informative each recorded action is for the dataset. We calibrate it so that the pre-AI aggregate data quality Q^*_{pre}—defined as the steady-state fixed point of the data quality recursion when all workers use full labor—equals the total forecasting precision of the best human expert:

$$Q^*_{\text{pre}} \;=\; \bar{\kappa} + \gamma^l(\bar{\kappa}) \;=\; 29.6 + 4.85 \;=\; 34.5.$$

The interpretation is that, before AI is available, the most informative posterior any worker achieves in equilibrium combines her accumulated expertise $\bar{\kappa}$ with the precision of one period's labor signal $\gamma^l(\bar{\kappa})$. Because the dataset aggregates many such posteriors across the cross-section, Q^*_{pre} can match—and in calibrations where workers are abundant, even slightly exceed—this benchmark, while remaining well below the upper bound $\gamma_\eta = 128$ set by the innovation precision. This calibration target yields $\gamma_D = 390$.[9]

[9] In ongoing companion work, I calibrate γ_D directly to AI benchmark performance scores from the test-based-tasks literature, providing an external discipline that does not rely on the best-expert anchor used here. The qualitative findings of this paper are robust to the alternative calibration; quantitative magnitudes shift modestly within the range reported in the γ_D comparative statics of Section 6.3.

Table 1: Benchmark calibration. Labor signal precision specified as $\gamma^l(\gamma^H) = \gamma_0^l + \xi\,(\gamma^H - \gamma_0^H)$, with $\gamma_0^H = \gamma_\eta(1-\rho^2)$ the no-learning fixed point.

Panel A: Pre-AI economy			
AR(1) persistence	ρ	0.946	
Innovation precision	γ_η	128.2	Pre-AI wage–experience profile
Labor signal intercept	γ_0^l	1.0	
Labor signal slope	ξ	0.239	
Discount factor	β	0.96	Standard
Exit rate	δ	1/30	30-year expertise-building career
Data recording precision	γ_D	390	AI quality equals best human expert
Panel B: AI technology			
Fresh AI signal precision	γ^a	0.08	Model collapse literature
AI/labor learning ratio	c	0.25	AI-assisted learning experiments
AI price	p_A	0.007	

We model the arrival of AI as a price shock: the cost of AI drops from prohibitively high (the pre-AI economy) to $p_A = 0.007$. At this price, AI is cheap enough that workers with low expertise find it attractive—the information gain exceeds the cost—but not so cheap that every worker adopts AI regardless of skill. The benchmark $p_A = 0.007$ lies in the intermediate region where the interaction between AI adoption, skill accumulation, and data quality produces meaningful heterogeneity.

4.4 Entry distributions

We consider two benchmark entry distributions:

- **Uniform entry:** F_0 is uniform on $[\gamma_0^H, \bar{\kappa}]$. New workers arrive with heterogeneous initial skill, representing an economy with diverse educational backgrounds.
- **Bottom entry:** F_0 is concentrated at γ_0^H. All new workers enter at the lowest skill level, as is natural for recent graduates entering knowledge-intensive occupations.

5 Quantitative Results

We now solve for the competitive equilibrium following the AI price shock. At $t = 0$, the cost of AI drops from prohibitively high to $p_A = 0.007$, and the economy transitions from the pre-AI steady state to a new post-AI equilibrium.

5.1 Steady state

Figure 3 summarizes the competitive equilibrium for both entry specifications.

Uniform entry (left column). The H map exhibits the three-fixed-point structure of Proposition 1: a low-skill trap at $\gamma^H \approx 16.4$ (where $\alpha^* = 1$), an unstable interior fixed point at $\gamma^H \approx 25.5$, and the full-labor steady state at $\bar{\kappa} = 29.6$ (where $\alpha^* = 0$). The post-AI skill distribution is bimodal (panel b): workers who enter near the trap remain there with full AI use, while those entering above the unstable threshold accumulate skill under labor and reach $\bar{\kappa}$. Mean AI intensity is $\bar{\alpha} = 0.73$, mean expertise falls from 25.5 (pre-AI) to 20.3 (post-AI), and aggregate data quality drops from $Q^*_{\text{pre}} = 34.5$ to $Q^* = 22.6$.

Bottom entry (right column). Every new worker enters at $\gamma_0^H = 13.5$ and adopts near-full AI ($\bar{\alpha} = 0.88$). Since the entry point lies below the unstable interior fixed point, the bulk of mass is absorbed into the low-skill trap at $\gamma^H \approx 16.4$, with a smaller mass of long-tenure workers reaching the expert level at $\bar{\kappa} = 29.6$ (panel b). The post-AI distribution is therefore bimodal under bottom entry as well as under uniform entry, but with much greater weight at the trap. Mean expertise falls from 25.4 (pre-AI) to 16.8 (post-AI), and aggregate data quality drops to $Q^* = 18.0$, about half of the pre-AI level.

5.2 Welfare

The welfare effects of AI depend on how welfare is measured. Table 2 reports four measures. The arrival of AI produces two competing forces, both visible in Figure 3. First, a positive *technology effect*: at every skill level, a worker can use AI to boost her current output—panel (d) shows that workers optimally adopt AI, especially at low skill, precisely because it improves their immediate performance. Second, a negative *distribution effect*: AI adoption slows skill accumulation, shifting the long-run skill distribution toward lower expertise—panel (b) shows the post-AI distribution

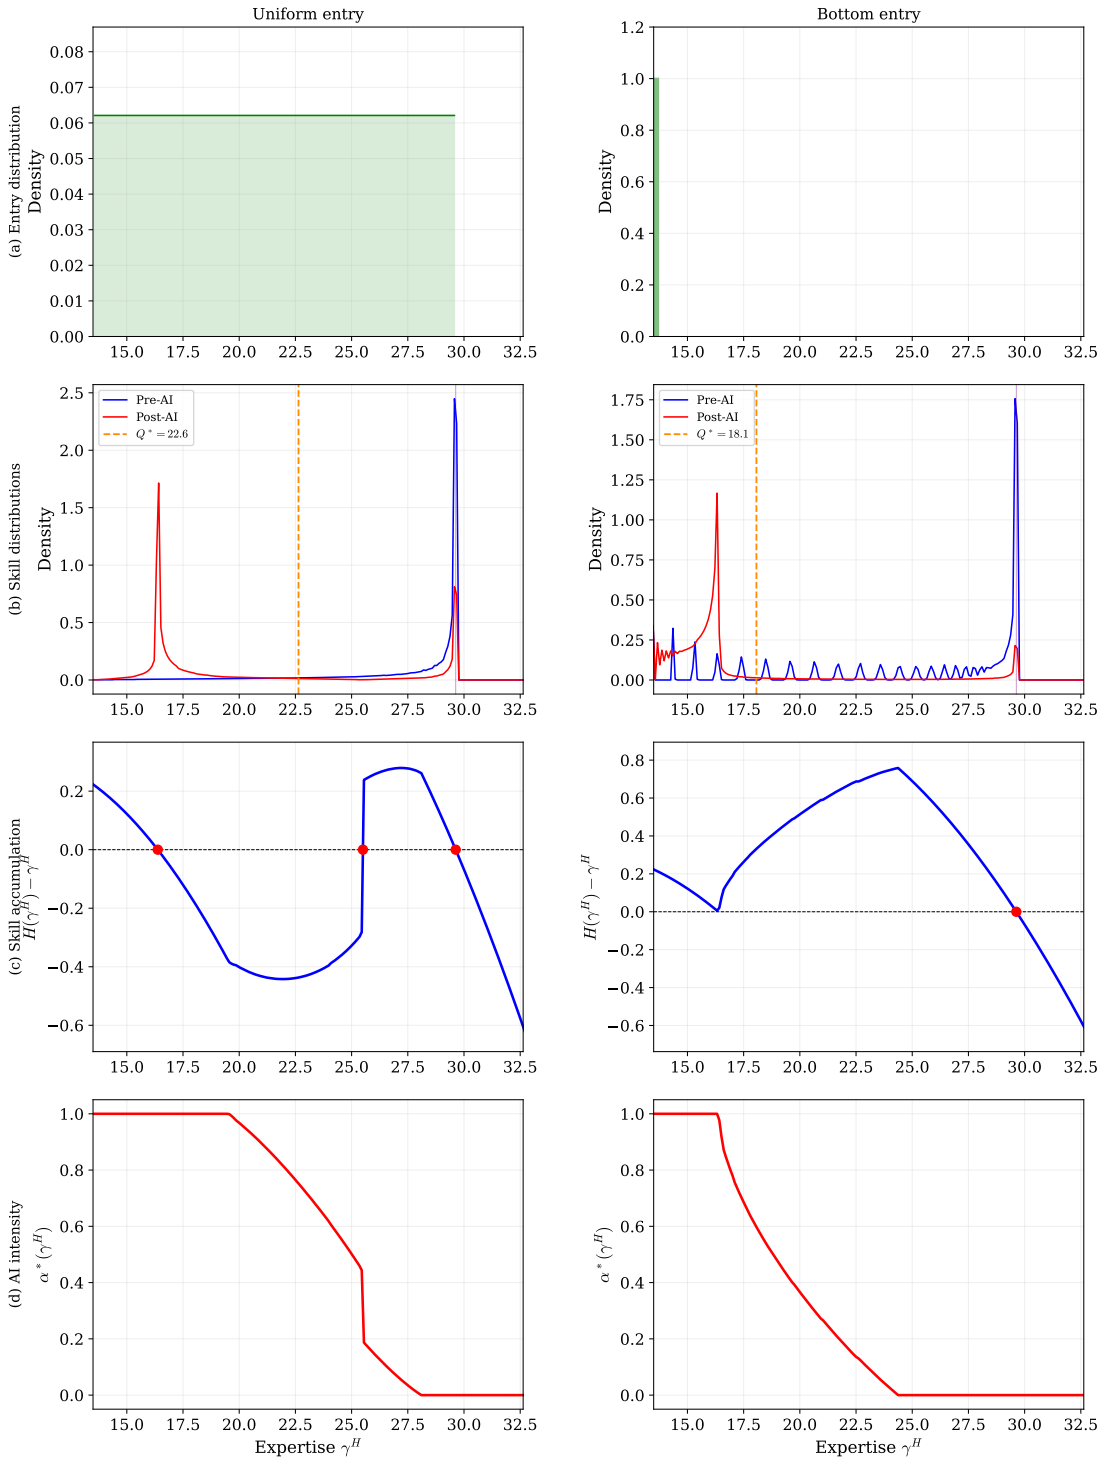

Figure 3: Competitive equilibrium steady state, $p_A = 0.007$. Left column: uniform entry. Right column: bottom entry. Rows show (a) entry distribution F_0, (b) pre-AI and post-AI skill distributions, (c) skill accumulation $H(\gamma^H) - \gamma^H$ with fixed points marked, and (d) optimal AI intensity $\alpha^*(\gamma^H)$.

concentrated at lower γ^H than the pre-AI distribution, with a pronounced mass of workers stuck in the slow zone near $\gamma^H \approx 30$ (uniform) or near the entry level (bottom). Whether the net welfare effect is positive or negative depends on which force dominates, and this in turn depends on the welfare measure used.

Table 2: Welfare effects of AI: pre-AI vs. post-AI

	Uniform entry	Bottom entry
Steady-state comparison		
Cross-sect. lifetime ($\int V\, dG$)	-5.8%	-12.9%
Cross-sect. flow ($\int u\, dG$)	$+0.3\%$	$+1.1\%$
New entrant lifetime ($\int V\, dF_0$)	$+7.0\%$	$+13.8\%$

Note: All entries are percentage changes from pre-AI to post-AI. Cross-sectional measures integrate over the respective steady-state distribution G. New entrant welfare integrates V over the entry distribution F_0. Initial generation welfare evaluates the post-AI transition value function V_0 at the pre-AI distribution G_0.

Note that the cross-sectional lifetime welfare, $\int V\, dG$, falls in the post-AI steady state (-5.8% under uniform entry, -12.9% under bottom entry), while the cross-sectional flow welfare, $\int u\, dG$, moves only modestly and stays slightly positive in both cases ($+0.3\%$ uniform, $+1.1\%$ bottom). Both measures reflect a positive *technology effect* (AI improves the flow payoff at every skill level whenever AI is used) and a negative *distribution effect* (the skill distribution shifts toward lower expertise). The sign of the net change differs across measures because the lifetime value V is much steeper in expertise than the flow payoff u. In the post-AI economy, the flow payoff $u(\gamma^H, \alpha^*(\gamma^H); Q) = -1/[\gamma^H + (1-\alpha^*)\gamma^l(\gamma^H) + \alpha^*(Q+\gamma^a)]$ increases only moderately with skill, while the lifetime value $V(\gamma^H)$ increases steeply: a low-skill worker is trapped at low skill for the remainder of her career, accumulating mediocre payoffs period after period, while a high-skill worker maintains high skill throughout. The career trajectory amplifies skill differences, making V much more sensitive to the position in the skill distribution than u. When AI shifts mass toward low expertise, the welfare cost is therefore much larger under V—large enough to turn the aggregate change clearly negative—while the flow measure barely moves.

New entrant lifetime welfare rises ($+7.0\%$ uniform, $+13.8\%$ bottom), because this measure holds the distribution fixed at F_0 and captures only the technology effect.

5.3 Transition dynamics

We compute perfect-foresight transitions from the pre-AI steady state to the post-AI steady state. At $t = 0$, the AI price shock arrives. Workers and the aggregate dataset adjust gradually as new cohorts enter and old cohorts retire. Figure 4 summarizes the results for both entry specifications.

Uniform entry (left column). Starting from $Q_0 = Q_{\text{pre}} = 34.5$, aggregate data quality declines monotonically to its post-AI level of 22.6 as workers adopt AI (panel c). Mean expertise drops from 25.5 to 20.3 (panel a), and mean AI intensity rises quickly before settling at $\bar{\alpha} = 0.73$ (panel b). The skill distribution (panel d) develops a persistent bimodal shape: experienced workers who entered before AI retain high expertise, while newer cohorts accumulate less skill due to heavy AI use.

Bottom entry (right column). When all entrants start at the lowest skill level, the transition is starker. AI intensity jumps immediately to $\bar{\alpha} \approx 0.88$ and remains high throughout (panel b). Mean skill drops steadily from 25.4 to 16.8 (panel a), and data quality collapses from 34.5 to 18.0 (panel c) as the pre-AI cohort retires and is replaced by AI-dependent entrants who contribute almost no novel information to the dataset. The skill distribution converges to a tight mass at low expertise (panel d). Compared to uniform entry, the bottom-entry economy suffers a larger decline in both skill and data quality, because the composition of entrants leaves no high-skill cohort to anchor the distribution.

Under both entry specifications, the bulk of the transition occurs within the first 50–100 periods, but full convergence takes longer. The speed of the transition is governed by three forces. First, worker turnover: with $\delta = 1/30$, the expected career is 30 years. Second, slow skill accumulation under AI: with $\tilde{\gamma}^a(\gamma^H) = 0.25\,\gamma^l(\gamma^H)$, workers using AI accumulate expertise at a constant fraction of the labor rate, so the absolute deskilling penalty $\gamma^l(\gamma^H) - \tilde{\gamma}^a(\gamma^H) = 0.75\,\gamma^l(\gamma^H)$ widens with experience as $\gamma^l(\gamma^H)$ rises. Third, the data quality feedback loop: as the skill distribution shifts, data quality Q changes, which alters optimal AI intensity, which further shifts the distribution. This feedback amplifies the transition but also slows convergence, as each small adjustment in Q propagates into all workers' decisions.

Under bottom entry, the path also exhibits high-frequency wiggles around its smooth trend. These reflect a genuine non-smoothness in the dynamics rather than numerical noise. The optimal AI intensity $\alpha^*(\gamma^H)$ is a step-like function (Figure 3, panel d): workers below the unstable fixed

point of the H-map sit at the corner $\alpha = 1$, while workers above it sit at the corner $\alpha = 0$. As the distribution evolves, mass crosses the threshold in discrete jumps—a small fraction of workers flips from full AI to full labor in a given period, and back in the next. This produces small jumps in Φ_t and, via the Kalman recursion, in Q_t. The fixed-point map on Q_t is therefore non-contractive at the threshold, and a damped iteration retains a residual two-period oscillation that we break by averaging over the last several iterates. The economic content of the path is not affected—the smoothed trajectory tracks the underlying mean dynamics—but the wiggles are a real feature of the model rather than an artifact of the solver.

6 Constrained Efficiency

The competitive equilibrium features an aggregate data quality externality: each worker takes Q as given when choosing α, but her action affects Φ and hence future Q. A constrained planner who internalizes this externality, but otherwise respects the model's primitives, would choose different AI intensity. We analyze this problem in two steps: first the steady-state planner, then the dynamic planner.

6.1 Steady-state planner

The steady-state planner chooses a stationary type-dependent policy $\alpha(\gamma^H)$ to maximize flow welfare:

$$\max_{\alpha(\cdot)} \int u\big(\gamma^H, \alpha(\gamma^H); Q\big)\, dG(\gamma^H),$$

subject to two constraints:

1. Distribution Stationarity: G is the invariant distribution of expertise under the policy $\alpha(\cdot)$, the skill transition map H, and worker turnover with exit rate δ and entry distribution F_0.

2. Aggregate data consistency: Q satisfies the steady-state version of (25), with Φ determined by the policy $\alpha(\cdot)$, the distribution G, and Q itself through (24).

This is an infinite-dimensional problem: the planner must choose a function $\alpha(\gamma^H)$ subject to fixed-point constraints on both the distribution G and data quality Q. Direct computation is intractable. A key simplification is that the externality operates only through the scalar Q: every

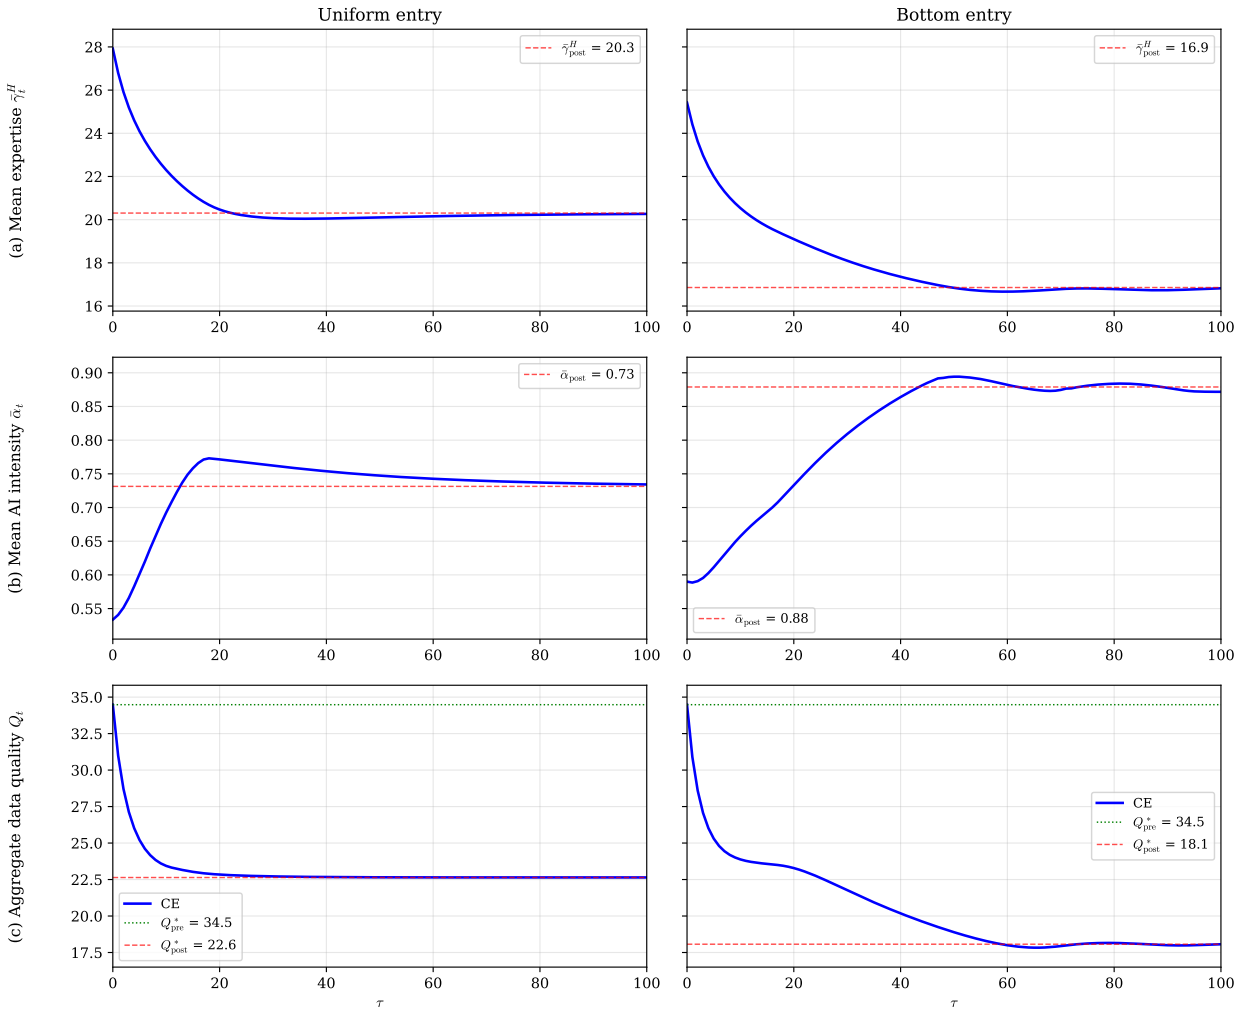

Figure 4: Transition dynamics from the pre-AI to the post-AI steady state. Left column: uniform entry. Right column: bottom entry. Rows show (a) mean expertise, (b) mean AI intensity, and (c) aggregate data quality Q_t. Y-axes are shared across columns within each row to make the comparison direct. Horizontal red dashed lines mark the post-AI steady-state values toward which each variable converges: $\bar{\gamma}^H_{\text{post}}$ in (a), $\bar{\alpha}_{\text{post}}$ in (b), and Q^*_{post} in (c). The horizontal green dotted line in panel (c) marks the pre-AI calibration anchor $Q^*_{\text{pre}} = \bar{\kappa} + \gamma^l(\bar{\kappa}) = 34.5$, the total forecasting precision of the best human expert (cf. Section 4).

worker's flow payoff and skill dynamics depend on other workers' actions *only* through aggregate data quality. This means the planner's problem can be decomposed. Rather than solving the full infinite-dimensional problem, we can construct an auxiliary Bellman equation that adds a single correction term to each worker's private payoff, parameterized by a scalar $\tilde{\mu}$—the shadow value of data quality. This shadow price represents the social value of a marginal unit of novel information contributed to the aggregate dataset. Searching over $\tilde{\mu}$ then recovers the planner's optimum.

Proposition 4 (Steady-State Auxiliary Bellman)**.** *The steady-state planner's problem is equivalent to the following two-step procedure:*

1. Inner problem. *For a given data quality price $\tilde{\mu} \geq 0$, solve for the multiplier $\Lambda(\gamma^H)$ satisfying*

$$\Lambda(\gamma^H) = \max_{\alpha \in [0,1]} \Big[\underbrace{u(\gamma^H, \alpha; Q)}_{\textit{private payoff}} + \underbrace{\tilde{\mu}\, \varphi(\gamma^H, \alpha; Q)}_{\textit{data subsidy}} + (1-\delta)\, \Lambda\big(H(\gamma^H, \alpha)\big) \Big], \tag{26}$$

   where $\varphi(\gamma^H, \alpha; Q) = \gamma_D[(1-\alpha)\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H)) + \alpha\gamma^a/(\gamma^H + Q + \gamma^a)]^2$ is the individual data contribution from Proposition 2. This yields a policy $\alpha^*(\gamma^H; \tilde{\mu})$, a stationary distribution $G(\tilde{\mu})$, and an aggregate-consistent $Q(\tilde{\mu})$.

2. Outer problem. *Search over $\tilde{\mu}$ to maximize* actual *welfare evaluated at the original payoff:*

$$\tilde{\mu}^* = \arg\max_{\tilde{\mu} \geq 0} \int u\big(\gamma^H,\, \alpha^*(\gamma^H; \tilde{\mu});\, Q(\tilde{\mu})\big)\, dG(\tilde{\mu}).$$

Here $\Lambda(\gamma^H)$ is the Lagrange multiplier on the stationarity constraint—the social shadow value of having a worker at skill γ^H in the stationary distribution. The proof is in Appendix B.[10]

To see where (26) comes from, consider two perturbations to the planner's allocation (the full derivation is in Appendix B).

The multiplier $\Lambda(\gamma^H)$ is the *social shadow value* of having a worker at skill γ^H in the stationary distribution—the marginal increase in the planner's welfare from adding one infinitesimal worker at γ^H. This shadow value accounts for everything the worker contributes: her own output, her effect on data quality, and the value of the skill level she will reach in future periods.

Suppose the planner marginally increases AI intensity for workers at skill γ^H. Three effects

[10] This is the *golden rule* planner—analogous to the golden rule in the Ramsey growth model. The dynamic planner (Section 6.4) discounts at $\beta_s(1-\delta)$ and converges to a *modified golden rule* with a less aggressive correction.

arise, and at the optimum they must sum to zero:

$$\underbrace{\frac{\partial u}{\partial \alpha}}_{\substack{\text{private marginal}\\ \text{payoff from AI}}} + \underbrace{\tilde{\mu}\,\frac{\partial \varphi}{\partial \alpha}}_{\substack{\text{externality}\\ \text{correction}}} + \underbrace{(1-\delta)\,\Lambda'\big(H(\gamma^H,\alpha)\big)\,\frac{\partial H}{\partial \alpha}}_{\substack{\text{effect on future}\\ \text{skill composition}}} = 0, \tag{27}$$

where $\tilde{\mu} = \mu \cdot \partial\Gamma/\partial\Phi$. The first term is the private marginal payoff from AI—the same trade-off a CE worker faces. The second term prices the worker's marginal effect on data quality: more AI means less novel data, which the planner values at $\tilde{\mu}$ per unit. The third term captures how changing α shifts the worker's next-period skill from H to a slightly different level; this has social shadow value $\Lambda'(H)$ and is realized only if the worker survives (probability $1-\delta$).

Now consider adding one extra worker at skill γ^H to the stationary distribution (see Appendix B for the full derivation). The stationarity constraint requires inflow to equal outflow at *each* skill level separately, not just in aggregate, so there is a type-specific shadow price $\Lambda(\gamma^H)$ at every γ^H. Adding one worker at γ^H violates stationarity at that skill level—effectively crowding out another worker of the same type—and the shadow cost of this is $\Lambda(\gamma^H)$. The benefit is the worker's per-period social contribution. At the optimum, the planner is indifferent—the cost exactly equals the benefit:

$$\underbrace{\Lambda(\gamma^H)}_{\substack{\text{shadow cost:}\\ \text{crowd out at } \gamma^H}} = \underbrace{u(\gamma^H,\alpha^*;Q)}_{\text{flow output}} + \underbrace{\tilde{\mu}\,\varphi(\gamma^H,\alpha^*;Q)}_{\text{data contribution}} + \underbrace{(1-\delta)\,\Lambda\big(H(\gamma^H,\alpha^*)\big)}_{\text{continuation value}}. \tag{28}$$

The right side is the benefit: her flow output, her contribution to data quality valued at the shadow price $\tilde{\mu}$, and a continuation value—with probability $(1-\delta)$ she survives and becomes a worker at $H(\gamma^H,\alpha^*)$, which is worth $\Lambda(H)$ to the planner. This indifference condition has a recursive structure and is exactly the auxiliary Bellman (26).

The planner's time preference β does not appear in (28) because the planner's objective $\int u\,dG$ evaluates a single stationary cross-section—there is no time dimension. The only reason the worker's future matters is that she transitions to a new skill level if she survives, hence the factor $(1-\delta)$ but no β.

The proposition transforms the planner's infinite-dimensional problem into a one-dimensional search over the scalar $\tilde{\mu}$. The auxiliary Bellman (26) adds a single term to the CE Bellman: $\tilde{\mu}\,\varphi(\gamma^H,\alpha;Q)$. This term subsidizes the worker for her contribution to the aggregate dataset. The function φ measures how much novel information the worker's action generates. At $\tilde{\mu}=0$, we recover the CE.

At $\tilde{\mu} > 0$, workers are rewarded for generating novel data, which tilts their choice toward labor.

Why is $\tilde{\mu}$ a scalar but the correction type-dependent? Because the data contribution φ depends on skill γ^H. At low skill, the labor innovation weight $\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))$ is large—an unskilled worker's labor action contains a high fraction of fresh information. At high skill, the innovation weight shrinks toward zero. The uniform price $\tilde{\mu}$ therefore generates a *larger effective subsidy* for low-skill workers, who are precisely the ones that the CE pushes most aggressively toward AI.

Table 3: Steady-state planner vs. competitive equilibrium

	Uniform entry		Bottom entry	
	CE	Planner	CE	Planner
Data quality (Q^*)	22.6	28.8	18.1	27.4
Mean AI intensity	0.73	0.40	0.88	0.39
Mean skill	20.3	23.8	16.8	21.7
Welfare gains				
Flow gain		+5.9%		+12.3%
Lifetime gain		+7.2%		+11.8%

Figure 5 compares the planner's allocation to the CE for both entry specifications, and Table 3 summarizes the key outcomes. Several findings emerge. First, the planner substantially reduces AI intensity: mean AI use falls from 0.73 to 0.40 (uniform) and from 0.88 to 0.39 (bottom). The planner does not eliminate AI—it still accounts for roughly 40% of task allocation—but the correction is concentrated at the early-mid-career range where the dynamic and informational gains from labor are strongest. Second, the skill trap is mitigated: the slow zone shrinks because the planner restricts AI use through this critical range, keeping the learning input higher (panel d). Third, data quality recovers substantially: from 22.6 to 28.8 (uniform, +27%) and from 18.1 to 27.4 (bottom, +51%). Fourth, welfare improves under both flow and lifetime measures: flow welfare rises by 5.9% (uniform) and 12.3% (bottom), and lifetime welfare rises by 7.2% and 11.8% respectively. The gains are much larger under bottom entry because the CE traps every entering worker in the slow AI zone, so the planner's correction has a disproportionately large effect.[11]

[11]The flow and lifetime gains differ because the lifetime measure also weights how the planner improves the prospects of fresh entrants. Under uniform entry the lifetime gain exceeds the flow gain; under bottom entry the ordering reverses. Appendix C derives an exact decomposition $\int V\,dG = (\int u\,dG - \beta\delta \int V\,dF_0)/(1-\beta)$ that pins the gap to the gain in entry-cohort lifetime value $\int V\,dF_0$, which is large under bottom entry (every entrant is at the planner's most aggressive intervention point) and small under uniform entry (some entrants would already reach the

The planner reduces AI intensity at every skill level relative to the CE under both entries (panel a). The optimal $\alpha^*(\gamma^H)$ is monotone declining in γ^H, with the strongest correction concentrated in the early-mid-career range $\gamma^H \in [15, 22]$ where labor signals become informative and skill accumulation dynamics are strongest.

The optimal intervention is not a uniform AI tax but a *skill-contingent* restriction that is most binding for workers in the early-mid stage of their careers and relaxes once expertise has accumulated and the data-supply role of the worker is established. We discuss implementation in Section 7.

Robustness of the benchmark results. Appendix D reports comparative statics with respect to the four AI-related parameters and the turnover rate δ, together with a structural extension allowing interior learning synergy between labor and AI. Two main results emerge.

First, the bimodal CE skill distribution—a low-skill mass at the AI trap coexisting with a high-skill mass at $\bar{\kappa}$—*strengthens* as AI becomes more capable or more attractive. Pushing γ^a up an order of magnitude makes AI more informationally valuable; workers lean on it more heavily in the CE, the trap mass grows, and the planner's welfare gain rises above the benchmark. Lowering the cost of AI has the same qualitative effect: the cheaper AI is, the more attractive it becomes for low-skill workers, the more pronounced the bimodality, and the larger the welfare gain. Shortening the effective career from 30 to 10 years also sharpens the trap, because shorter horizons make workers internalize less of the future skill cost of AI use. The interior synergy extension—which allows a moderate mix of labor and AI to build skill faster than either pure mode—leaves the bimodal CE intact, since the deskilling penalty at high α is enough to sustain the trap even when low-α collaboration is mildly complementary. Bimodality therefore grows with AI's informational quality, with AI's market penetration, and with the speed of cohort turnover.

In the opposite direction, bimodality dissolves only as the deskilling channel itself disappears—when c approaches one, AI and labor build skill at identical rates. Even there, the planner's allocation differs meaningfully from the CE. At $c = 0.75$, where the bimodal shape has already faded, the CE skill distribution concentrates near $\gamma^H \approx 25$ while the planner shifts mass up to $\bar{\kappa} \approx 30$, with welfare gains in the high single digits. The planner continues to internalize the data-quality externality regardless of whether AI is a worse teacher: AI-assisted actions remain mostly echoes

expert level under CE).

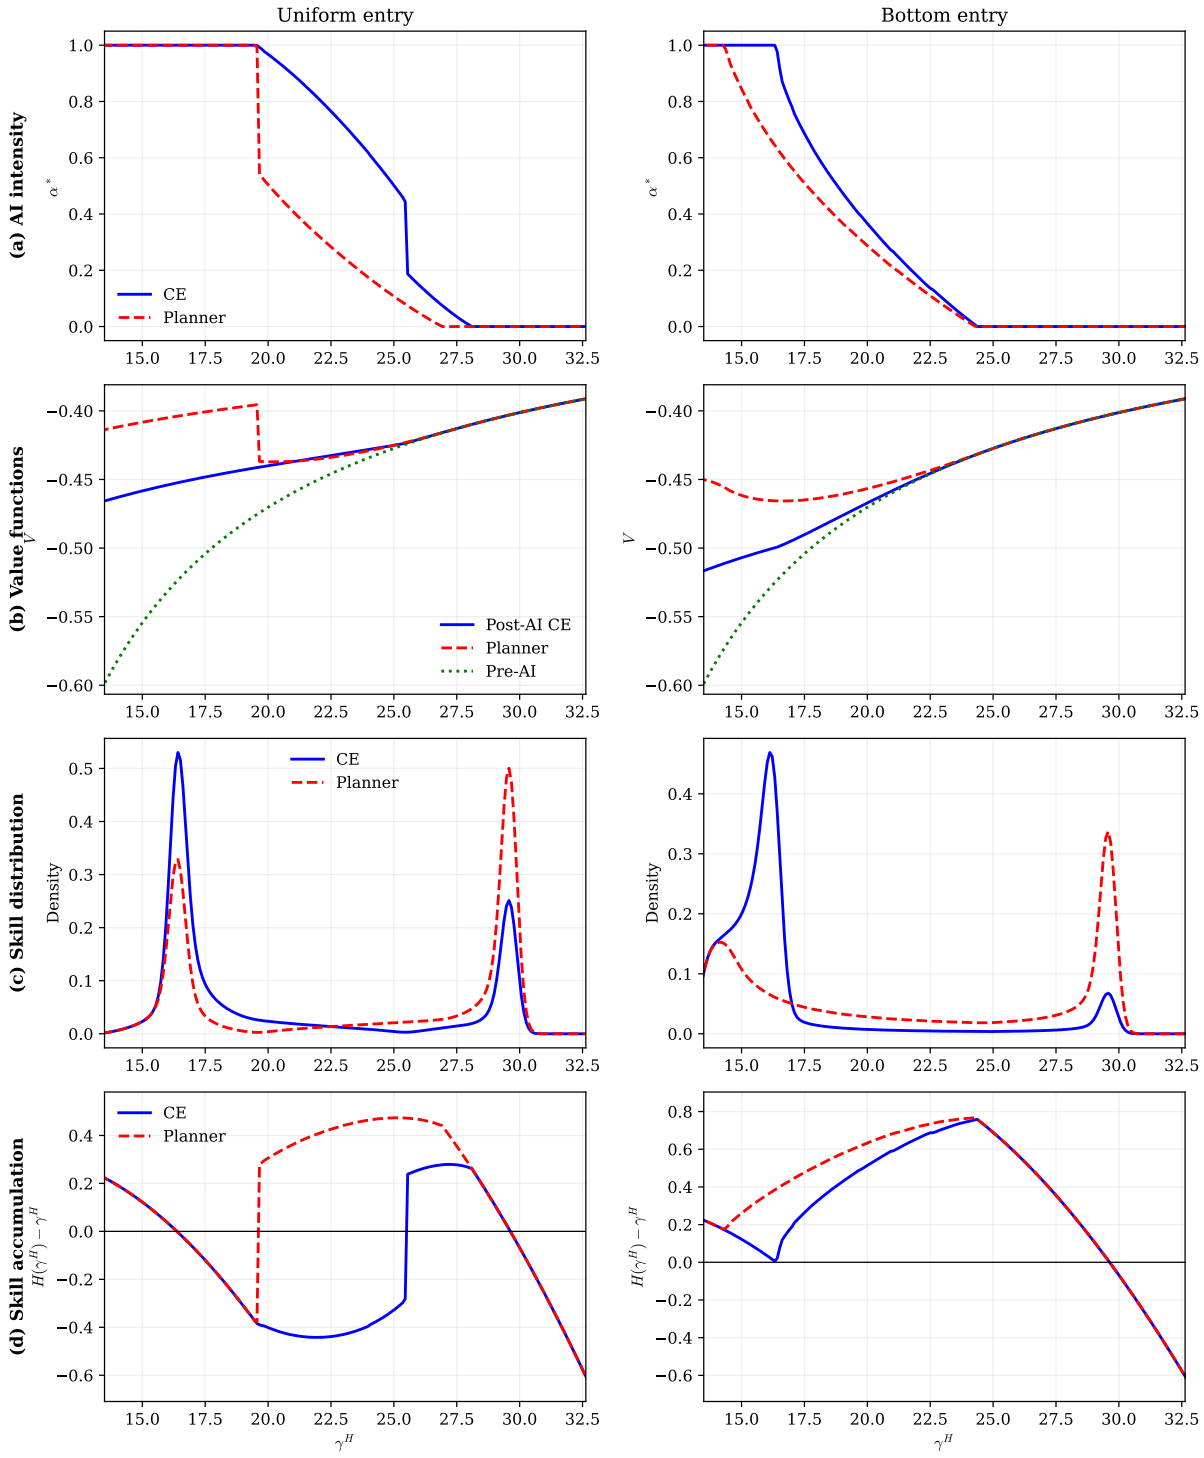

Figure 5: Steady-state planner vs. competitive equilibrium. Columns: uniform entry (left), bottom entry (right). Rows show (a) AI intensity $\alpha^*(\gamma^H)$, (b) lifetime value (with green dotted line showing the pre-AI value), (c) skill distribution, and (d) skill accumulation $H(\gamma^H) - \gamma^H$. The CE (blue solid) shows the classic three-fixed-point structure: a low-skill trap, an unstable interior point, and the full-labor steady state at $\bar{\kappa}$. The planner (red dashed) shifts substantial mass from the low-skill trap to the high-skill steady state but a residual mass of fully-AI-dependent workers remains.

of the existing dataset, so restricting them stays valuable even when the skill-trap mechanism is switched off.

Second, the planner's welfare gain is consistently large—ranging from the mid single digits to the mid double digits in flow terms across the entire space of perturbations, and never collapsing toward zero. The robustness has a structural source: the gain reflects two distinct externalities—the skill-trap channel and the data-quality channel—which respond differently to parameter changes. The skill-trap channel is largest precisely when AI is most attractive: cheap, capable, or used heavily. The data-quality channel operates independently of skill formation, so the planner's correction stays valuable even when the deskilling channel is switched off. As a result, improving AI technology raises the gain rather than lowering it: the externality scales with how much workers rely on AI, so the more advanced AI becomes, the larger the welfare loss from leaving its data externality unaddressed. The benchmark result that internalizing the data externality delivers economically meaningful welfare gains is therefore not a knife-edge feature of the calibration.

6.2 Data Contribution by Skill

Proposition 3 holds the cross-sectional distribution fixed and shows that the integrand of Φ^* in equation (24) is decreasing in each worker's a. We now examine how this per-type contribution varies across skill levels in equilibrium, once $a^*(\gamma^H)$ is determined endogenously.

Figure 6 plots the integrand at the steady-state policy and steady-state Q^* for the competitive equilibrium and the planner, alongside the full-labor benchmark $\gamma_D[\gamma^l(\gamma^H)/(\gamma^H+\gamma^l(\gamma^H))]^2$ that each type would contribute under $a \equiv 0$. The benchmark is increasing in γ^H in the relevant range: under Assumption 1, $\gamma^l(\gamma^H)$ rises with expertise, so a labor action from an experienced worker carries a higher-precision novel signal than a labor action from a novice. Pure-labor data production is therefore concentrated at the higher end of the skill distribution.

Three patterns emerge. First, both the CE and planner curves lie below the full-labor benchmark across the support—workers in both economies use some AI, which dilutes the per-type contribution. The vertical gap between actual contribution and benchmark at each γ^H measures the data loss caused by the worker's chosen a^*. Second, the planner's curve sits consistently above

the CE curve, with the largest gap in the early-mid-career range $\gamma^H \in [16, 22]$. By restricting AI intensity at every skill level, the planner lifts the per-type contribution toward the full-labor frontier, and the lift is largest precisely where the CE concentrates workers at $a^* \approx 1$. Third, the two curves merge at high expertise: workers near $\bar{\kappa}$ use almost no AI in either economy, so both contributions approach the benchmark.

The shape of the CE curve reflects the step structure of $a^*(\gamma^H)$. Under uniform entry, the CE policy jumps sharply from $a^* = 1$ to much lower intensities at the tipping point ($\gamma^H \approx 19$), producing a corresponding jump in the per-type contribution. The planner's curve is smoother because the planner's policy declines smoothly with expertise, without a hard threshold. Under bottom entry, the planner-CE gap extends down to the entry skill $\gamma_0^H = 13.5$, since every entrant arrives within the range where the planner's correction is binding.

The cross-sectional message is that the data externality is broad-based rather than concentrated at any one skill level: the planner does not relocate data production from one type to another but rather *lifts* the per-type contribution at every type, with the lift concentrated where AI over-adoption is most severe.[12] The dual benefit of restricting AI for early-career workers—the contemporaneous data-contribution gain documented here, plus the dynamic skill-accumulation gain documented in Section 6.1—operates through distinct but complementary channels.

6.3 Does improving AI eliminate the externality?

A natural question is whether the data quality externality is a transitional friction that vanishes as AI technology improves, or a persistent feature. AI performance depends on the interaction of data, model parameters, and compute—what the scaling laws literature calls the "AI triad" (Kaplan et al., 2020; Hoffmann et al., 2022). Recent estimates suggest that high-quality human-generated text data may be exhausted within a few years (Villalobos et al., 2024), and prominent researchers have argued that data has become the binding constraint on AI progress (Sutskever, 2024). In our model, γ_D governs one dimension of this triad: how efficiently AI extracts novel information from workers' recorded actions. Increasing γ_D captures advances in data processing—

[12] Because the labor signal precision $\gamma^l(\gamma^H)$ rises with expertise (Assumption 1), the planner's skill investment additionally serves as an investment in future data-generation capacity: a more experienced worker performing labor contributes a higher-precision novel signal, so building expertise today raises tomorrow's data-production frontier. The concavity of empirical wage–experience profiles disciplines how quickly $\gamma^l(\gamma^H)$ can grow.

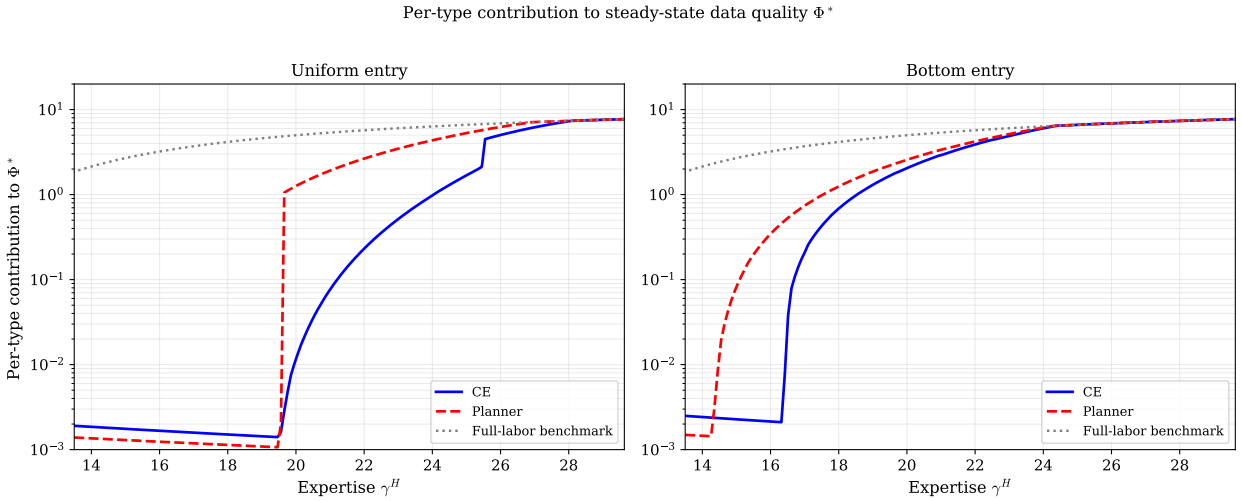

Figure 6: Per-type contribution to steady-state data quality, the integrand of equation (24) evaluated at the steady-state policy $\alpha^*(\gamma^H)$ and steady-state Q^*. Left: uniform entry. Right: bottom entry. Blue solid: CE; red dashed: planner; grey dotted: full-labor benchmark $\gamma_D[\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))]^2$, each type's contribution under $\alpha \equiv 0$. The benchmark rises with γ^H because experienced workers extract richer labor signals. Both CE and planner sit below the benchmark; the planner sits consistently above the CE, with the largest gap in the early-mid-career range. Log scale on the vertical axis.

better algorithms, larger context windows, more efficient training—while holding fixed the *supply* of novel human-generated data. We now ask what happens to the competitive equilibrium and the planner's allocation as γ_D increases.

We vary γ_D from the baseline value of 390 up to 7,805 (a twenty-fold increase), holding all other parameters fixed. At each level, we compute the CE and the planner's optimal steady state under bottom entry. Table 4 reports the results. As γ_D rises twenty-fold, data quality in the CE remains

Table 4: Steady-state outcomes as AI data efficiency (γ_D) improves. Bottom entry.

	Competitive equilibrium		Planner		Flow gain
γ_D	Data quality	AI intensity	Data quality	AI intensity	(planner over CE)
390	18.1	0.88	27.4	0.39	+12.3%
1,171	18.1	0.96	36.1	0.51	+20.9%
1,951	18.1	0.96	41.4	0.55	+25.1%
2,732	18.1	0.98	45.1	0.58	+28.0%
3,902	18.1	0.99	49.7	0.60	+31.0%
5,853	18.1	0.99	53.7	0.65	+34.3%
7,805	18.1	0.99	59.7	0.61	+36.6%

essentially flat at 17. The bottleneck is not data processing but data *supply*: nearly all workers use AI almost exclusively, so their actions are overwhelmingly echoes of the existing dataset. No amount of processing efficiency can extract novel information that was never generated.

The planner's welfare gain rises sharply—flow gains range from +12.3% at the baseline to +36.6% at the highest γ_D. The economic mechanism is complementarity between data processing and data generation. γ_D enters the data-quality law of motion (25) multiplicatively: a higher γ_D does not just improve AI's baseline performance, it raises the marginal return to each unit of novel human-generated data, so each labor signal feeds a more powerful pipeline. The CE breaks this complementarity through the externality—workers do not internalize how their labor contributes to the shared dataset, so they over-adopt AI and starve the technology of its most valuable input. The planner restores it by keeping enough workers in labor to feed the pipeline; as γ_D grows, the value of this correction grows with it. The pattern echoes the AI scaling laws literature (Kaplan et al., 2020; Hoffmann et al., 2022): AI performance depends multiplicatively on the interaction of data, model parameters, and compute, and improving one factor delivers its full benefit only when the others keep pace. Here the binding "other factor" is the supply of novel human-generated information.

Quantitatively, the planner maintains substantially lower AI intensity than the CE (which approaches full AI as γ_D grows), so the two economies diverge further as technology improves. The planner's flow of novel data is amplified by higher γ_D into dramatically higher data quality (Q rises from 27.4 at $\gamma_D = 390$ to 59.7 at $\gamma_D = 7{,}805$, while the CE's Q stays stuck at 18.1). Figure 7 illustrates this at $\gamma_D = 7{,}805$.

6.4 Dynamic planner

The dynamic planner maximizes $\sum_{t=0}^{\infty} \beta_s^t \int u(\gamma^H, \alpha_t(\gamma^H); Q_t)\, dG_t(\gamma^H)$ subject to the distribution dynamics, the data quality law of motion (25), and the initial condition from the pre-AI economy. By the same logic as the steady-state problem, the dynamic planner's allocation can be characterized by a time-varying auxiliary Bellman equation with a sequence of shadow prices $\{\tilde{\mu}_t\}$ on data quality (Proposition 5 in Appendix B). The key difference from the steady-state case is that the discount is $\beta_s(1-\delta)$ rather than $(1-\delta)$: since the distribution and data quality evolve over time,

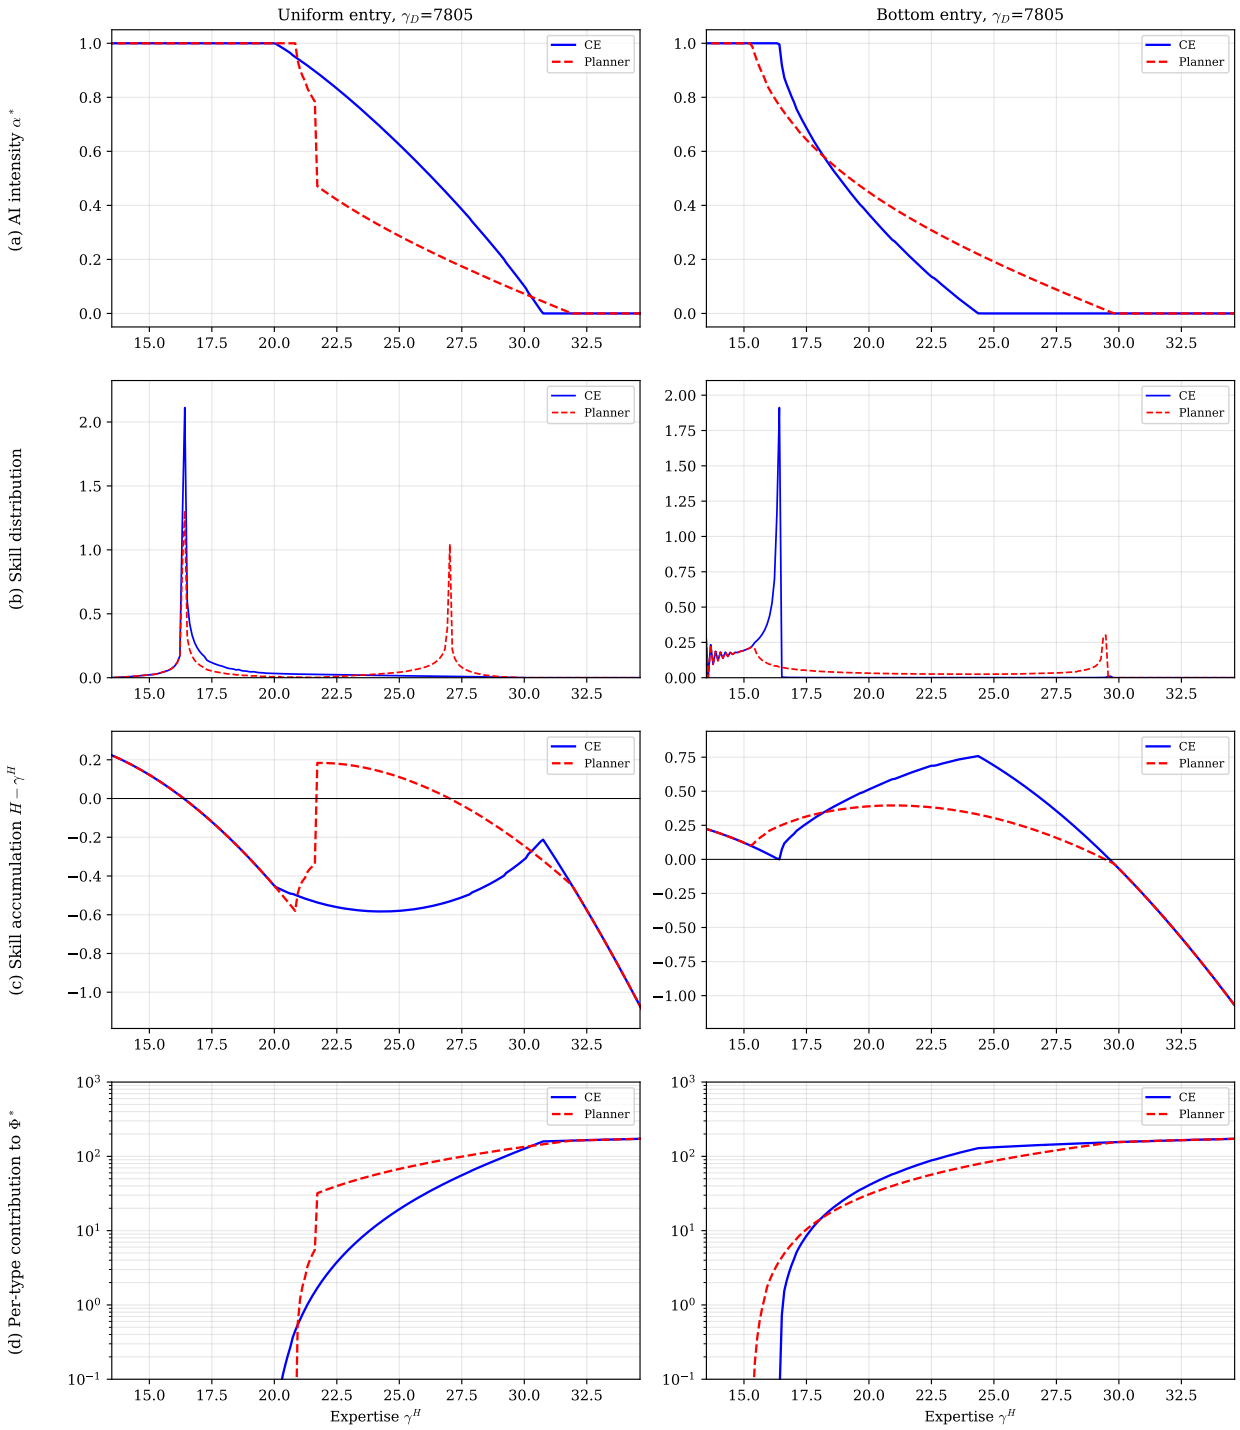

Figure 7: Planner vs. CE at high data efficiency ($\gamma_D = 7{,}805$, twenty times the baseline). Left column: uniform entry. Right column: bottom entry. Despite the twenty-fold improvement in data processing, the CE barely improves because workers generate almost no novel data. The planner maintains moderate AI intensity, preserving the flow of novel information and achieving substantially higher data quality.

the planner's time preference no longer cancels.

The economic tradeoff along the transition is between short-run cost and long-run gain. Reducing AI intensity lowers current output precision for workers who switch from AI to labor: at date 0, a low-skill worker using AI achieves precision $\gamma^H + Q_0 + \gamma^a$, while the same worker using labor achieves only $\gamma^H + \gamma^l(\gamma^H)$—a substantial gap when Q_0 is large. But this sacrifice preserves the flow of novel data into the aggregate dataset, maintaining higher data quality and enabling faster skill accumulation for future cohorts. The shadow price of data quality $\tilde{\mu}_t$ is highest early in the transition, when the planner's correction can reshape the entire future path, and declines toward a constant as the economy approaches steady state.

Figure 8 shows the transition paths under both entry specifications. Three features stand out.

First, the planner reduces AI intensity at every date and every skill level relative to the CE (panel b). The reduction is largest along the transition, when the shadow price of data quality is highest, and stabilizes at the terminal steady state.

Second, the planner preserves data quality well above the CE path (panel c). Under uniform entry, the planner's transition data quality settles at approximately 27 versus the CE's 24; under bottom entry, 22 versus 18. Mean expertise is correspondingly higher throughout (panel a), because better data quality sustains faster skill accumulation.[13]

The transition welfare gain, computed as the lifetime raw-utility value at $t = 0$ integrated over the pre-AI distribution G_0, is $+1.5\%$ under uniform entry and $+0.3\%$ under bottom entry—substantially smaller than the steady-state gains in Table 3. The gap reflects the cost of redirection: switching workers from AI to labor at $t = 0$ lowers their current output precision, and this cost must be paid up front, while the data-quality benefits accrue only gradually as new cohorts displace old ones in the stationary distribution. With discount factor $\beta_s(1-\delta) \approx 0.928$, near-term losses receive substantial weight relative to far-future gains. The much larger steady-state gains in

[13]The transition's long-run data quality is the *modified golden rule* of the dynamic planner's problem, not the steady-state planner's golden rule reported in Table 3 ($Q^* = 30$ uniform, 27 bottom). The distinction parallels the Ramsey–Cass–Koopmans growth model. The steady-state planner maximizes the cross-sectional flow integral $\int u\, dG$ with no time dimension, so the relevant discount is just the survival factor $(1-\delta)$. The dynamic planner, by contrast, must trade current "consumption"—output precision today—against future "capital"—data quality tomorrow—and discounts at $\beta_s(1-\delta)$. Just as the Ramsey planner's capital stock settles strictly below the golden-rule level (because at golden-rule capital the marginal cost of foregone consumption exceeds the discounted marginal benefit of additional capital), dynamic data quality settles strictly below the SS planner's Q^*. Data quality plays the role of capital in this analogy, and forgoing AI use to feed novel signals into the aggregate dataset is the analogue of forgoing current consumption to invest in physical capital.

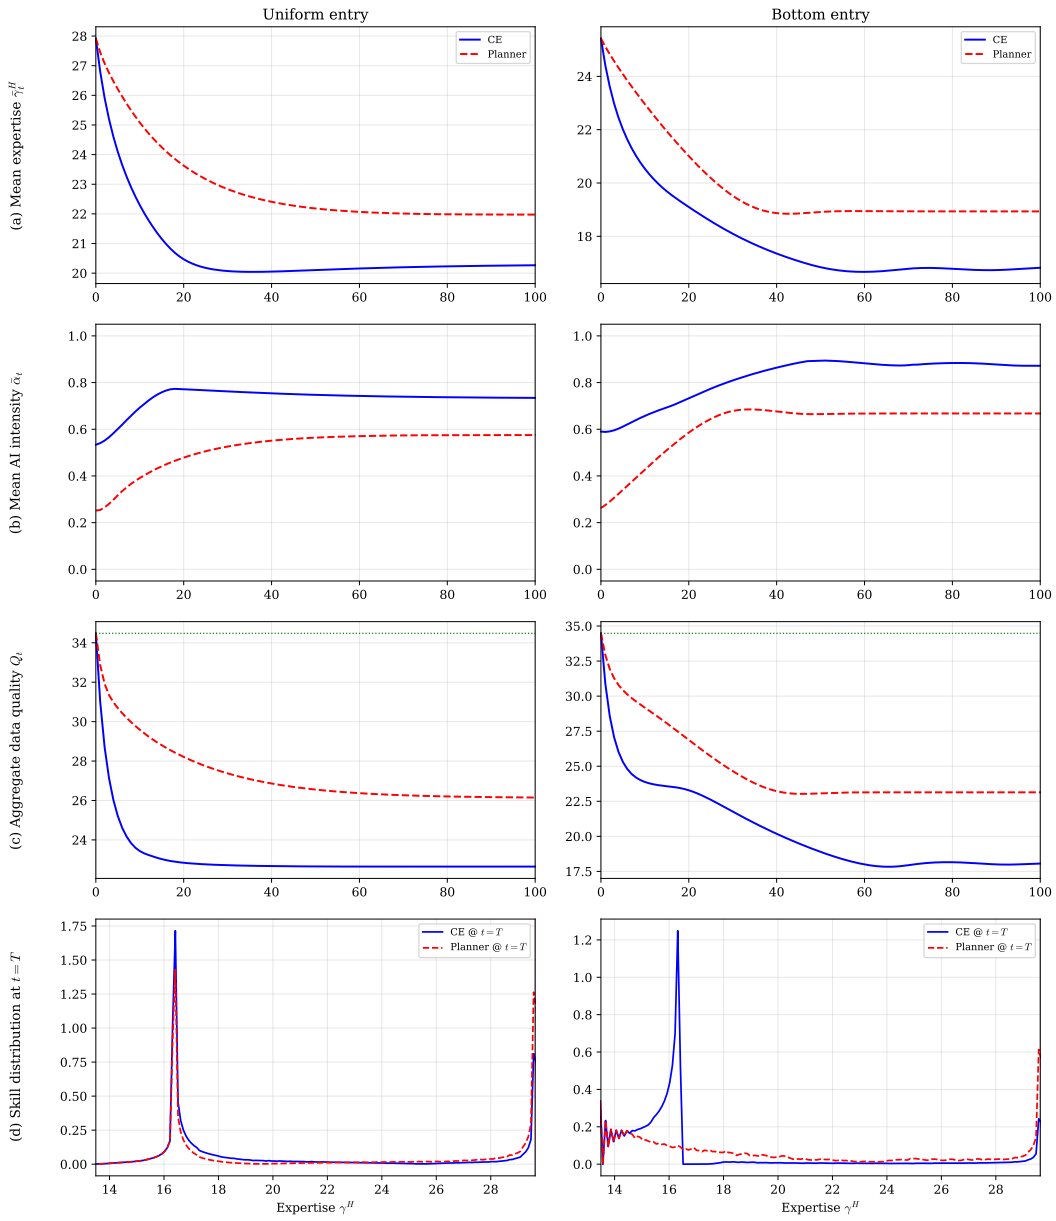

Figure 8: Dynamic planner vs. competitive equilibrium along the transition. Left column: uniform entry. Right column: bottom entry. Rows show (a) data quality, (b) mean skill, (c) mean AI intensity, and (d) AI intensity $\alpha^*(\gamma^H)$ at $t = 0$ (solid) and at the terminal steady state (dotted). The planner reduces AI intensity at every date and every skill level, preserving higher data quality and mean skill.

Table 3 (+5.9% and +12.3%) reflect the long-run cohort-weighted distribution after the transition has played out.

7 Conclusion

This paper has developed a model of AI adoption with endogenous skill formation and a data quality externality, producing three main findings.

First, the interaction between learning-by-doing and a data quality externality generates an *AI skill trap*. In theory, the model can produce multiple steady states with persistent divergence across workers; in the calibrated economy, it produces a bimodal skill distribution under every entry distribution we consider, with a labor-investing high-skill group coexisting with a fully AI-dependent low-skill group. The mechanism is the same in both cases: AI is effectively underpriced because workers do not bear the cost that their AI use imposes on the quality of the shared dataset.

Second, the constrained planner's optimal policy shifts mass out of the AI-dependent low-skill group but does not eliminate the group itself. The planner restricts AI use most for early- and mid-career workers, where the gap between actual and full-labor data contribution is largest under the CE and where the dynamic return to skill accumulation is highest; it allows moderate AI use for high-skill workers; and it concedes the lowest-skill region to full AI, because for those workers the labor-investment return cannot overcome the accumulated deskilling cost. The planner's stationary distribution remains bimodal: substantial mass has been lifted from the trap into the high-skill group, but a residual low-skill mass persists. Bimodality is therefore a robust feature of an AI economy with a data externality, not an artifact of decentralization—the externality governs the size of the AI-dependent low-skill group, not its existence. Internalizing the externality nonetheless delivers double-digit flow welfare gains in the worst-case entry distribution. The planner does not relocate data production across types but *lifts* the per-type contribution at every skill level, with the lift concentrated where AI over-adoption is most severe.

Third, improving AI technology does not resolve the externality. Even a twenty-fold improvement in data-processing efficiency barely moves the competitive equilibrium, because workers locked into AI contribute almost no novel information for the improved technology to exploit.

The planner's instrument—a shadow price on data quality—maps to a corrective tax on AI

use. The optimal correction is skill-contingent, not uniform: it restricts AI most for novice workers, both to keep them on a labor-learning trajectory toward higher skill and to extract data contributions that are otherwise lost to AI delegation. In practice, this resembles an experience-dependent AI access rule, implementable through training mandates that require minimum levels of hands-on work for junior workers—a policy that many professional fields already impose through apprenticeship and residency requirements.

Several limitations suggest directions for future work. The model features a single task type; extending it to heterogeneous occupations would allow analysis of how AI reshapes the composition of human capital investment. The model treats AI as improving only through the data channel, while in practice AI labs also improve models through architecture and compute; the interaction between endogenous data quality and exogenous technological progress is a natural extension. Finally, the model assumes a single aggregate dataset, while in practice firms may internalize some of the externality through proprietary data. Whether market structure mitigates or exacerbates the externality is an open question.

Our analysis focuses on workers' over-adoption of AI, but firms may over-adopt as well. Hosseini Maasoum and Lichtinger (2025) document that AI-adopting firms sharply reduce entry-level hiring, with junior headcount declining by 7.7% while senior employment is unaffected. This firm-side response mirrors the worker-side mechanism in our model: just as individual workers underinvest in skill formation by relying on AI, firms underinvest in their talent pipeline by substituting AI for junior roles. Both channels erode the stock of human expertise, and the interaction between them is a natural direction for future research.

References

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

_ **and Pascual Restrepo**, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

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

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **and** _ , "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

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

_ **, Tianyi Lin, Asuman Ozdaglar, and James Siderius**, "How AI Aggregation Affects Knowledge," NBER Working Paper 35036, National Bureau of Economic Research 2026.

Adams, Jonathan J., Min Fang, Zheng Liu, and Yajie Wang, "The Rise of AI Pricing: Trends, Driving Forces, and Implications for Firm Performance," *Journal of Monetary Economics*, 2026, *157*.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction," *Journal of Economic Perspectives*, 2019, *33* (2), 31–50.

Alemohammad, Sina, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoohi, and Richard G. Baraniuk, "Self-Consuming Generative Models Go MAD," *arXiv preprint arXiv:2307.01850*, 2023.

Amodei, Dario, "Machines of Loving Grace: How AI Could Transform the World for the Better," Essay, October 2024 2024.

_ , "The Adolescence of Technology," Essay, January 2026 2026.

ao Guerreiro, Jo Sergio Rebelo, and Pedro Teles, "Should Robots Be Taxed?," *Review of Economic Studies*, 2022, *89* (1), 279–311.

Autor, David, "Applying AI to Rebuild Middle Class Jobs," *NBER Working Paper*, 2024, (32140).

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Bai, Yuntao, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon et al., "Constitutional AI: Harmlessness from AI Feedback," *arXiv preprint arXiv:2212.08073*, 2022.

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Barcaui, André, "ChatGPT as a Cognitive Crutch: Evidence from a Randomized Controlled Trial on Knowledge Retention," *Computers and Education: Artificial Intelligence*, 2025, *8*, 100362.

Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakçı, and Rei Mariman, "Generative AI Can Harm Learning," *Proceedings of the National Academy of Sciences*, 2025.

Becker, Gary S., *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press, 1964.

Ben-Porath, Yoram, "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 1967, *75* (4), 352–365.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Briesch, Martin, Dominik Sobania, and Franz Rothlauf, "Large Language Models Suffer from Their Own Output: An Analysis of the Self-Consuming Training Loop," *arXiv preprint arXiv:2311.16822*, 2023.

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

Budzyń, Krzysztof, Marcin Romańczyk, Diana Kitala, Paweł Kołodziej, Marek Bugajski, Hans Olov Adami, Johannes Blom, Marek Buszkiewicz, Natalie Grace Halvorsen, Hassan Cesare, Tomasz Romańczyk, Øyvind Holme, Krzysztof Jarus, Shona Fielding, Melina A. Kunar, Maria Pellise, Nastazja Dagny Pilonis, Michał F. Kamiński, Mette Kalager, Michael Bretthauer, and Yuichi Mori, "Endoscopist Deskilling Risk after Exposure to Artificial Intelligence

in Colonoscopy: A Multicentre, Observational Study," *The Lancet Gastroenterology & Hepatology*, 2025.

Cai, Zhifeng, "Generative AI and Data Quality: Implications for Productivity, Labor Displacement, and Policy," *Working Paper, Rutgers University*, 2024.

Christiano, Paul F., Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei, "Deep Reinforcement Learning from Human Preferences," in "Advances in Neural Information Processing Systems," Vol. 30 2017.

Cui, Jingyi, Gabriel Dias, and Justin Ye, "Signaling in the Age of AI: Evidence from Cover Letters," *Working Paper, Yale University*, 2025.

Cui, Zheyuan Kevin, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz, "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers," *Management Science*, 2025.

del Rio-Chanona, R. Maria, Nadzeya Laurentsyeva, and Johannes Wachs, "Large Language Models Reduce Public Knowledge Sharing on Online Q&A Platforms," *PNAS Nexus*, 2024, *3* (9), pgae400.

Doshi, Anil R. and Oliver P. Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances*, 2024, *10* (28), eadn5290.

Dubey, Abhimanyu, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman et al., "The Llama 3 Herd of Models," *arXiv preprint arXiv:2407.21783*, 2024.

Fajgelbaum, Pablo D., Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty Traps," *Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Farboodi, Maryam and Laura Veldkamp, "A Growth Model of the Data Economy," *Review of Economic Studies*, 2025. Forthcoming.

Galdin, Anais and Jesse Silbert, "Making Talk Cheap: Generative AI and Labor Market Signaling," *Working Paper, Princeton University*, 2025.

Gerstgrasser, Matthias, Rylan Schaeffer, Apratim Dey, Rafael Rafailov, Henry Sleight, John Hughes, Tomasz Korbak, Rajashree Agrawal, Dhruv Pai, Andrey Gromov, Daniel A. Roberts, Diyi Yang, David L. Donoho, and Sanmi Koyejo, "Is Model Collapse Inevitable? Breaking the Curse of Recursion by Accumulating Real and Synthetic Data," *arXiv preprint arXiv:2404.01413*, 2024.

Goos, Maarten, Alan Manning, and Anna Salomons, "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 2014, *104* (8), 2509–2526.

_ **and** _ , "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 2007, *89* (1), 118–133.

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 1980, *70* (3), 393–408.

Gunasekar, Suriya, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi et al., "Textbooks Are All You Need," *arXiv preprint arXiv:2306.11644*, 2023.

Hémous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality," *American Economic Journal: Macroeconomics*, 2022, *14* (1), 179–223.

Hoffmann, Jordan, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark et al., "Training Compute-Optimal Large Language Models," *Advances in Neural Information Processing Systems*, 2022, *35*.

Ide, Enrique, "Automation, AI, and the Intergenerational Transmission of Knowledge," *Working Paper*, 2025. arXiv:2507.16078.

Jaimovich, Nir and Henry E. Siu, "Job Polarization and Jobless Recoveries," *Review of Economics and Statistics*, 2020, *102* (1), 129–147.

Jošt, Gregor, Viktor Taneski, and Sašo Karakaťič, "The Impact of Large Language Models on Programming Education and Student Learning Outcomes," *Applied Sciences*, 2024, *14* (10), 4115.

Kambourov, Gueorgui and Iourii Manovskii, "Occupational Specificity of Human Capital," *International Economic Review*, 2009, *50* (1), 63–115.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei, "Scaling Laws for Neural Language Models," *arXiv preprint arXiv:2001.08361*, 2020.

Korinek, Anton and Donghyun Suh, "Scenarios for the Transition to AGI," *NBER Working Paper*, 2024, (32255).

_ **and Joseph E. Stiglitz**, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in "The Economics of Artificial Intelligence: An Agenda," University of Chicago Press, 2019, pp. 349–390.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Life Cycle Wage Growth across Countries," *Journal of Political Economy*, 2018, *126* (2), 797–849.

Lipton, Zachary C., "The Mythos of Model Interpretability," *Communications of the ACM*, 2018, *61* (10), 36–43.

Maasoum, Seyed Mahdi Hosseini and Guy Lichtinger, "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," 2025. SSRN Working Paper No. 5425555.

METR, "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity," arXiv:2507.09089 2025.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 2022, *90* (6), 2645–2683.

Ouyang, Long, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray et al., "Training Language Models to Follow Instructions with Human Feedback," in "Advances in Neural Information Processing Systems," Vol. 35 2022, pp. 27730–27744.

Puri, Indira and Laura Veldkamp, "Artificial intelligence and cognitive inequality," *Journal of Monetary Economics*, 2025, *157*, 103884.

Rudin, Cynthia, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence*, 2019, *1* (5), 206–215.

Shen, Judy Hanwen and Alex Tamkin, "How AI Impacts Skill Formation," *arXiv preprint arXiv:2601.20245*, 2026.

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

Sutskever, Ilya, "Sequence to Sequence Learning with Neural Networks: A Pre-NeurIPS Talk," NeurIPS 2024 Keynote 2024. December 2024, Vancouver.

Veldkamp, Laura, *Information Choice in Macroeconomics and Finance*, Princeton University Press, 2011.

Villalobos, Pablo, Jaime Sevilla, Lennart Heim, Tamay Besiroglu, Marius Hobbhahn, and Anson Ho, "Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data," *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.

Vives, Xavier, "How Fast Do Rational Agents Learn?," *Review of Economic Studies*, 1993, *60* (2), 329–347.

A Proofs

A.1 Proof of Lemma 1 (Optimal AI Intensity)

When $\beta = 0$, the worker maximizes the flow payoff (16) net of the AI cost $\alpha\, p_A$. Because $\gamma^l(\gamma^H)$ does not depend on α, the objective is concave in α: the information gain $1/[\gamma^H + (1-\alpha)\gamma^l(\gamma^H) + \alpha(Q+\gamma^a)]^2$ has a strictly negative second derivative in α whenever $Q + \gamma^a \neq \gamma^l(\gamma^H)$, and the maintained condition $Q + \gamma^a > \gamma^l(\gamma^H)$ ensures this. Setting the derivative with respect to α equal

to zero gives

$$\frac{Q+\gamma^a-\gamma^l(\gamma^H)}{\left[\gamma^H+(1-\alpha)\gamma^l(\gamma^H)+\alpha(Q+\gamma^a)\right]^2}=p_A\,. \tag{29}$$

Solving for the interior solution:

$$\alpha^*(\gamma^H;Q)=\frac{\sqrt{\dfrac{Q+\gamma^a-\gamma^l(\gamma^H)}{p_A}}-\gamma^H-\gamma^l(\gamma^H)}{Q+\gamma^a-\gamma^l(\gamma^H)}\,. \tag{30}$$

This is interior when $0<\alpha^*<1$. The boundary conditions define the zone thresholds:

- Full AI ($\alpha^*=1$): $\gamma^H\le T(\gamma^H)-(Q+\gamma^a)$, where $T(\gamma^H)\equiv\sqrt{(Q+\gamma^a-\gamma^l(\gamma^H))/p_A}$.
- Full labor ($\alpha^*=0$): $\gamma^H\ge T(\gamma^H)-\gamma^l(\gamma^H)$.

With type-dependent $\gamma^l(\cdot)$ these are no longer linear inequalities in γ^H; they remain well-defined because the right-hand side of each is continuous and $T(\gamma^H)-(Q+\gamma^a)$ falls faster than γ^H rises (under Assumption 1, $T'(\gamma^H)=-\gamma^{l\prime}(\gamma^H)/(2\,T\,p_A)\le 0$).

Strict monotonicity of α^*. Define $G(\alpha,\gamma^H)\equiv(Q+\gamma^a-\gamma^l(\gamma^H))/D(\alpha,\gamma^H)^2-p_A$, where $D(\alpha,\gamma^H)\equiv\gamma^H+(1-\alpha)\gamma^l(\gamma^H)+\alpha(Q+\gamma^a)$. The interior optimum satisfies $G(\alpha^*,\gamma^H)=0$. By the implicit function theorem, $d\alpha^*/d\gamma^H=-(\partial G/\partial\gamma^H)/(\partial G/\partial\alpha)$. The denominator

$$\frac{\partial G}{\partial\alpha}=-\frac{2(Q+\gamma^a-\gamma^l(\gamma^H))^2}{D^3}\;<\;0$$

under $Q+\gamma^a>\gamma^l(\gamma^H)$. The numerator

$$\frac{\partial G}{\partial\gamma^H}=-\frac{\gamma^{l\prime}(\gamma^H)}{D^2}-\frac{2(Q+\gamma^a-\gamma^l(\gamma^H))\,[1+(1-\alpha)\gamma^{l\prime}(\gamma^H)]}{D^3}\;<\;0$$

because both terms are non-positive and the second is strictly negative under $Q+\gamma^a>\gamma^l(\gamma^H)$. Hence $d\alpha^*/d\gamma^H<0$. Combined with the boundary thresholds above, this delivers the three regimes of Lemma 1. The constant-precision benchmark $\gamma^l\equiv\bar\gamma^l$ is the special case $\gamma^{l\prime}\equiv 0$, in which (30) reduces to the closed-form policy of Cai (2024). □

A.2 Proof of Proposition 1 (Existence of the Skill Trap)

Permanent-policy steady states. The full-AI steady state $\underline{\kappa}$ solves $\underline{\kappa}=H(\underline{\kappa};Q)$ at $\alpha\equiv 1$, i.e.,

$$\frac{1}{\underline{\kappa}}=\frac{\rho^2}{\underline{\kappa}+\tilde\gamma^a(\underline{\kappa})}+\frac{1}{\gamma_\eta},$$

and the full-labor steady state $\bar\kappa$ solves the same equation with $\tilde\gamma^a$ replaced by γ^l. Both fixed points exist and are unique when $\tilde\gamma^a(\cdot)$ and $\gamma^l(\cdot)$ are continuous and bounded on the relevant range;

under Assumption 1 both functions are weakly increasing, and the right-hand side of each fixed-point equation is monotone in γ^H, so a single crossing of the 45° line is guaranteed. Assumption 1 ($\tilde{\gamma}^a(\gamma^H) < \gamma^l(\gamma^H)$ pointwise) implies $\underline{\kappa} < \bar{\kappa}$.

Price thresholds. For given Q with $Q + \gamma^a > \gamma^l(\gamma^H)$ on the relevant range, define

$$\bar{p}_A(Q) \equiv \frac{Q + \gamma^a - \gamma^l(\underline{\kappa})}{(\underline{\kappa} + Q + \gamma^a)^2}, \qquad \underline{p}_A(Q) \equiv \frac{Q + \gamma^a - \gamma^l(\bar{\kappa})}{(\bar{\kappa} + \gamma^l(\bar{\kappa}))^2}. \tag{31}$$

The thresholds correspond to the FOC (29) evaluated at $\alpha = 1, \gamma^H = \underline{\kappa}$ (boundary of the full-AI zone) and at $\alpha = 0, \gamma^H = \bar{\kappa}$ (boundary of the full-labor zone), respectively. We need $\underline{p}_A(Q) \leq \bar{p}_A(Q)$ for the trap to exist, which requires both the numerator and denominator of $\bar{p}_A / \underline{p}_A$ to favor $\bar{p}_A$. The numerator inequality $Q + \gamma^a - \gamma^l(\underline{\kappa}) \geq Q + \gamma^a - \gamma^l(\bar{\kappa})$ follows from $\gamma^l(\underline{\kappa}) \leq \gamma^l(\bar{\kappa})$ (Assumption 1 and $\underline{\kappa} < \bar{\kappa}$). The denominator inequality $(\underline{\kappa} + Q + \gamma^a)^2 \leq (\bar{\kappa} + \gamma^l(\bar{\kappa}))^2$ is equivalent to

$$Q \;\leq\; (\bar{\kappa} - \underline{\kappa}) + \gamma^l(\bar{\kappa}) - \gamma^a,$$

which holds whenever Q is not too large—an interpretable restriction parallel to the constant-γ^l benchmark of Cai (2024). Together these two inequalities give $\underline{p}_A(Q) \leq \bar{p}_A(Q)$, so the price interval is non-empty over a range of Q values.

Stability of the low fixed point. When $p_A \leq \bar{p}_A(Q)$, the myopic policy at $\gamma^H = \underline{\kappa}$ satisfies $\alpha^* = 1$. Since $H(\underline{\kappa}; Q)|_{\alpha=1} = \underline{\kappa}$ by construction, this is a stable fixed point: in the full-AI zone $\alpha^* \equiv 1$ so $L(\gamma^H, 1) = \gamma^H + \tilde{\gamma}^a(\gamma^H)$ and

$$L'(\gamma^H) = 1 + \tilde{\gamma}^{a\prime}(\gamma^H) = 1 + c\,\gamma^{l\prime}(\gamma^H),$$

which under Assumption 1 is bounded; $H' = K \cdot L' < 1$ provided the Kalman gain $K = \rho^2\gamma_\eta^2 / (\rho^2\gamma_\eta + L)^2$ is small enough, which holds whenever $\rho < 1$ and γ_η is bounded.

Amplification on the collaboration zone. On the interior of the collaboration zone, $L(\gamma^H, \alpha^*) = \gamma^H + (1 - \alpha^*)\gamma^l(\gamma^H) + \alpha^*\tilde{\gamma}^a(\gamma^H) = \gamma^H + \gamma^l(\gamma^H) - \alpha^*(1 - c)\gamma^l(\gamma^H)$ using $\tilde{\gamma}^a = c\gamma^l$. Differentiating along $\alpha^*(\gamma^H)$:

$$L'(\gamma^H) = 1 + \big[1 - \alpha^*(1 - c)\big]\gamma^{l\prime}(\gamma^H) - (1 - c)\gamma^l(\gamma^H)\,\frac{d\alpha^*}{d\gamma^H}. \tag{32}$$

From the proof of Lemma 1, the implicit-function expression for $d\alpha^*/d\gamma^H$ at the interior optimum

is

$$\frac{d\alpha^*}{d\gamma^H} = -\frac{1 + (1-\alpha^*)\gamma^{l\prime}(\gamma^H) + \gamma^{l\prime}(\gamma^H)/[2\,T(\gamma^H)\,p_A]}{Q + \gamma^a - \gamma^l(\gamma^H)},$$

where $T(\gamma^H) = \sqrt{(Q + \gamma^a - \gamma^l(\gamma^H))/p_A}$. Substituting into (32) and grouping terms,

$$L'(\gamma^H) = \underbrace{1 + \frac{(1-c)\gamma^l(\gamma^H)}{Q + \gamma^a - \gamma^l(\gamma^H)}}_{\text{substitution} > 1} + \underbrace{\Lambda(\gamma^H, \alpha^*)\,\gamma^{l\prime}(\gamma^H)}_{\text{profile} \geq 0}, \tag{33}$$

where the profile coefficient

$$\Lambda(\gamma^H, \alpha^*) \equiv \left[1 - \alpha^*(1-c)\right] + \frac{(1-c)\gamma^l(\gamma^H)\left[(1-\alpha^*) + 1/(2T(\gamma^H)p_A)\right]}{Q + \gamma^a - \gamma^l(\gamma^H)} \;\geq\; 0.$$

The substitution component recovers the constant-precision formula of (21) with $\gamma^l \to \gamma^l(\gamma^H)$, and the profile component is non-negative under Assumption 1, with strict positivity whenever $\gamma^l(\cdot)$ is strictly increasing.

Conclusion. As $Q + \gamma^a \to \gamma^l(\gamma^H)$ from above for some γ^H in the collaboration zone, the substitution component of $L'(\gamma^H)$ diverges to $+\infty$, so $H'(\gamma^H) = K \cdot L'(\gamma^H) > 1$ at such γ^H. The collaboration zone is non-empty whenever $\underline{p}_A(Q) \leq p_A \leq \bar{p}_A(Q)$. By the deskilling assumption $\bar{\kappa} > \underline{\kappa}$, and by Assumption 1 the threshold map $Q + \gamma^a - \gamma^l(\gamma^H)$ is continuous in γ^H, so for any Q in the open interval $(\gamma^l(\underline{\kappa}) - \gamma^a,\ \gamma^l(\bar{\kappa}) - \gamma^a + (\bar{\kappa} - \underline{\kappa}))$ there is a non-empty subinterval of γ^H on which $H' > 1$. Continuity of H together with $H' < 1$ in the full-AI zone (low fixed point at $\underline{\kappa}$) and $H' < 1$ for large γ^H (Kalman attenuation dominates as $\alpha^* \to 0$) implies H crosses the 45° line from below at least once above $\underline{\kappa}$, giving a second stable fixed point. The high fixed point need not lie in the full-labor zone; it can occur at an interior $\alpha^* > 0$ when the calibrated dataset is rich enough that experts still find some AI use beneficial. □

A.3 Proof of Proposition 2 (Aggregate Data Quality)

We derive the law of motion for $Q_t = 1/\mathrm{Var}(\theta_t|\Omega_{t-1})$ by computing the Bayesian posterior precision about θ_t after observing the $\bar{N}$ new records $\{D_t^i\}_{i=1}^{\bar{N}}$, then forwarding one period through the AR(1) dynamics of θ. The argument proceeds in three steps and parallels the derivation in Cai (2024) for the short-lived-agent benchmark, adapted to our continuous cross-sectional distribution of expertise.

Step 1: Recorded actions as noisy signals about θ_t.

Worker i takes an action with AI intensity α_t^i. Her posterior mean for θ_t, from equation (23), is

$$a_t^i \;=\; \underbrace{\Big[(1-\alpha_t^i)\tfrac{\gamma_t^{H,i}}{\gamma_t^{H,i}+\gamma^l(\gamma_t^{H,i})}+\alpha_t^i\tfrac{\gamma_t^{H,i}+Q_t}{\gamma_t^{H,i}+Q_t+\gamma^a}\Big]E(\theta_t|\mathcal{H}_t^i\cup\Omega_{t-1})}_{\text{“known” component}} \;+\; \varphi_i\,\theta_t \;+\; \xi_t^i, \tag{34}$$

where

$$\varphi_i \;\equiv\; (1-\alpha_t^i)\,\frac{\gamma^l(\gamma_t^{H,i})}{\gamma_t^{H,i}+\gamma^l(\gamma_t^{H,i})} \;+\; \alpha_t^i\,\frac{\gamma^a}{\gamma_t^{H,i}+Q_t+\gamma^a}$$

is the *innovation weight* on the fresh signal, and

$$\xi_t^i \;=\; (1-\alpha_t^i)\,\frac{\gamma^l(\gamma_t^{H,i})}{\gamma_t^{H,i}+\gamma^l(\gamma_t^{H,i})}\,\varepsilon_t^{l,i} \;+\; \alpha_t^i\,\frac{\gamma^a}{\gamma_t^{H,i}+Q_t+\gamma^a}\,\varepsilon_t^{A,i}$$

is the signal-noise residual, with $\varepsilon_t^{l,i}\sim\mathcal{N}(0,1/\gamma^l(\gamma_t^{H,i}))$ and $\varepsilon_t^{A,i}\sim\mathcal{N}(0,1/\gamma^a)$ independent across workers and across the two signal types. Under Assumption 2, the prior component $E(\theta_t|\mathcal{H}_t^i)$ is perfectly recoverable from Ω_{t-1} (the dataset “knows” what the worker knows from her history), so the entire “known” bracket in (34) is measurable with respect to Ω_{t-1}.

The recorded datum is $D_t^i = a_t^i + \varepsilon_t^{D,i}$, where the recording noise $\varepsilon_t^{D,i}\sim\mathcal{N}(0,\bar{N}/\gamma_D)$ is independent across workers. The scaling $\bar{N}/\gamma_D$ is the convention that each individual record uses a $1/\bar{N}$ share of the total data-recording precision γ_D—so that γ_D itself is the precision per unit mass of workers, a finite object in the continuum limit we take below.

Subtracting the known bracket in (34) from D_t^i and dividing through by φ_i yields a rescaled observation

$$\hat{D}_t^i \;\equiv\; \frac{D_t^i-\text{known}}{\varphi_i} \;=\; \theta_t \;+\; \frac{\xi_t^i}{\varphi_i} \;+\; \frac{\varepsilon_t^{D,i}}{\varphi_i}, \tag{35}$$

whose conditional variance given θ_t is the sum of two independent pieces:

$$\mathrm{Var}(\hat{D}_t^i\mid\theta_t) \;=\; \underbrace{\frac{1}{\varphi_i^2}\Big[(1-\alpha_t^i)^2\tfrac{(\gamma^l(\gamma_t^{H,i}))^2}{(\gamma_t^{H,i}+\gamma^l(\gamma_t^{H,i}))^2}\tfrac{1}{\gamma^l(\gamma_t^{H,i})}+(\alpha_t^i)^2\tfrac{(\gamma^a)^2}{(\gamma_t^{H,i}+Q_t+\gamma^a)^2}\tfrac{1}{\gamma^a}\Big]}_{\text{signal-noise variance, } O(1)\text{ in }\bar{N}} \;+\; \underbrace{\frac{1}{\varphi_i^2}\cdot\frac{\bar{N}}{\gamma_D}}_{\text{recording-noise variance}}. \tag{36}$$

Step 2: Aggregation and the $\bar{N}\to\infty$ limit.

Let G_t denote the cross-sectional distribution of expertise in period t, so that the $\bar{N}$ recorded workers are drawn i.i.d. from G_t (with policy $\alpha_t(\cdot)$). Group workers by expertise: conditional on θ_t, the empirical mean of $\hat{D}_t^i$ over workers with expertise in a small neighborhood $[\gamma^H,\gamma^H+d\gamma^H)$ is, by the LLN,

$$\bar{X}_t(\gamma^H) \;=\; \frac{1}{\bar{N}\,g_t(\gamma^H)\,d\gamma^H}\sum_{i:\,\gamma_t^{H,i}\in[\gamma^H,\gamma^H+d\gamma^H)}\hat{D}_t^i \;=\; \theta_t \;+\; \bar{u}_t(\gamma^H),$$

with conditional variance

$$\mathrm{Var}\big(\bar{u}_t(\gamma^H)\big|\theta_t\big) \;=\; \frac{1}{\bar{N}\, g_t(\gamma^H)\, d\gamma^H}\,\mathrm{Var}(\hat{D}_t^i \mid \theta_t).$$

Substituting (36), the first (signal-noise) piece contributes

$$\frac{1}{\bar{N}\, g_t(\gamma^H)\, d\gamma^H}\cdot\frac{O(1)}{\varphi(\gamma^H)^2} \;\xrightarrow[\bar{N}\to\infty]{}\; 0,$$

because it is $O(1/\bar{N})$, while the second (recording-noise) piece contributes

$$\frac{1}{\bar{N}\, g_t(\gamma^H)\, d\gamma^H}\cdot\frac{\bar{N}}{\gamma_D\,\varphi(\gamma^H)^2} \;=\; \frac{1}{\gamma_D\,\varphi(\gamma^H)^2\, g_t(\gamma^H)\, d\gamma^H},$$

independent of $\bar{N}$. The two explicit $\bar{N}$ factors cancel: the averaging (division by $\bar{N}g_t(\gamma^H)d\gamma^H$) exactly offsets the per-record recording variance ($\bar{N}/\gamma_D$), leaving a finite precision per unit mass.

Taking reciprocals, the local precision of $\bar{X}_t(\gamma^H)$ about θ_t in the limit is $\gamma_D\,\varphi(\gamma^H)^2\, g_t(\gamma^H)\, d\gamma^H$. Summing (integrating) over all expertise levels, and using that the $\bar{X}_t(\gamma^H)$ are conditionally independent across non-overlapping neighborhoods, the *total* precision gain about θ_t from period-t data is

$$\Phi_t \;=\; \gamma_D \int \varphi\big(\gamma^H,\, \alpha_t(\gamma^H);\, Q_t\big)^2\, dG_t(\gamma^H), \tag{37}$$

which is equation (24).

Step 3: Forward update.

The dataset's prior precision about θ_t carried over from Ω_{t-1} is Q_t, by definition. Adding the new precision Φ_t gives the period-t posterior precision:

$$\frac{1}{\mathrm{Var}(\theta_t \mid \Omega_t)} \;=\; Q_t + \Phi_t.$$

The underlying state follows the AR(1) $\theta_{t+1} = \rho\,\theta_t + \eta_{t+1}$ with $\mathrm{Var}(\eta_{t+1}) = 1/\gamma_\eta$, so the one-period-ahead posterior variance is

$$\mathrm{Var}(\theta_{t+1} \mid \Omega_t) \;=\; \rho^2\,\mathrm{Var}(\theta_t \mid \Omega_t) + \frac{1}{\gamma_\eta} \;=\; \frac{\rho^2}{Q_t+\Phi_t} + \frac{1}{\gamma_\eta},$$

and taking reciprocals gives

$$\frac{1}{Q_{t+1}} \;=\; \frac{\rho^2}{Q_t+\Phi_t} + \frac{1}{\gamma_\eta},$$

which is equation (25). □

A.4 Proof of Proposition 3 (Data Corruption)

The integrand of Φ_t in (24) is the squared combined innovation weight:

$$\left[(1-\alpha)\,\frac{\gamma^l(\gamma^H)}{\gamma^H+\gamma^l(\gamma^H)}+\alpha\,\frac{\gamma^a}{\gamma^H+Q+\gamma^a}\right]^2.$$

Since $\gamma^a < \gamma^l(\gamma^H)$ and $Q+\gamma^a > \gamma^l(\gamma^H)$ pointwise, we have $\gamma^a/(\gamma^H+Q+\gamma^a) < \gamma^l(\gamma^H)/(\gamma^H+\gamma^l(\gamma^H))$ for all $\gamma^H \geq 0$. Therefore the expression inside the brackets is strictly decreasing in α at every γ^H, and so is its square. Increasing α for any worker reduces Φ_t, and Q_{t+1} is strictly increasing in Φ_t through (25). □

B Proof of Propositions 4 and 5 (Auxiliary Bellman Equivalence)

B.1 Steady-state planner (Proposition 4)

The planner's problem. The planner chooses a policy $\alpha(\gamma^H)$, a distribution G, and aggregate data quality Q to maximize

$$\max_{\alpha(\cdot),\,G,\,Q}\int u\big(\gamma^H,\alpha(\gamma^H);Q\big)\,dG(\gamma^H)$$

subject to two constraints. The first is stationarity of the distribution: for every Borel set B,

$$G(B)=(1-\delta)\int \mathbf{1}\{H(\gamma^H,\alpha(\gamma^H))\in B\}\,dG(\gamma^H)+\delta\,F_0(B). \tag{38}$$

This says: the mass of workers in B next period (left side) equals the mass of surviving workers who transition into B under the policy α (first term on the right) plus the mass of new entrants arriving in B (second term). The second constraint is aggregate data quality consistency:

$$Q=\Gamma\big(Q,\,\Phi(\alpha,G,Q)\big),\qquad \Phi=\int\varphi(\gamma^H,\alpha(\gamma^H);Q)\,dG. \tag{39}$$

Lagrangian. Attach a multiplier $\Lambda(\gamma^{H\prime})$ to the stationarity constraint (38) for each skill level $\gamma^{H\prime}$, and a scalar multiplier μ to the data quality constraint (39). The stationarity constraint (38) is an equality of two measures: the left side is G and the right side is $T(G)\equiv(1-\delta)\int\mathbf{1}\{H(\gamma^H,\alpha)\in\cdot\}\,dG+\delta F_0$. Multiplying the constraint $T(G)-G=0$ by Λ and integrating over skill levels:

$$\mathcal{L}=\underbrace{\int u\,dG}_{\text{objective}}+\underbrace{\int\Lambda\,d\big[T(G)-G\big]}_{\text{stationarity}}+\underbrace{\mu\big[Q-\Gamma(Q,\Phi)\big]}_{\text{data quality}}. \tag{40}$$

We now expand $\int\Lambda\,d[T(G)-G]$ into three terms:

First term: $\int\Lambda\,d[T(G)]$ contains the inflow piece. For each worker at γ^H, the transition map

H sends her to $H(\gamma^H, \alpha)$ with survival probability $(1-\delta)$, plus entrants arrive from F_0 at rate δ. Integrating Λ against this inflow measure gives

$$\int \Lambda \, d[T(G)] = (1-\delta) \int \Lambda\big(H(\gamma^H, \alpha)\big) \, dG(\gamma^H) + \delta \int \Lambda \, dF_0.$$

The first part evaluates Λ at each worker's transition destination $H(\gamma^H, \alpha)$: this is the social value of where workers end up. The second part is the social value of entrants.

Second term: $-\int \Lambda \, dG = -\int \Lambda(\gamma^H) \, dG(\gamma^H)$. This is the social cost of the existing distribution: each worker at γ^H "uses up" a slot whose shadow value is $\Lambda(\gamma^H)$.

Combining:

$$\mathcal{L} = \int \big[\underbrace{u(\gamma^H, \alpha; Q)}_{\text{flow payoff}} + \underbrace{(1-\delta)\,\Lambda\big(H(\gamma^H, \alpha)\big)}_{\text{value of transition}} - \underbrace{\Lambda(\gamma^H)}_{\text{cost of slot}}\big] \, dG + \delta \int \Lambda \, dF_0 + \mu\big[Q - \Gamma(Q, \Phi)\big]. \quad (41)$$

Each worker at γ^H contributes three terms to the Lagrangian: her flow payoff u, the social value of the skill level she transitions to (weighted by the survival probability $1-\delta$), minus the shadow cost of the slot she occupies in the stationary distribution. The $\delta \int \Lambda \, dF_0$ term captures the value of entrants and does not depend on the planner's choices.

First-order condition for α. Differentiating (41) with respect to $\alpha(\gamma^H)$ at a point where $G(\gamma^H) > 0$:

$$\frac{\partial u}{\partial \alpha} + (1-\delta)\,\Lambda'\big(H\big)\,\frac{\partial H}{\partial \alpha} - \mu\,\frac{\partial \Gamma}{\partial \Phi}\,\frac{\partial \varphi}{\partial \alpha} = 0.$$

The last term arises because changing α at γ^H changes the data contribution φ, which changes Φ, which affects Q through Γ. Defining $\tilde{\mu} = \mu \cdot \partial\Gamma/\partial\Phi$ and rearranging:

$$\frac{\partial u}{\partial \alpha} + \tilde{\mu}\,\frac{\partial \varphi}{\partial \alpha} + (1-\delta)\,\Lambda'(H)\,\frac{\partial H}{\partial \alpha} = 0.$$

This is equation (27) in the main text.

First-order condition for G. Now ask: what happens when we add an infinitesimal mass of workers at skill γ^H? Differentiating (41) with respect to $G(\gamma^H)$, three effects arise.

(a) Direct welfare. The new worker produces flow payoff $u(\gamma^H, \alpha^*; Q)$. This is the direct benefit of having an additional worker at this skill level.

(b) Distribution composition. The new worker occupies a slot at γ^H in the stationary distribution, which has shadow cost $\Lambda(\gamma^H)$—this is what it "costs" the planner to place a worker here. With probability $(1-\delta)$ she survives and transitions to $H(\gamma^H, \alpha^*)$, which has shadow value $\Lambda(H)$—this is the benefit of the worker reaching a new skill level. The net distribution effect is $(1-\delta)\Lambda(H) - \Lambda(\gamma^H)$.

(c) Data quality. The new worker contributes $\varphi(\gamma^H, \alpha^*; Q)$ to the aggregate data flow $\Phi = \int \varphi \, dG$. This changes Φ, which changes Q through the law of motion $\dot{Q} = \Gamma(Q, \Phi)$. The composite derivative through this chain is

$$\frac{\partial}{\partial G(\gamma^H)}\big[\mu\big(\Gamma(Q,\Phi) - \dot{Q}\big)\big] = \mu \, \frac{\partial \Gamma}{\partial \Phi} \, \frac{\partial \Phi}{\partial G(\gamma^H)} = \mu \, \frac{\partial \Gamma}{\partial \Phi} \, \varphi(\gamma^H, \alpha^*; Q) = \tilde{\mu} \, \varphi(\gamma^H, \alpha^*; Q),$$

where $\tilde{\mu} = \mu \cdot \partial\Gamma/\partial\Phi$ converts the aggregate multiplier μ on the Q constraint into a per-worker data contribution price. Note that this composite derivative fully accounts for the chain $G \to \Phi \to Q$: all further effects of Q on Λ or on other parts of the Lagrangian are already handled by the separate FOC for Q (which pins down μ).

Setting the sum of effects (a), (b), and (c) to zero:

$$\underbrace{u}_{\text{(a)}} + \underbrace{(1-\delta)\,\Lambda(H) - \Lambda(\gamma^H)}_{\text{(b)}} + \underbrace{\tilde{\mu}\,\varphi}_{\text{(c)}} = 0.$$

Rearranging:

$$\Lambda(\gamma^H) = u(\gamma^H, \alpha^*; Q) + \tilde{\mu}\,\varphi(\gamma^H, \alpha^*; Q) + (1-\delta)\,\Lambda\big(H(\gamma^H, \alpha^*)\big).$$

This is equation (28) in the main text. In words: the social value of a worker at γ^H (left side) equals her total per-period contribution—private output u plus data contribution $\tilde{\mu}\varphi$—plus the discounted social value of the skill she will reach next period, $(1-\delta)\Lambda(H)$. Combining with the FOC for α gives the auxiliary Bellman (26). □

B.2 Dynamic planner

Proposition 5 (Dynamic Auxiliary Bellman)**.** *The dynamic planner's problem is equivalent to the following transformed problem. There exists a sequence of shadow prices $\{\tilde{\mu}_t\}_{t=0}^{\infty}$ such that the planner's optimal policy $\{\alpha_t^{plan}(\gamma^H)\}$ and the associated paths $\{Q_t, G_t\}$ satisfy:*

(i) *At each date t, the policy $\alpha_t(\gamma^H)$ solves the auxiliary Bellman*

$$\Lambda_t(\gamma^H) = \max_{\alpha \in [0,1]} \Big[u(\gamma^H, \alpha; Q_t) + \tilde{\mu}_t \, \varphi(\gamma^H, \alpha; Q_t) + \beta_s (1-\delta)\, \Lambda_{t+1}\big(H(\gamma^H, \alpha)\big) \Big]. \tag{42}$$

(ii) *The distribution evolves according to the distribution dynamics and data quality according to the law of motion under the policy $\alpha_t(\cdot)$.*

(iii) *The shadow prices satisfy the backward recursion*

$$\tilde{\mu}_{t-1} = \beta_s \int \frac{\partial u}{\partial Q} \, dG_t + \tilde{\mu}_t \left[\frac{\partial \Gamma}{\partial Q} + \frac{\partial \Gamma}{\partial \Phi} \frac{\partial \Phi}{\partial Q} \right], \tag{43}$$

where the derivatives are evaluated at the date-t allocation.

Lagrangian. The dynamic planner maximizes $\sum_{t=0}^{\infty} \beta_s^t \int u \, dG_t$ subject to date-by-date distribution dynamics and data quality dynamics. Attach multipliers $\lambda_t(\gamma^H)$ to the distribution constraint at each (t, γ^H) and μ_t to the data quality constraint at each t. By the same logic as the steady-state case, the FOC for $\alpha_t(\gamma^H)$ gives

$$\beta_s^t \frac{\partial u}{\partial \alpha} + \beta_s^t \tilde{\mu}_t \frac{\partial \varphi}{\partial \alpha} + (1-\delta) \lambda_{t+1}(H) \frac{\partial H}{\partial \alpha} = 0,$$

where $\tilde{\mu}_t = \mu_t (\partial \Gamma_t / \partial \Phi_t) / \beta_s^t$. The β_s^t premultiplies u and $\tilde{\mu}_t \varphi$ because these are date-t payoffs weighted by β_s^t in the objective. The continuation term uses λ_{t+1} (the date-$(t+1)$ multiplier) without β_s^t because the stationarity constraint at $t+1$ is not weighted by β_s.

Rescaling. Define $\Lambda_t(\gamma^H) = \lambda_t(\gamma^H)/\beta_s^t$. The FOC for $G_t(\gamma^H)$ gives

$$\Lambda_t(\gamma^H) = u + \tilde{\mu}_t \varphi + \frac{\beta_s^{t+1}}{\beta_s^t} (1-\delta) \Lambda_{t+1}(H) = u + \tilde{\mu}_t \varphi + \beta_s (1-\delta) \Lambda_{t+1}(H).$$

The ratio $\beta_s^{t+1}/\beta_s^t = \beta_s$ is what introduces the planner's time discount into the continuation value. This is the auxiliary Bellman (42). □

C Welfare Arithmetic: Flow vs. Lifetime

The welfare gain from planner intervention reported in Table 3 differs slightly between the flow measure and the lifetime measure, and the direction of the gap reverses across the two entry distributions: under uniform entry the lifetime gain (+7.2%) exceeds the flow gain (+5.9%), while under bottom entry the flow gain (+12.3%) exceeds the lifetime gain (+11.8%). This appendix derives the exact decomposition and explains the reversal.

Three welfare measures. Throughout the paper we report three measures of welfare in any given steady state:

(i) cross-sectional flow welfare $\int u \, dG$;

(ii) cross-sectional lifetime welfare $\int V \, dG$;

(iii) new-entrant lifetime welfare $\int V \, dF_0$.

We now derive an identity that ties these three measures together exactly.

Derivation of the identity. Let $u(\gamma^H, \alpha; Q)$ denote the flow payoff and let $V(\gamma^H)$ be the lifetime value satisfying the Bellman

$$V(\gamma^H) \;=\; u\big(\gamma^H, \alpha^*(\gamma^H); Q^*\big) \;+\; \beta(1-\delta)\, V\big(H(\gamma^H, \alpha^*(\gamma^H))\big). \tag{44}$$

The factor $(1-\delta)$ is the survival probability: with probability δ the worker exits and her continuation value is zero (she is replaced by a fresh entrant who is a separate agent). Integrating (44) over the stationary cross-section G gives

$$\int V\, dG \;=\; \int u\, dG \;+\; \beta(1-\delta) \int V\big(H(\gamma^H)\big)\, dG. \tag{45}$$

Let $T_* G$ denote the push-forward of G under H—the distribution of next-period skill among surviving workers. Stationarity of the cross-section requires

$$G \;=\; (1-\delta)\, T_* G \;+\; \delta\, F_0, \tag{46}$$

since a fraction $1-\delta$ of next-period workers are survivors transitioned by H and a fraction δ are fresh entrants. Solving (46) for $T_* G$ and integrating V against it,

$$\int V\, dT_* G \;=\; \frac{\int V\, dG \;-\; \delta \int V\, dF_0}{1-\delta}.$$

Substituting into (45) and collecting terms,

$$\int V\, dG \;=\; \frac{\int u\, dG \;-\; \beta\delta \int V\, dF_0}{1-\beta}. \tag{47}$$

This pins the three welfare measures to each other: any two determine the third, and the discount factor that converts flow to lifetime is $1/(1-\beta)$, not $1/[1-\beta(1-\delta)]$. The entry term $\beta\delta \int V\, dF_0$ enters because workers die at rate δ and are replaced by entrants drawn from F_0; the cross-sectional lifetime measure incorporates this replacement, while the entrant value $\int V\, dF_0$ does not.

Limiting case: no turnover. When $\delta = 0$ workers are infinitely lived. The entry term vanishes from (47), leaving

$$\int V\, dG \;=\; \frac{\int u\, dG}{1-\beta}.$$

Cross-sectional lifetime welfare is then a pure rescaling of cross-sectional flow welfare. Any percentage change in $\int u\, dG$ translates one-for-one into the same percentage change in $\int V\, dG$, so flow gain $\equiv$ lifetime gain. *The entire flow-vs-lifetime gap reported in this paper is therefore an entry-cohort phenomenon*: it arises because new workers continually enter at F_0, and a planner who improves the entrant's lifetime value $\int V\, dF_0$ pulls the cross-sectional lifetime welfare $\int V\, dG$ along with it through the $-\beta\delta \int V\, dF_0$ term in (47).

Decomposing the welfare gain. Write u_C, W_C, E_C for the CE values of $\int u\, dG$, $\int V\, dG$, $\int V\, dF_0$, and u_P, W_P, E_P for the corresponding planner values. The flow gain is $\Delta u/|u_C|$ with $\Delta u = u_P - u_C$. Substituting (47) into the lifetime gain gives

$$\frac{\Delta W}{|W_C|} = \frac{\Delta u - \beta\delta\, \Delta E}{|u_C - \beta\delta\, E_C|}, \tag{48}$$

where $\Delta E = E_P - E_C$. The lifetime gain has two adjustments relative to the flow gain. The numerator subtracts $\beta\delta\, \Delta E$: a planner who improves the lifetime value of fresh entrants by $\Delta E > 0$ shrinks the lifetime numerator. The denominator subtracts $\beta\delta\, E_C$: when $E_C < 0$ (entrants face a negative lifetime value under the CE), the lifetime denominator $|u_C - \beta\delta\, E_C|$ is smaller than $|u_C|$.

Why the gap reverses across entry distributions. Table 5 reports the components of (48) at the calibrated economy.

Table 5: Decomposition of flow and lifetime welfare gains.

	Uniform entry	Bottom entry
$u_C = \int u\, dG_{\text{CE}}$	-0.0314	-0.0363
$E_C = \int V_{\text{CE}}\, dF_0$	-0.4349	-0.5168
$E_P = \int V_{\text{plan}}\, dF_0$	-0.4166	-0.4500
Flow gain $\Delta u/\lvert u_C\rvert$	$+5.85\%$	$+12.31\%$
Entry-cohort gain $\Delta E/\lvert E_C\rvert$	$+4.27\%$	$+12.93\%$
Lifetime gain (48)	$+7.16\%$	$+11.81\%$

The crucial row is *entry-cohort gain*—how much the planner improves the lifetime value of a fresh entrant. Under bottom entry, every entrant arrives at $\gamma_0^H = 13.5$ where the planner's correction operates with maximum force; the planner pulls each entrant out of the trap she would otherwise be locked into, raising entrant lifetime value by $+12.9\%$, slightly above the cross-sectional flow gain of $+12.3\%$. The drag $\beta\delta\, \Delta E$ in the numerator of (48) therefore exceeds the proportional shrinkage of the denominator, and the lifetime gain falls below the flow gain. Under uniform entry, fresh entrants are spread across $[\gamma_0^H, \bar{\kappa}]$, including some who would already reach the expert level even under the CE; the planner's improvement in entrant lifetime value is only $+4.3\%$, well below the cross-sectional flow gain of $+5.9\%$. The numerator drag is small, the denominator

shrinkage dominates, and the lifetime gain rises above the flow gain.

The gap reversal is therefore driven by which channel the planner reaches more strongly: the cross-sectional flow welfare (the uniform-entry case, where the planner reshapes the stationary distribution but doesn't dramatically affect entrants on average) or the entry-cohort lifetime value (the bottom-entry case, where every entrant is at the planner's most aggressive intervention point). Identity (47) holds to four decimal places at both the CE and the planner allocation, under both entry distributions.

D Robustness

D.1 Skill trap across AI prices

We report full CE-and-planner outcomes as the AI price varies around the benchmark $p_A = 0.007$ to $\{0.010, 0.015, 0.020\}$, holding all other parameters at their baseline values and re-pinning γ_D so that $Q_{\text{pre}} = \bar{\kappa} + \gamma^l(\bar{\kappa})$ is unchanged. Table 6 reports the results; Figure 9 shows the full allocations.

Table 6: Robustness: varying the AI price p_A

	Uniform entry				Bottom entry			
	$p_A = 0.007$	0.010	0.015	0.020	$p_A = 0.007$	0.010	0.015	0.020
Competitive equilibrium								
Data quality Q_{CE}	22.6	26.8	31.6	32.9	18.1	23.0	28.0	29.9
Mean AI intensity	0.73	0.47	0.12	0.05	0.88	0.64	0.30	0.17
Mean skill	20.3	23.0	26.6	27.4	16.8	19.2	22.5	23.8
Planner								
Data quality Q_{plan}	28.8	31.3	32.7	33.2	27.4	28.7	30.1	30.9
Mean AI intensity	0.40	0.15	0.07	0.04	0.39	0.27	0.17	0.11
Flow gain	+5.9%	+3.1%	+0.4%	+0.1%	+12.3%	+5.8%	+1.5%	+0.6%
Lifetime gain	+7.2%	+4.0%	+0.8%	+0.3%	+11.8%	+6.6%	+2.1%	+1.0%

The baseline column ($p_A = 0.007$) reproduces the headline planner-vs-CE comparison from Section 6; the higher columns trace what happens as the AI price rises above its calibrated value. A higher price endogenously pushes workers out of full delegation: α_{CE} falls from 0.73 at the baseline to 0.05 at $p_A = 0.020$ under uniform entry, and from 0.88 to 0.17 under bottom entry. CE data quality recovers in step, and the planner's optimal subsidy $\tilde{\mu}^*$ shrinks to internalize the smaller residual externality. Flow welfare gains decline monotonically with the price—from $+5.9\%$ to $+0.1\%$ under uniform entry, and from $+12.3\%$ to $+0.6\%$ under bottom entry—confirming that the externality is largest precisely where AI is cheapest, the same region in which the skill trap is most binding. At a high enough price, p_A alone implements something close to the planner's allocation; at the calibrated price, it does not.

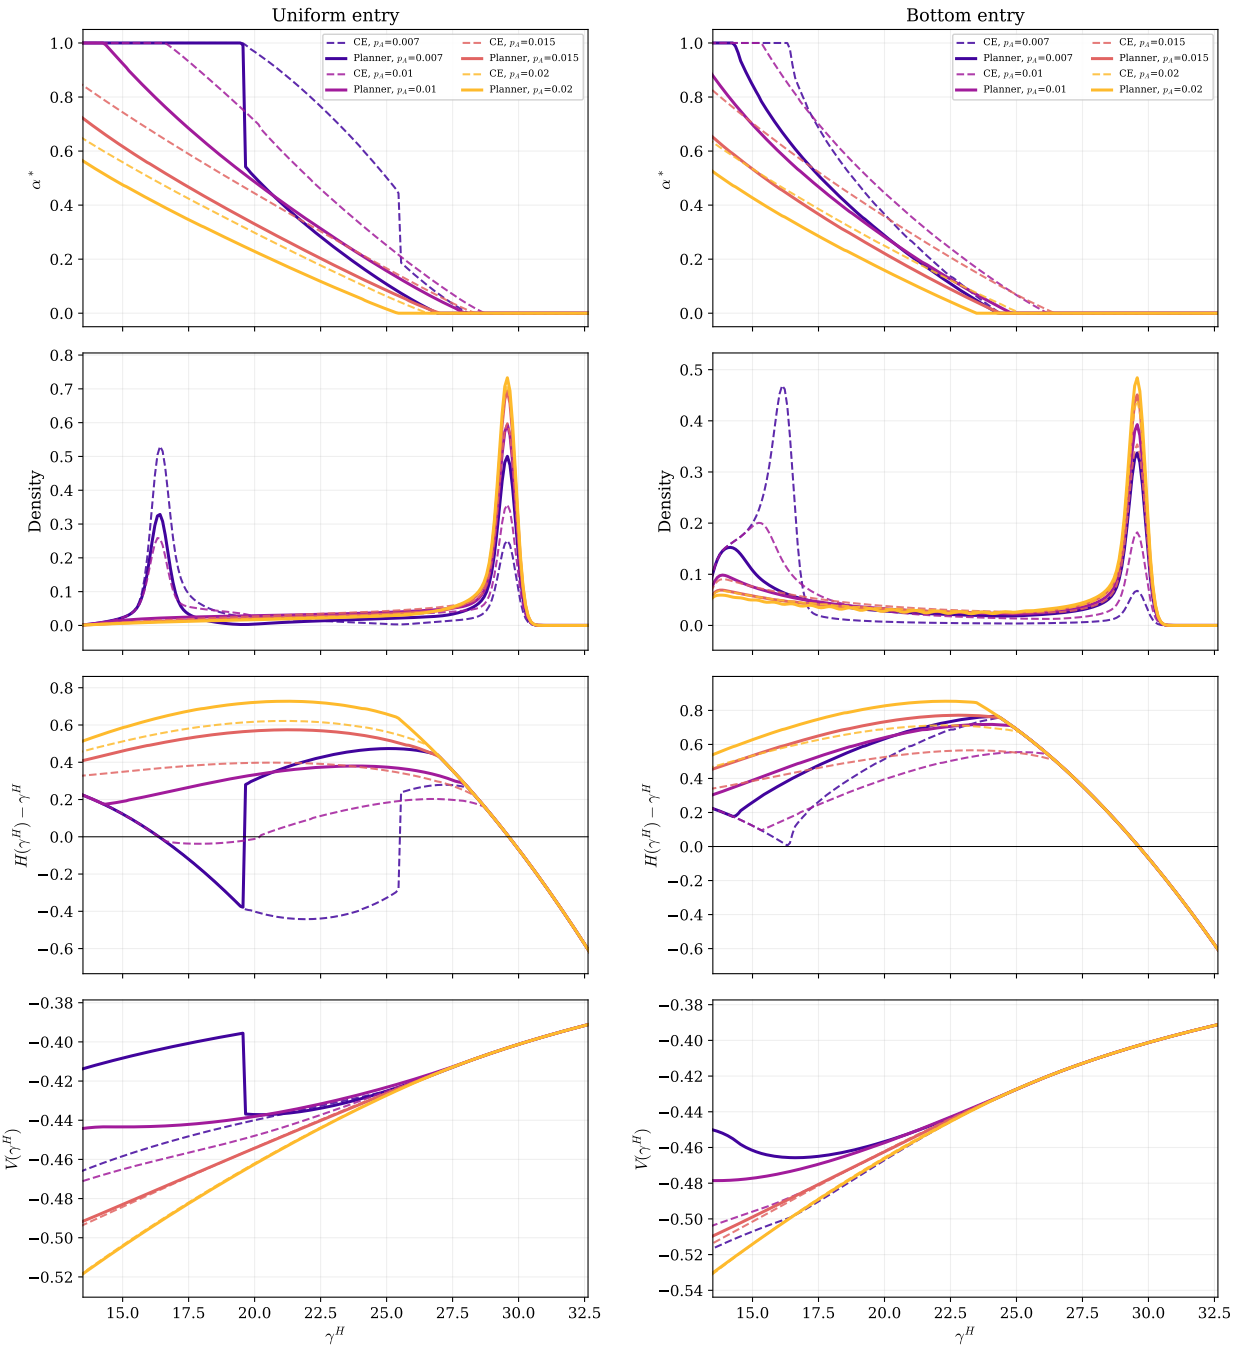

Figure 9: Planner allocations across the AI price $p_A \in \{0.007, 0.010, 0.015, 0.020\}$. Rows: AI intensity $\alpha^*(\gamma^H)$; stationary skill density; expected one-step skill change $H(\gamma^H) - \gamma^H$; lifetime value $V(\gamma^H)$. Columns: uniform entry (left), bottom entry (right). Black solid curves show the CE allocation at the baseline ($p_A = 0.007$); coloured curves show the planner's policy at each value of p_A (light to dark = 0.007, 0.010, 0.015, 0.020).

D.2 Varying the AI/labor learning ratio (c)

The deskilling assumption $\tilde{\gamma}^a(\gamma^H) < \gamma^l(\gamma^H)$ (equivalently $c < 1$) is central to the skill-trap mechanism. We vary the ratio $c = \tilde{\gamma}^a(\gamma^H)/\gamma^l(\gamma^H)$ from the benchmark 0.25 up to 1.00, including the limiting case $c = 1$ in which AI delivers exactly the same learning as labor and the deskilling assumption fails outright. We hold all other parameters at their baseline values and re-pin γ_D to keep Q_{pre} fixed. Table 7 reports the results.

Table 7: Robustness: varying the AI/labor learning ratio c

	Uniform entry				Bottom entry			
	$c = 0.25$	0.50	0.75	1.00	$c = 0.25$	0.50	0.75	1.00
Competitive equilibrium								
Data quality Q_{CE}	22.6	20.4	18.5	17.6	18.1	18.9	17.9	17.2
Mean AI intensity	0.73	0.79	0.71	0.75	0.88	0.84	0.79	0.80
Mean skill	20.3	22.2	25.3	27.9	16.8	19.7	22.7	25.4
Planner								
Data quality Q_{plan}	28.8	28.2	27.1	25.8	27.4	27.0	26.7	26.0
Mean AI intensity	0.40	0.41	0.43	0.46	0.39	0.46	0.50	0.52
Flow gain	+5.9%	+6.9%	+7.0%	+6.4%	+12.3%	+9.0%	+8.8%	+8.0%
Lifetime gain	+7.2%	+6.8%	+5.1%	+4.4%	+11.8%	+6.8%	+5.2%	+4.0%

Under bottom entry the flow welfare gain declines as c rises (from +12.3% at the benchmark to +8.0% at $c = 1$): faster AI-assisted learning shrinks the deskilling penalty for low-skill workers and so the part of the externality attributable to skill loss diminishes. Strikingly, however, even at $c = 1$—where the deskilling assumption fails completely and labor and AI build skill at identical rates—the planner still produces single-digit-percent welfare gains. This is because c does not enter Φ: the data-corruption channel survives independently of the learning channel. The planner's lower mean AI intensity (0.46–0.52) preserves a higher flow of novel data even when no skill-acquisition argument applies. The skill trap, by contrast, does dissolve at $c = 1$ in the sense

that mean skill becomes equal under CE and planner (27.9 uniform, 25.4 bottom in both columns); the welfare gain is therefore entirely a data-quality gain, not a skill-distribution gain. This isolates the role of the data externality from the role of deskilling.

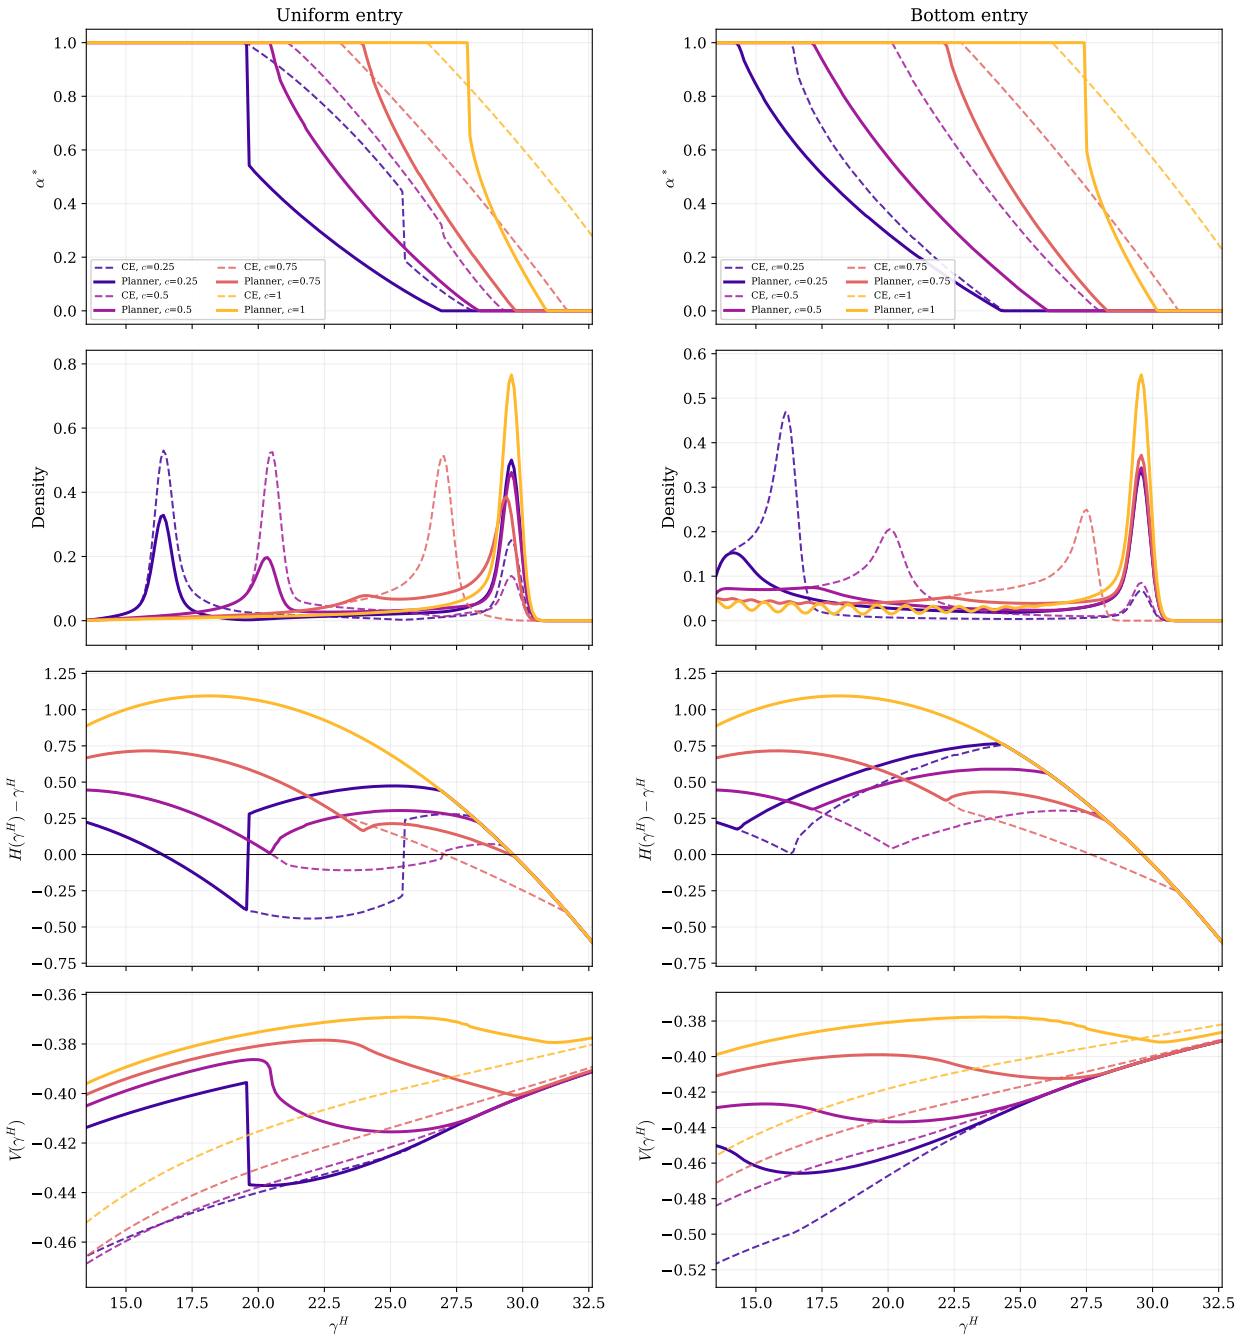

Figure 10: Planner allocations across $c \in \{0.25, 0.50, 0.75, 1.00\}$. Same panel layout as Figure 9; black solid curves show the CE at the baseline ($c = 0.25$); coloured curves show the planner's policy at each value of c (light to dark).

D.3 Varying the turnover rate (δ)

The benchmark $\delta = 1/30 \approx 0.033$ matches the 30-year cognitive career horizon implied by Lagakos et al. (2018). We test sensitivity at $\delta \in \{1/30, 1/20, 1/15, 1/10\}$, i.e., effective career lengths from 30 down to 10 years. γ_D is re-pinned and the discount factor $\beta(1-\delta)$ updates accordingly.

Table 8: Robustness: varying the turnover rate δ

	Uniform entry				Bottom entry			
	$\delta = 1/30$	$1/20$	$1/15$	$1/10$	$\delta = 1/30$	$1/20$	$1/15$	$1/10$
Competitive equilibrium								
Data quality Q_{CE}	22.6	21.9	21.3	20.3	18.1	16.5	15.2	13.5
Mean AI intensity	0.73	0.75	0.77	0.80	0.88	0.92	0.96	1.00
Mean skill	20.3	20.2	20.2	20.2	16.8	16.0	15.4	14.7
Planner								
Data quality Q_{plan}	28.8	28.4	27.6	26.7	27.4	25.7	24.3	22.2
Mean AI intensity	0.40	0.41	0.46	0.50	0.39	0.45	0.49	0.55
Flow gain	+5.9%	+5.9%	+6.0%	+6.4%	+12.3%	+13.1%	+13.6%	+13.7%
Lifetime gain	+7.2%	+7.6%	+7.7%	+8.2%	+11.8%	+12.8%	+13.4%	+13.4%

Shorter careers (higher δ) push CE workers toward fuller AI delegation under bottom entry—α_{CE} rises from 0.88 to 1.00 as δ rises from $1/30$ to $1/10$—because workers internalize less of the future return to expertise. CE data quality falls in step. The planner's welfare gain is robust across the entire career-length range: it stays in the $+6\%$ band under uniform entry and rises modestly from $+12.3\%$ to $+13.7\%$ under bottom entry. The data externality is a pricing failure that does not depend on the agent's horizon—workers at any career length take aggregate Q as given. The bimodal stationary distribution and the skill trap survive at every value of δ tested.

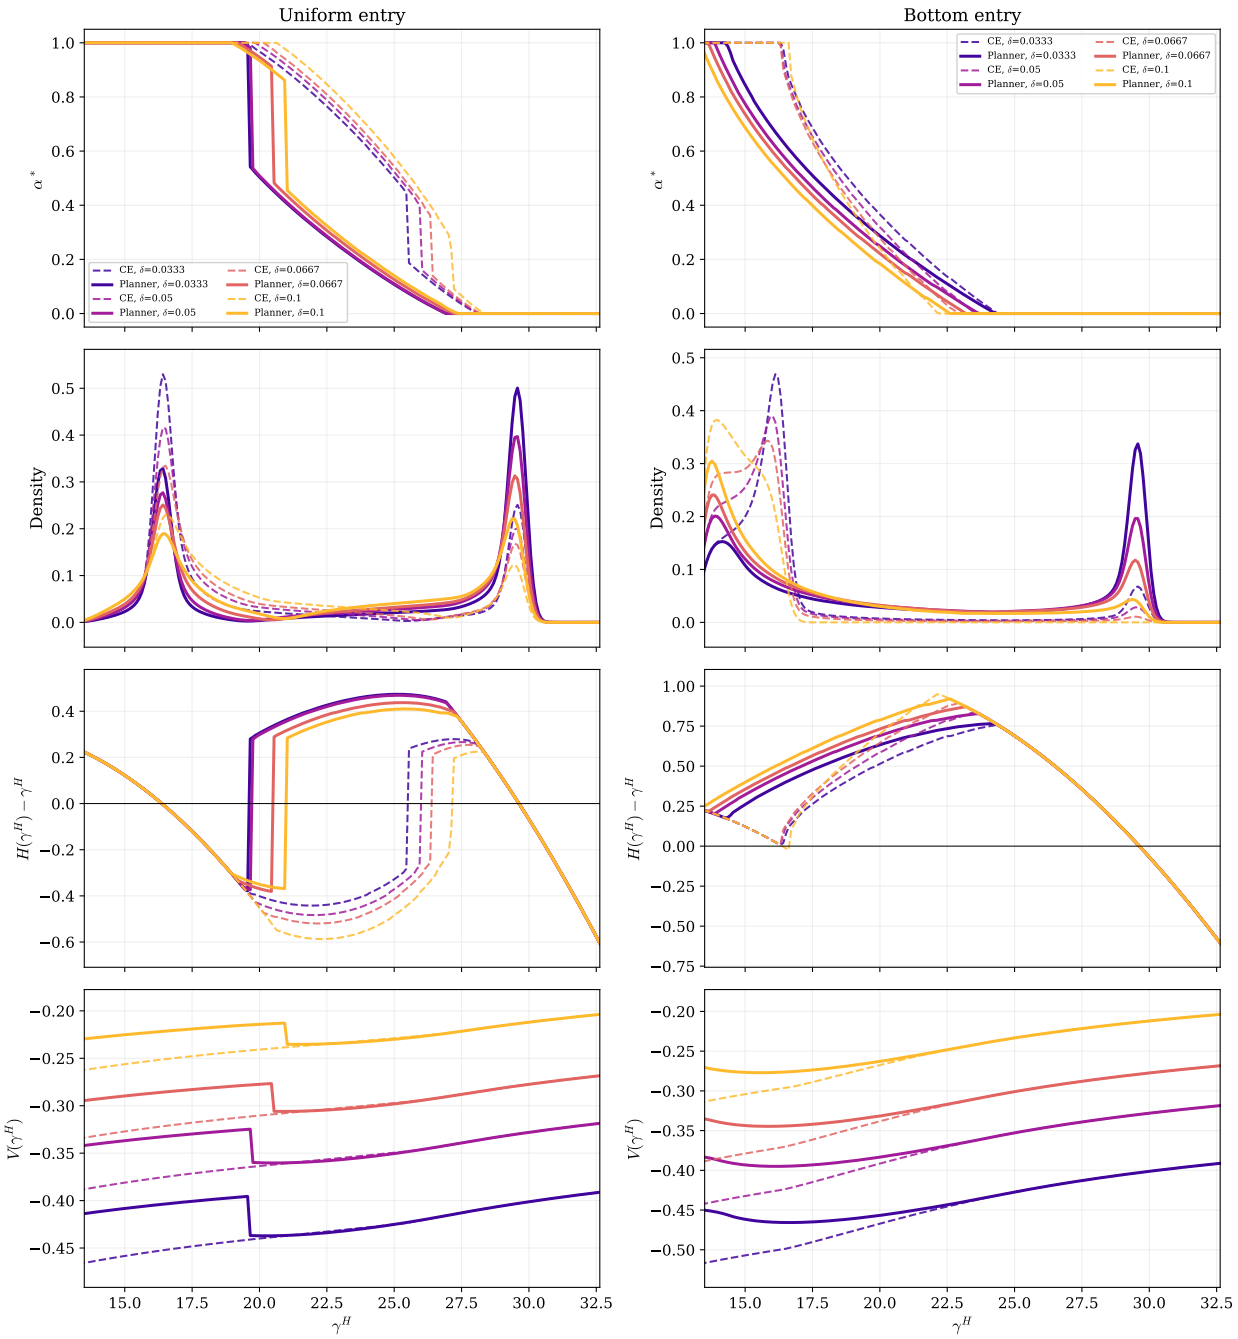

Figure 11: Planner allocations across $\delta \in \{1/30, 1/20, 1/15, 1/10\}$ (30-, 20-, 15-, 10-year careers). Same panel layout as Figure 9; black solid curves show the CE at the baseline ($\delta = 1/30$); coloured curves show the planner's policy at each value of δ (light to dark).

D.4 Varying the AI signal precision (γ^a)

γ^a governs how much novel information AI generates. The benchmark $\gamma^a = 0.08$ is small relative to the entry-level labor signal $\gamma^l(\gamma_0^H) = 1.0$ (and even smaller relative to $\gamma^l(\gamma^H)$ at higher expertise, since γ^l rises with γ^H), so AI produces little fresh signal per period. We vary $\gamma^a \in \{0.08, 0.20, 0.50, 1.00\}$, ending at $\gamma^a = 1.0$ where AI's fresh signal precision equals the entry-level labor signal precision—a regime that contradicts the model collapse evidence but provides a sanity check on whether the externality survives even when AI is informationally as rich as a novice's labor.

Table 9: Robustness: varying the AI signal precision γ^a

	Uniform entry				Bottom entry			
	$\gamma^a = 0.08$	0.20	0.50	1.00	$\gamma^a = 0.08$	0.20	0.50	1.00
Competitive equilibrium								
Data quality Q_{CE}	22.6	22.5	22.3	22.0	18.1	18.0	17.7	17.2
Mean AI intensity	0.73	0.73	0.75	0.78	0.88	0.89	0.91	0.94
Mean skill	20.3	20.3	20.1	19.8	16.8	16.7	16.5	16.1
Planner								
Data quality Q_{plan}	28.8	29.1	28.8	29.0	27.4	27.3	27.3	27.2
Mean AI intensity	0.40	0.38	0.40	0.40	0.39	0.39	0.40	0.42
Flow gain	+5.9%	+5.9%	+6.1%	+6.3%	+12.3%	+12.5%	+13.0%	+13.8%
Lifetime gain	+7.2%	+7.2%	+7.4%	+7.7%	+11.8%	+12.0%	+12.4%	+13.2%

The welfare gains do not shrink with γ^a—and in fact rise modestly across a 12-fold variation, from +5.9% to +6.3% under uniform entry and from +12.3% to +13.8% under bottom entry. Two channels explain why γ^a is not the parameter that drives the externality. First, AI's contribution to Φ_t carries innovation weight $\gamma^a/(\gamma^H + Q + \gamma^a)$, which remains small whenever Q is large; with $Q \approx$ 22–29 in the calibration, even $\gamma^a = 1.0$ leaves the AI weight an order of magnitude below the labor weight $\gamma^l(\gamma^H)/(\gamma^H + \gamma^l(\gamma^H))$. The dataset's existing precision dilutes any fresh

AI signal, so labor remains substantially more informative per action across the entire range of γ^a considered. Second, γ^a governs AI's *forecast* precision but not its *learning* rate: skill accumulation in (10) depends on $\tilde{\gamma}^a(\gamma^H)$, the rate at which AI-assisted work builds expertise, while γ^a is the precision of the contemporaneous AI signal entering the worker's action. Raising γ^a thus makes AI a better forecaster but no better a teacher; the gap $\gamma^l(\gamma^H) - \tilde{\gamma}^a(\gamma^H)$ that drives skill loss when α rises is unchanged. If anything, higher γ^a worsens the CE outcome—workers find AI more attractive (mean α rises from 0.73 to 0.78 under uniform entry, 0.88 to 0.94 under bottom entry), Q falls, and mean skill falls. The planner allocation barely moves, so the gap widens. The externality is therefore robust to—and even slightly amplified by—improvements in raw AI precision; what would shrink the planner's correction is a higher $\tilde{\gamma}^a$ (AI as a teacher), not a higher γ^a (AI as a forecaster). The bimodal stationary distribution and the skill trap survive at every value tested.

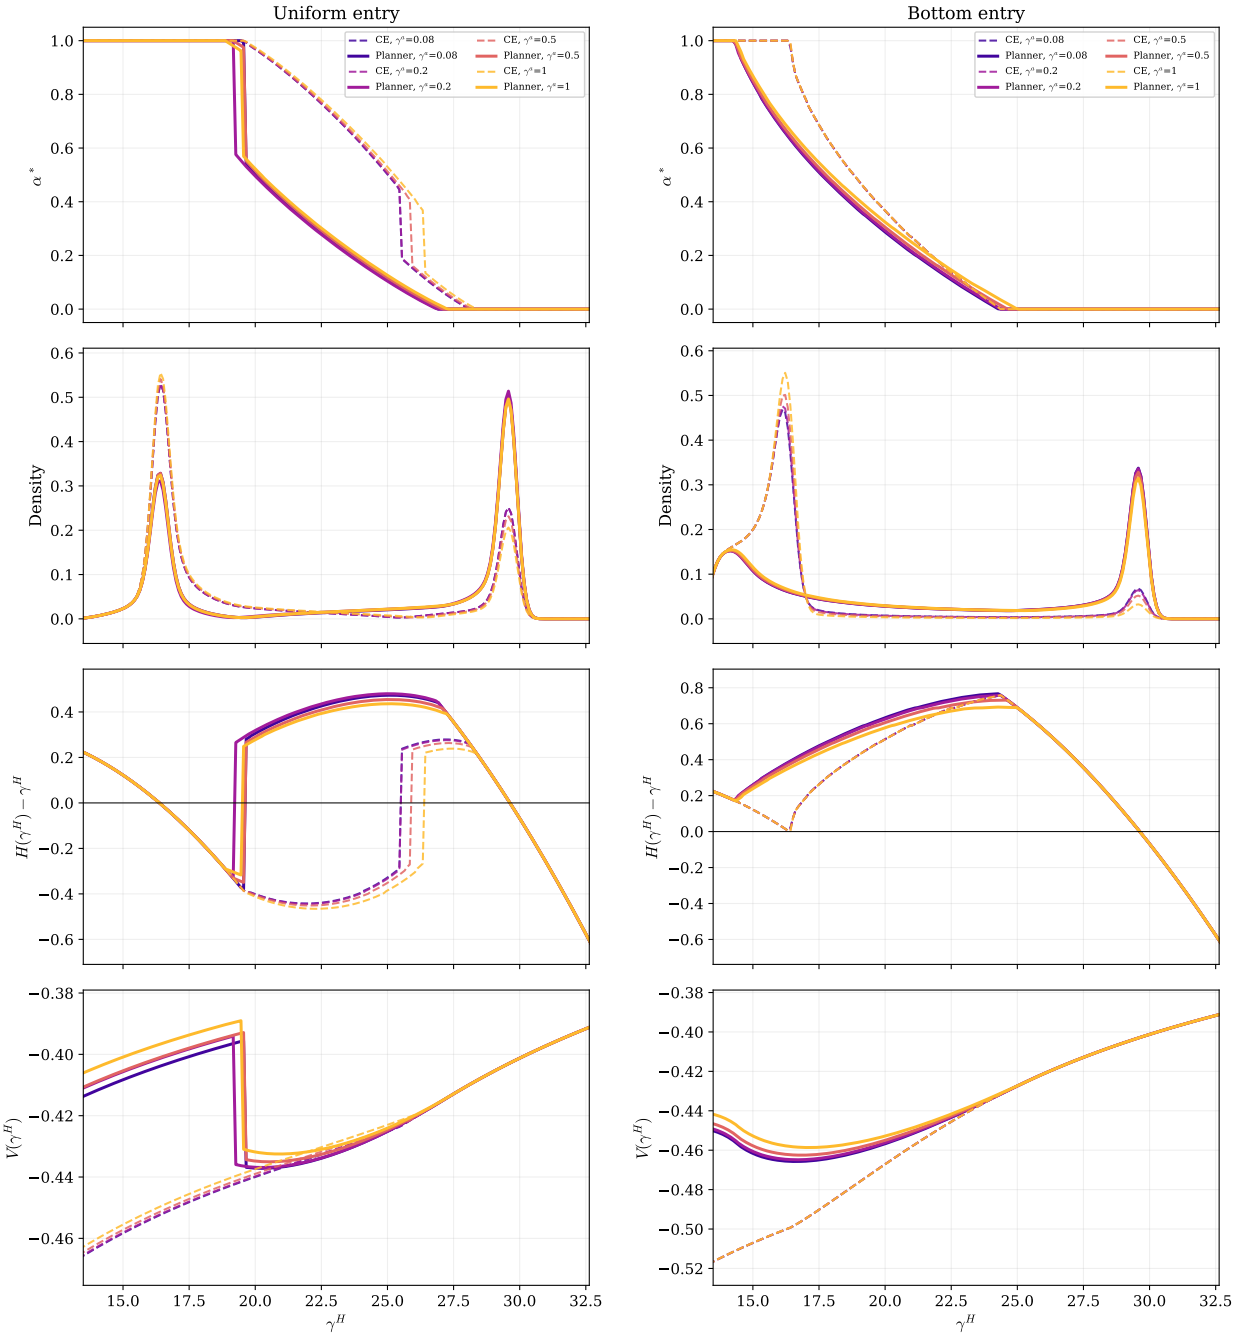

Figure 12: Planner allocations across $\gamma^a \in \{0.08, 0.20, 0.50, 1.00\}$. Same panel layout as Figure 9; black solid curves show the CE at the baseline ($\gamma^a = 0.08$); coloured curves show the planner's policy at each value of γ^a (light to dark).

D.5 Learning Synergy

The model treats the learning input as linear in AI intensity, $\ell(\alpha;\gamma^H) = (1-\alpha)\gamma^l(\gamma^H) + \alpha\tilde{\gamma}^a(\gamma^H)$. This is the simplest specification consistent with the deskilling premise of the paper—under (11), $\tilde{\gamma}^a(\gamma^H) = c\,\gamma^l(\gamma^H)$ with $c < 1$, so pure labor produces stronger learning signals than pure AI use—but it imposes monotonicity on the interior of $\alpha \in (0,1)$, ruling out the possibility that a moderate mix of labor and AI accelerates learning beyond either endpoint.

The empirical case for an interior synergy is weak. The literature on AI's effect on independent skill acquisition is dominated by negative findings: AI access impairs learning in educational settings (Bastani et al., 2025; Jošt et al., 2024; Barcaui, 2025), slows skill acquisition among software developers (Shen and Tamkin, 2026), and reduces diagnostic performance among physicians who rely on AI (Budzyń et al., 2025). The most rigorous recent productivity study of experienced AI users finds delays rather than gains (METR, 2025). A smaller "scaffolded use" literature in education suggests that guided AI use can complement learning, but no clean external moment is available to discipline a magnitude. We therefore treat learning synergy as a robustness parameter and study how its strength affects the externality result, rather than calibrating it to a target.

We augment the learning input with an asymmetric concavity term:

$$\ell(\alpha;\gamma^H) = (1-\alpha)\,\gamma^l(\gamma^H) + \alpha\,\tilde{\gamma}^a(\gamma^H) + \lambda\,\alpha(1-\alpha)^2. \tag{49}$$

The synergy term $\lambda\,\alpha(1-\alpha)^2$ vanishes at both endpoints—so $\ell(0;\gamma^H) = \gamma^l(\gamma^H)$ and $\ell(1;\gamma^H) = \tilde{\gamma}^a(\gamma^H)$ are preserved—and is asymmetric in α: it peaks at $\alpha = 1/3$ rather than at $\alpha = 1/2$ and dies quickly as α approaches one. This shape is meant to capture an empirically natural feature: a worker who delegates almost everything to AI gets "stuck"—there is no synergy bonus at very high α—while moderate use of AI can deliver a learning bonus when combined with substantial labor effort. A symmetric form $\lambda\alpha(1-\alpha)$ would predict that even at $\alpha = 0.99$ the worker still receives a non-trivial bonus, which is harder to defend.

The data-innovation weight φ keeps its standard form from (24); we abstract from a parallel synergy term in φ because such a term would mechanically magnify the externality—it would enter only through Φ_t and would therefore leave the CE policy unchanged while the planner internalizes it—without offering any new economic insight.

We solve the model at $\lambda \in \{0, 0.5, 1.0, 2.0\}$ for both entry distributions; $\lambda = 0$ recovers the

baseline of Table 3. The headline value $\lambda = 2.0$ produces an interior learning peak at $a^* \approx 0.18$ at the entry expertise $\gamma^H = \gamma_0^H$, with a $\sim 11\%$ gain over pure labor—an aggressive but illustrative value chosen to make the synergy mechanism visible against the dominant deskilling channel. $\lambda = 0.5$ is essentially below the threshold for an interior peak ($\lambda < 1$ implies no interior maximum at $\gamma^H = \gamma_0^H$ given $c = 0.25$), and $\lambda = 1.0$ produces a small interior peak ($\sim 0.8\%$) close to the corner. Figure 13 shows $\ell(a)$ at each λ. Because $\gamma^l(\gamma^H)$ rises with γ^H while the synergy term does not scale with γ^l, synergy is relatively largest at the entry expertise and fades toward the upper end of the skill range, consistent with the scaffolded-use intuition that novices are the most likely beneficiaries of guided AI.

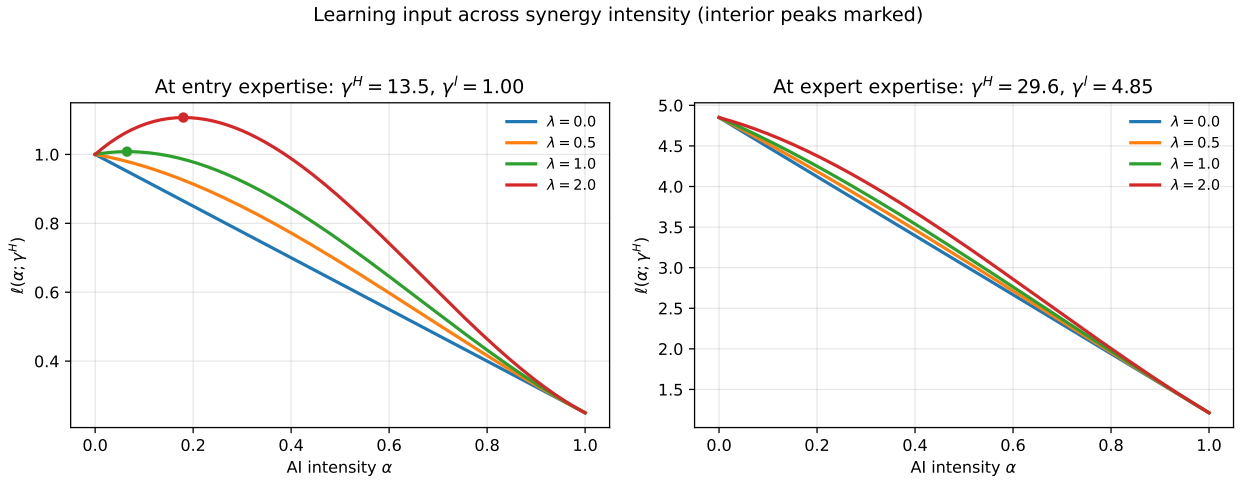

Figure 13: Learning input $\ell(a; \gamma^H)$ across learning-synergy intensity λ under the asymmetric form (49). *Left:* entry expertise $\gamma^H = \gamma_0^H = 13.5$, where $\gamma^l(\gamma_0^H) = 1.00$. Interior peaks (marked) emerge for $\lambda \gtrsim 1$, located at low a (roughly $1/3$ at $\lambda \to \infty$) and dying as $a \to 1$. *Right:* expert expertise $\gamma^H = \bar{\kappa} = 29.6$, where $\gamma^l(\bar{\kappa}) = 4.85$. Synergy is small relative to γ^l here, so curves are nearly linear. The endpoint values $\ell(0; \gamma^H) = \gamma^l(\gamma^H)$ and $\ell(1; \gamma^H) = \tilde{\gamma}^a(\gamma^H) = c\,\gamma^l(\gamma^H)$ are common across λ by construction.

Three observations stand out from Table 10. First, the externality survives across the entire range of λ: planner intervention delivers a positive welfare gain in every case, ranging from $+4.9\%$ (uniform, $\lambda = 2$) to $+12.4\%$ (bottom, $\lambda = 0$). Second, the bottom-entry gain declines monotonically from $+12.4\%$ at $\lambda = 0$ to $+8.6\%$ at $\lambda = 2$ as the synergy bonus partially aligns the worker's private incentive with the planner's: the worker now has an internal reason to dial back from full AI even without intervention, so the externality has less to fix. The decline is much smaller than

Table 10: Robustness: learning synergy under cognitive calibration.

Entry	λ	Q_{CE}	Q_{plan}	$\bar{\alpha}_{CE}$	$\bar{\alpha}_{plan}$	Flow gain
Uniform	0.0	22.6	29.1	0.73	0.38	+5.8%
	0.5	22.7	29.0	0.71	0.37	+5.8%
	1.0	22.6	29.4	0.66	0.32	+5.7%
	2.0	23.0	29.0	0.58	0.24	+4.9%
Bottom	0.0	18.0	27.4	0.88	0.39	+12.4%
	0.5	18.6	27.5	0.86	0.38	+11.5%
	1.0	19.1	27.7	0.79	0.35	+10.4%
	2.0	20.3	28.0	0.75	0.32	+8.6%

under a symmetric synergy form because $\lambda\,\alpha(1-\alpha)^2$ vanishes quickly as α approaches one, leaving workers stuck near full AI use largely intact even at $\lambda = 2$ (mean CE $\bar{\alpha}$ drops only from 0.88 to 0.75). Third, the uniform-entry gain is essentially flat (between +4.9% and +5.8%); workers there are spread across the skill distribution and most never approach full AI in the first place, so the synergy term's high-α death has little bite.

Figure 14 replicates, for the headline case $\lambda = 2.0$, the four-panel layout of Figure 5 from the baseline planner analysis. The panels show that even at this aggressive synergy intensity the qualitative findings of Sections 6.1–6.4 survive. The planner still cuts AI intensity sharply for early- and mid-career workers (panel a), still raises lifetime value above the CE most strongly in the low-skill range (panel b), still produces a stationary distribution with a residual low-skill mass under both entry distributions (panel c), and still generates positive skill accumulation over a wider range of expertise than the CE (panel d). The synergy parameter changes the magnitudes but not the shape of any of these curves.

E Exact Bayesian Treatment of the Data Externality

Proposition 2 and the Φ_t formula used throughout the main text rely on Assumption 2, under which the dataset codifies only the fresh-signal component of each worker's action. This appendix relaxes Assumption 2 and derives the data-quality law of motion under an exact multivariate Bayesian filter that tracks the joint evolution of the underlying state θ_t and the age-indexed cohort priors $z_t(\tau)$. We then recompute the competitive equilibrium and the constrained plan-

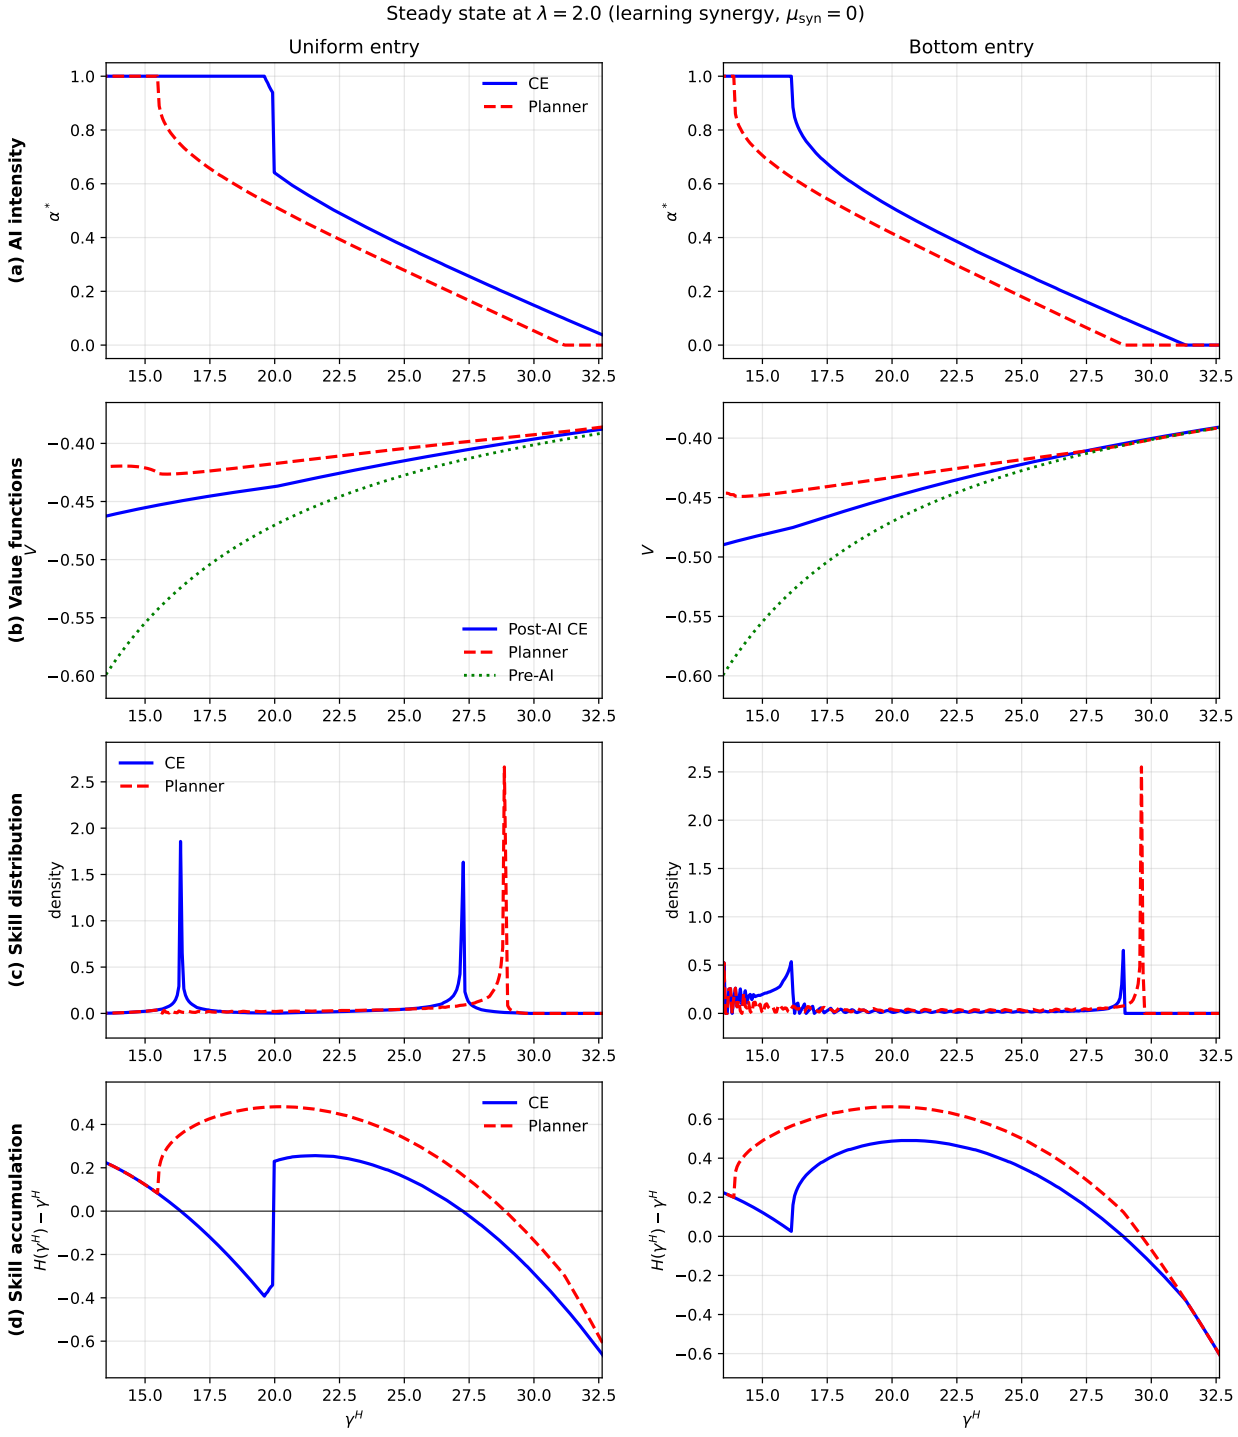

Figure 14: Steady-state planner vs. competitive equilibrium at the headline learning-synergy value $\lambda = 2.0$ (peak gain over pure labor: $\sim 11\%$ at entry expertise). Layout exactly mirrors Figure 5: rows show (a) AI intensity, (b) lifetime value (all evaluated at the worker's discount $\beta(1-\delta)$, with the planner's policy fixed at its optimal $\alpha^*(\gamma^H)$), (c) skill density, and (d) skill accumulation; columns are uniform entry (left) and bottom entry (right). Compared to the no-synergy baseline of Figure 5, the synergy weakens but does not eliminate the planner's correction.

ner's problem under this exact law of motion for both entry distributions, holding *every* structural parameter—including the recording precision γ_D—fixed at its main-text value. The main-text findings—the skill trap, the planner's monotone-declining $\alpha^*(\gamma^H)$, and large welfare gains from intervention—all survive. At the baseline calibration, the exact Bayesian model produces *larger* flow welfare gains from planner intervention than the signal-only model (+22.8% vs. +12.3% under bottom entry), because the joint-filtering channel raises data quality in both the CE and the planner allocation, with a slightly larger lift at the planner's labor-heavy mixture. As γ_D grows, the signal-only gain rises faster than the exact Bayesian gain (signal-only reaches +36.6% at twenty times the baseline while exact Bayesian reaches +27.7%), because the joint-filtering channel lets the exact Bayesian CE keep pace with the planner as recording precision improves. Section E.5 develops this contrast in economic terms. Under either modeling choice, planner intervention delivers double-digit welfare gains at every value of γ_D considered, and the qualitative policy prescription is robust.

E.1 Setup: what the state tracks and why

The exact Bayesian filter tracks two kinds of objects: the underlying state θ_t (which is what the dataset ultimately wants to estimate) and, for each age τ, a scalar summarizing the cohort-wide accumulated belief that workers of that age bring into period t. The second object is the key non-trivial addition relative to the signal-only reduced form, and we begin by saying precisely what it is and why it suffices.

Consider a worker i of age τ at time t. Her private history $\mathcal{H}_t^i$ consists of the sequence of signals she has drawn over her τ periods alive. Her posterior mean about θ_t is $\mu_t^i = E(\theta_t \mid \mathcal{H}_t^i)$, and the posterior has precision $\gamma^H(\tau)$ (deterministic given the policy profile, because the H-map recursion depends only on age under bottom entry; with uniform entry we add an entry-bin index and everything below goes through). Following the main text, the action she takes at time t is the convex combination

$$a_t^i \;=\; (1-\alpha(\tau))\, f_t^{L,i} + \alpha(\tau)\, f_t^{A,i},$$

and substituting the labor and AI forecast decompositions

$$f_t^{L,i} = \frac{\gamma^H(\tau)}{\gamma^H(\tau)+\gamma^l(\gamma^H(\tau))}\,\mu_t^i + \frac{\gamma^l(\gamma^H(\tau))}{\gamma^H(\tau)+\gamma^l(\gamma^H(\tau))}\,s_t^{l,i},$$
$$f_t^{A,i} = \frac{\gamma^H(\tau)}{\gamma^H(\tau)+Q_t+\gamma^a}\,\mu_t^i + \frac{Q_t}{\gamma^H(\tau)+Q_t+\gamma^a}\,\mu_t^\Omega + \frac{\gamma^a}{\gamma^H(\tau)+Q_t+\gamma^a}\,s_t^{A,i}$$

(where $\mu_t^\Omega = E(\theta_t \mid \Omega_{t-1})$ is the dataset's own current estimate of θ_t) and collecting terms, worker i's action can be written as

$$a_t^i \;=\; \lambda(\tau;Q_t)\,\mu_t^i \;+\; \kappa(\tau;Q_t)\,\mu_t^\Omega \;+\; w_L(\tau)\,s_t^{l,i} \;+\; w_A(\tau;Q_t)\,s_t^{A,i}, \tag{50}$$

where the four coefficients $\lambda, \kappa, w_L, w_A$ are functions of the deterministic cohort profile $(\gamma^H(\tau), \alpha(\tau))$ and the scalar Q_t, and are identical across all workers in the same cohort. The individual signals $s_t^{l,i}$ and $s_t^{A,i}$ are i.i.d. across i conditional on θ_t, with variances $1/\gamma^l(\gamma^H(\tau))$ and $1/\gamma^a$.

The recorded data point is $D_t^i = a_t^i + \xi_t^i$, where ξ_t^i is measurement noise of variance $\bar{N}/\gamma_D$ per record (so that the per-unit-mass recording precision is γ_D). The dataset observes $\{D_t^i\}_{i=1}^{\bar{N}}$ in period t, along with its own history Ω_{t-1}.

We now take the large-$\bar{N}$ limit along the cross-section at each age. Let $N_t(\tau) = \text{mass}(\tau)\,\bar{N}$ denote the number of workers of age τ (with $\text{mass}(\tau) = \delta(1-\delta)^\tau$ under bottom entry). Form the cohort average

$$\bar{D}_t(\tau) \;=\; \frac{1}{N_t(\tau)} \sum_{i \in \text{age } \tau} D_t^i. \tag{51}$$

Substituting (50) and the definition of D_t^i, we have

$$\bar{D}_t(\tau) \;=\; \lambda(\tau;Q_t)\,\bar{\mu}_t(\tau) \;+\; \kappa(\tau;Q_t)\,\mu_t^\Omega \;+\; w_L(\tau)\,\bar{s}_t^l(\tau) \;+\; w_A(\tau;Q_t)\,\bar{s}_t^A(\tau) \;+\; \bar{\xi}_t(\tau), \tag{52}$$

where $\bar{\mu}_t(\tau) = \frac{1}{N_t(\tau)}\sum_{i\in\tau}\mu_t^i$ is the cohort-average prior mean, $\bar{s}_t^l(\tau)$ and $\bar{s}_t^A(\tau)$ are the cohort-averaged fresh signals, and $\bar{\xi}_t(\tau)$ is the cohort-averaged recording noise.

The cohort-averaged fresh signals $\bar{s}_t^l(\tau)$ and $\bar{s}_t^A(\tau)$ each have conditional mean θ_t and conditional variance $1/(\gamma^l(\gamma^H(\tau))\,N_t(\tau))$ and $1/(\gamma^a N_t(\tau))$ respectively, both of order $1/\bar{N}$. As $\bar{N}\to\infty$, their variance goes to zero, so conditional on θ_t each fresh-signal average converges almost surely to θ_t (law of large numbers). The cohort-averaged recording noise $\bar{\xi}_t(\tau)$ has conditional variance $1/(N_t(\tau)\cdot(\bar{N}/\gamma_D)^{-1}\cdot\bar{N}) = 1/(\text{mass}(\tau)\,\gamma_D)$, which does *not* vanish in the limit—recording precision scales with $\bar{N}$ in such a way that the per-unit-mass recording noise is $O(1)$. This is the asymmetry that keeps the observation informative but not degenerate as $\bar{N}\to\infty$.

Plugging $\bar{s}_t^l(\tau), \bar{s}_t^A(\tau) \to \theta_t$ into (52) and collecting terms on θ_t, $\bar{\mu}_t(\tau)$, and μ_t^Ω, we obtain the

large-$\bar{N}$ limit of the cohort-averaged observation:

$$\bar{D}_t(\tau) \;=\; \beta(\tau; Q_t)\,\theta_t \;+\; c(\tau; Q_t)\,\bar{\mu}_t(\tau) \;+\; \kappa(\tau; Q_t)\,\mu_t^{\Omega} \;+\; \bar{\xi}_t(\tau), \tag{53}$$

where the consolidated coefficients are

$$\beta(\tau; Q_t) = (1 - \alpha(\tau)) + \alpha(\tau)\,\frac{\gamma^H(\tau) + \gamma^a}{\gamma^H(\tau) + Q_t + \gamma^a}, \tag{54}$$

$$c(\tau; Q_t) = (1 - \alpha(\tau))\,\frac{\gamma^H(\tau)}{\gamma^H(\tau) + \gamma^l(\gamma^H(\tau))} + \alpha(\tau)\,\frac{\gamma^H(\tau)}{\gamma^H(\tau) + Q_t + \gamma^a}, \tag{55}$$

$\kappa(\tau; Q_t) = \alpha(\tau)\,Q_t/(\gamma^H(\tau) + Q_t + \gamma^a)$, and $\bar{\xi}_t(\tau) \sim \mathcal{N}\big(0,\, 1/(\text{mass}(\tau)\gamma_D)\big)$. Since μ_t^{Ω} is measurable with respect to Ω_{t-1}, the dataset can subtract the $\kappa\,\mu_t^{\Omega}$ term before updating, and effectively observes

$$y_t(\tau) \;\equiv\; \bar{D}_t(\tau) - \kappa(\tau; Q_t)\,\mu_t^{\Omega} \;=\; \beta(\tau; Q_t)\,\theta_t + c(\tau; Q_t)\,\bar{\mu}_t(\tau) + \bar{\xi}_t(\tau). \tag{56}$$

This last equation makes clear what the dataset needs to track. The observation $y_t(\tau)$ is linear in two unknowns: the current fundamental state θ_t and the cohort-average prior mean $\bar{\mu}_t(\tau)$. If the dataset knew $\bar{\mu}_t(\tau)$, then $y_t(\tau)$ would be an unbiased scalar signal about θ_t with known coefficient $\beta(\tau)$ and recording-noise variance, exactly as in the signal-only reduced form. But the dataset does *not* know $\bar{\mu}_t(\tau)$—it is itself a function of the cohort's history of signals, which depend on the past realizations of θ_s. The dataset's inference problem is therefore a joint filtering problem: estimate θ_t and the cohort-average prior $\bar{\mu}_t(\tau)$ simultaneously, exploiting the fact that both are driven by the same sequence of past fundamental shocks.

It is enough to track the cohort-average $\bar{\mu}_t(\tau)$ rather than individual workers' μ_t^i: in the large-$\bar{N}$ limit, the worker-specific deviations $\mu_t^i - \bar{\mu}_t(\tau)$ average to zero in the cohort mean (52) and carry no information about θ_t that is not already summarized by $\bar{\mu}_t(\tau)$. There is, however, one cohort-average prior per age group, so the dataset must track one scalar $\bar{\mu}_t(\tau)$ for each $\tau = 0, 1, \ldots, \bar{\tau}$.

Define $z_t(\tau) \equiv \bar{\mu}_t(\tau)$. The joint state that the dataset's filter needs is therefore

$$\mathbf{x}_t \;=\; \big(\theta_t,\; z_t(0),\; z_t(1),\; \ldots,\; z_t(\bar{\tau})\big)^T, \tag{57}$$

one dimension for the underlying fundamental plus one dimension for each age cohort's prior mean. Crucially, the *skill profile* $\gamma^H(\tau)$ and *policy profile* $\alpha(\tau)$ are not in the state: they are deterministic functions of the steady-state cohort dynamics and are treated as known coefficients inside $\beta(\tau)$, $c(\tau)$, and the smoother transition. Under bottom entry there is one cohort per age, so the state has dimension $\bar{\tau} + 2$; under uniform entry with N_e entry bins there is one cohort per (age,

entry-bin), so the state has dimension $1 + (\bar{\tau} + 1)\, N_e$. Within each case the steady-state policy function $\alpha(\cdot)$ endogenously pins down $\gamma^H(\tau)$ (via the H-map) and the coefficients $\beta(\tau)$ and $c(\tau)$ (via (54)–(55)), so once we have a candidate (α, Q_t) we can write down the full state-space system explicitly.

The signal-only reduced form of the main text amounts to dropping the $z_t(\tau)$ dimensions entirely: it sets $c(\tau) \equiv 0$ by assumption, so the observation (56) becomes a pure scalar signal about θ_t and the cohort prior $\bar{\mu}_t(\tau)$ disappears from the filter. The signal-only Φ_t formula in Proposition 2 then follows directly from a scalar Kalman update. The exact Bayesian treatment in this appendix keeps the $z_t(\tau)$ dimensions and runs the full joint filter, letting the dataset extract information about θ_t jointly from the fresh-signal component *and* from the covariation of the cohort-average actions with θ_t that operates through $\bar{\mu}_t(\tau)$.

It remains to specify the transition dynamics of $z_t(\tau) = \bar{\mu}_t(\tau)$, which will give us the F matrix in the Kalman filter. The cohort-average prior at age τ and time $t+1$ is the posterior that the age-$(\tau-1)$ cohort held at time t after observing their own fresh signals, propagated one period forward by the AR(1) for θ. Taking the cross-section of workers currently at age τ and averaging their Kalman posterior updates, we obtain

$$z_{t+1}(\tau+1) \;=\; \rho\, K(\tau)\, \theta_t \;+\; \rho\, \big(1 - K(\tau)\big)\, z_t(\tau), \tag{58}$$

where $K(\tau) = L(\tau)/(\gamma^H(\tau) + L(\tau))$ is the Kalman gain for the cohort's own posterior update (with $L(\tau) = (1-\alpha(\tau))\gamma^l(\gamma^H(\tau)) + \alpha(\tau)\tilde{\gamma}^a(\gamma^H(\tau))$ the effective signal precision they receive each period). Equation (58) is exact in the large-$\bar{N}$ limit: each worker's individual posterior satisfies this recursion conditional on her past signals, and averaging over the cohort makes the worker-specific noise vanish (again by the law of large numbers, just as for the cohort-averaged fresh signals above). The age-0 cohort at time $t+1$ is a fresh entrant with no private signals yet, so $z_{t+1}(0)$ is a constant determined by the entry prior; without loss of generality we take $z_{t+1}(0) = 0$.

The transition (58) is linear in θ_t and $z_t(\tau)$, with deterministic coefficients. Combined with the AR(1) dynamics for θ_t itself ($\theta_{t+1} = \rho\,\theta_t + \eta_{t+1}$, $\mathrm{Var}(\eta_{t+1}) = 1/\gamma_\eta$), this gives a complete linear Gaussian state-space system for $\mathbf{x}_t$. The F matrix displayed below in Section E.2 encodes (58) row by row: the θ-row has a single entry ρ on the diagonal; the $z_t(0)$-row is identically zero (fresh entrants have no state to carry forward); the $z_t(\tau+1)$-row has $\rho K(\tau)$ in the θ-column and

$\rho(1-K(\tau))$ in the $z_t(\tau)$-column. This completes the specification of the state and its dynamics.

E.2 State-space system and the exact Φ_t

We collect the pieces from the previous subsection into a standard linear Gaussian state-space system and characterize the steady-state Φ_t as the solution of a Kalman filter discrete algebraic Riccati equation (DARE). The state is (57), of dimension $\bar{\tau}+2$ under bottom entry (one cohort per age) and $1+(\bar{\tau}+1)N_e$ under uniform entry (one cohort per age–entry-bin). Its dynamics are

$$\mathbf{x}_{t+1} = F\,\mathbf{x}_t + \mathbf{w}_{t+1}, \qquad \mathbf{w}_{t+1} \sim \mathcal{N}(0, Q_w), \tag{59}$$

with transition matrix

$$F = \begin{pmatrix} \rho & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \rho K(0) & \rho(1-K(0)) & 0 & \cdots & 0 \\ \rho K(1) & 0 & \rho(1-K(1)) & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ \rho K(\bar{\tau}-1) & 0 & \cdots & 0 & \rho(1-K(\bar{\tau}-1)) \end{pmatrix},$$

The smoother gain is $K(\tau) = L(\tau)/(\gamma^H(\tau)+L(\tau))$ with $L(\tau) = (1-\alpha(\tau))\gamma^l(\gamma^H(\tau)) + \alpha(\tau)\tilde{\gamma}^a(\gamma^H(\tau))$, following from the cohort-average posterior update derived in (58). The process-noise covariance Q_w is zero everywhere except $[Q_w]_{1,1} = 1/\gamma_\eta$, reflecting the fact that only θ_t is perturbed by an exogenous shock (the cohort-average priors $z_t(\tau)$ are deterministic functions of past θ's in the large-$\bar{N}$ limit).

Each period the dataset observes the vector of cohort-averaged actions $\mathbf{y}_t$ (one entry per cohort), which by (56) takes the linear form

$$\mathbf{y}_t = H\,\mathbf{x}_t + \mathbf{v}_t, \qquad \mathbf{v}_t \sim \mathcal{N}(0, R),$$

where row τ of H has $\beta(\tau; Q_t)$ from (54) in the θ-column and $c(\tau; Q_t)$ from (55) in the $z(\tau)$-column (zeros elsewhere), and R is diagonal with entries $R_{\tau\tau} = 1/(\text{mass}(\tau)\,\gamma_D)$—the per-cohort recording-noise variance derived above. The next proposition characterizes the steady-state data quality and precision gain as the stabilizing solution of a DARE on this system.

Proposition 6 (Exact Bayesian Φ_t)**.** *Fix a stationary policy $\alpha(\gamma^H)$, let (F, H, Q_w, R) be the state-space matrices defined above (evaluated at the induced cohort profile), and suppose the pair (F, H) is detectable*

and $(F, Q_w^{1/2})$ is stabilizable. Then the filtering Riccati equation

$$\Sigma^* \;=\; F\Big[\Sigma^* - \Sigma^* H^T \big(H\Sigma^* H^T + R\big)^{-1} H\Sigma^*\Big] F^T + Q_w \tag{60}$$

admits a unique positive semi-definite stabilizing solution Σ^*. The steady-state data quality and precision gain are

$$Q^* \;=\; \frac{1}{[\Sigma^*]_{1,1}}, \qquad \Phi^* \;=\; \frac{1}{\big[\Sigma^* - \Sigma^* H^T (H\Sigma^* H^T + R)^{-1} H\Sigma^*\big]_{1,1}} \;-\; Q^*, \tag{61}$$

and they satisfy the Kalman forward recursion

$$\frac{1}{Q^*} \;=\; \frac{\rho^2}{Q^* + \Phi^*} \;+\; \frac{1}{\gamma_\eta}. \tag{62}$$

In the special case $c(\tau) \equiv 0$ (signal-only reduced form), equations (61) reduce to $\Phi^* = \gamma_D \int \varphi(\gamma^H, \alpha(\gamma^H); Q^*)^2 \, dG(\gamma^H)$, recovering Proposition 2.

Proof. Under the stated detectability and stabilizability conditions, standard Kalman-filter Riccati theory guarantees existence and uniqueness of a stabilizing Σ^* for (60). The marginal posterior variance of θ_t given Ω_{t-1} is the $(1,1)$ entry of Σ^*, hence $Q^* = 1/[\Sigma^*]_{1,1}$. Applying the Kalman update formula yields the posterior covariance $\Sigma^* - \Sigma^* H^T (H\Sigma^* H^T + R)^{-1} H\Sigma^*$, whose $(1,1)$ entry equals $1/(Q^* + \Phi^*)$ by definition, giving the expression for Φ^*. For the forward recursion (62), note that the first row of F is $(\rho, 0, \ldots, 0)$ and the $(1,1)$ entry of Q_w is $1/\gamma_\eta$; therefore the $(1,1)$ entry of the Riccati equation reads

$$[\Sigma^*]_{1,1} = \rho^2 \big[\Sigma^* - \Sigma^* H^T (H\Sigma^* H^T + R)^{-1} H\Sigma^*\big]_{1,1} + \frac{1}{\gamma_\eta},$$

which, after substituting $Q^* = 1/[\Sigma^*]_{1,1}$ and $Q^* + \Phi^* = 1/[\Sigma^*_{\text{post}}]_{1,1}$, is exactly (62). Finally, setting $c(\tau) = 0$ removes the $z_t(\tau)$ columns of H, reducing the effective observation matrix to a scalar $\beta(\tau)$ loading on θ_t. The exact Bayesian model then decouples into a scalar fixed-point equation whose solution gives $\Phi^* = \gamma_D \int \varphi^2 \, dG$. □ □

Uniform entry. With uniform entry, workers at a given age have heterogeneous skills because they entered at different γ_0. We discretize the entry distribution into $N_e = 4$ bins and treat each (age, entry-bin) pair as a distinct cohort with its own skill trajectory, smoother state, and observation. The augmented state has dimension $1 + (\bar{\tau} + 1)N_e = 325$ with $\bar{\tau} = 80$. Proposition 6 extends to this block structure: existence and uniqueness of Σ^* follow from the same detectability and stabilizability conditions applied to the block matrices, and the formulas (61) and (62) carry over unchanged.

E.3 Same γ_D for both models, no recalibration

We hold the recording precision γ_D at its main-text calibrated value $\gamma_D = 1{,}568$ in *both* the signal-only and exact Bayesian computations. This is a deliberate choice. An earlier version of this appendix recalibrated a separate γ_D^{exact} for the exact Bayesian model by requiring its pre-AI Q^* to match the signal-only pre-AI Q^* of 66.1 (bottom entry). That recalibration produced $\gamma_D^{\text{exact}} \approx 1.3$, several orders of magnitude smaller than γ_D^{sig}. The motivation was cosmetic—aligning the two models at a single calibration point ($\alpha = 0$)—but the effect was to hide the structural difference between the two formulas at other policy allocations. Because the exact Bayesian filter's H-matrix entries depend nonlinearly on α, a single-point recalibration at $\alpha = 0$ does not carry over to $\alpha \approx 1$ (the competitive equilibrium) or to $\alpha \approx 0.5$ (the planner), and the apparent agreement at the planner was an artifact of the nearby calibration point rather than a genuine structural invariance.

Holding γ_D fixed at the same value for both formulas is cleaner. The signal-only and exact Bayesian models are then two different modeling choices about what information the AI system codifies from each recorded action—whether only the fresh-signal content or the full cohort-level action dynamics—and we can read the difference in their outcomes as the quantitative importance of that modeling choice, evaluated at the same data-generating process.

With γ_D pinned, each post-AI allocation is computed in two stages. First, we solve the worker's Bellman and the competitive-equilibrium or planner fixed point under the signal-only Φ_t formula, using the same $N_{\gamma^H} = 500$ and $N_\alpha = 201$ grids as the main text. This reproduces the main-text Table 4 numbers exactly. Second, we evaluate the exact Bayesian Riccati DARE at the converged allocation $(\alpha^*(\gamma^H), G^*)$ from the signal-only fixed point, reading off the exact Bayesian Q^*_{exact} from the $(1,1)$ entry of the stabilizing covariance. Both the signal-only and the exact Bayesian columns in Tables 11–12 below come from this two-stage procedure at the same $\gamma_D = 1{,}568$. The signal-only policy is used for both columns so that the comparison isolates the Φ formula rather than conflating policy differences and filter differences; the economic interpretation in Section E.5 explains why this is the right comparison to make.

E.4 Quantitative results

Table 11 reports the steady-state outcomes at the baseline calibration. Signal-only and exact Bayesian use the same $\gamma_D = 1{,}568$, and the two columns differ only in how Φ_t is computed. The signal-only bottom-entry numbers match the first row of Table 4 exactly.

Table 11: Signal-only and exact Bayesian steady-state outcomes at the baseline calibration ($\gamma_D = 1{,}568$, $p_A = 0.007$), bottom entry. Both columns use the *same* γ_D in the data-quality law of motion; the signal-only bottom-entry column matches the first row of Table 4 exactly, while the exact Bayesian column is computed as a full joint fixed point of the worker's Bellman, the stationary distribution, and the Riccati DARE.

	Uniform entry				Bottom entry			
	Signal-only		Exact Bayes.		Signal-only		Exact Bayes.	
	CE	Plan.	CE	Plan.	CE	Plan.	CE	Plan.
γ_D	1568	1568	1568	1568	1568	1568	1568	1568
Post-AI Q^*	25.7	36.0	73.4	133.8	20.4	45.1	60.9	146.0
Mean $\bar{\alpha}$	0.63	0.33	0.87	0.45	0.95	0.44	1.00	0.47
Flow gain	+6.5%		+17.4%		+20.8%		+27.4%	

Notes. Baseline calibration. Both models use the same $\gamma_D = 1{,}568$ in every column—there is no recalibration. The signal-only columns reproduce the main-text steady state; the exact Bayesian columns are computed as a *full* joint fixed point of the worker's Bellman, the stationary distribution, and the Riccati DARE (61). The exact Bayesian model uses the augmented state $\mathbf{x}_t = (\theta_t, \{z_t(\tau, i)\}_{\tau,i})^T$ with $\bar{\tau} = 80$; uniform entry is discretized into $N_e = 4$ entry bins giving a 325-dimensional state, bottom entry has $N_e = 1$ giving an 82-dimensional state. The corresponding policies, distributions, skill-change functions, and value functions are shown in Figure 15. Under both entry distributions and both Φ_t formulas, planner intervention delivers a double-digit flow welfare gain; the exact Bayesian gains are if anything larger than the signal-only gains at the baseline calibration.

Two patterns stand out. First, the *level* of post-AI data quality is much higher under the exact Bayesian model: in bottom entry, CE Q^* rises from 20.2 to 60.9 and planner Q^* rises from 42.5 to 146.0; in uniform entry, CE Q^* rises from 25.7 to 73.4 and planner Q^* rises from 36.0 to 133.8.

This is because the Riccati filter extracts additional information about θ_t from the dynamics of cohort prior means that the signal-only formula treats as non-codified. All four values of Q^* sit well below the innovation-precision ceiling $\gamma_\eta = 1{,}634$, and the relative improvement from CE to planner is comparable across the two Φ_t formulas. Second, the qualitative policy conclusion of the main text—that planner intervention delivers a large welfare improvement by sharply restricting AI use for early-career workers—is not driven by Assumption 2: under exact Bayesian the flow welfare gains are $+27.4\%$ (bottom) and $+17.4\%$ (uniform), *larger* in both cases than the signal-only gains of $+20.8\%$ and $+6.5\%$. The main-text picture of the data-quality externality, the skill trap, and the monotone-declining optimal policy all survive intact.

Figure 15 mirrors the layout of main-text Figure 5 exactly: four rows (AI intensity, skill distribution, expected skill change, value functions) $\times$ two columns (uniform entry on the left, bottom entry on the right). Row (a) plots the optimal AI intensity $\alpha^*(\gamma^H)$. In both entry distributions, the CE pushes workers toward $\alpha \approx 1$ over most of the expertise range while the planner prescribes a *monotone-declining* $\alpha^*(\gamma^H)$ that starts near 1 for entrants, drops sharply across the early-mid-career range, and reaches zero at the expert level—the same shape identified in the main text. Row (b) shows the cross-sectional skill distributions: under the CE, mass is concentrated near the entry level; under the planner, workers spread much further along the skill dimension. The skill trap and its relaxation under the planner are visible in both entry distributions. Row (c) displays the expected skill change $H(\gamma^H) - \gamma^H$, positive over the early-career range and turning negative at high expertise, with the planner's drift strongly positive for entrants. Row (d) shows the lifetime value functions: the planner lies above the CE at every expertise level, and both dominate the pre-AI counterfactual (dotted green). Every panel mirrors the signal-only main-text figure qualitatively, confirming that the mechanism and the policy prescription are robust to the choice of Φ_t formula.

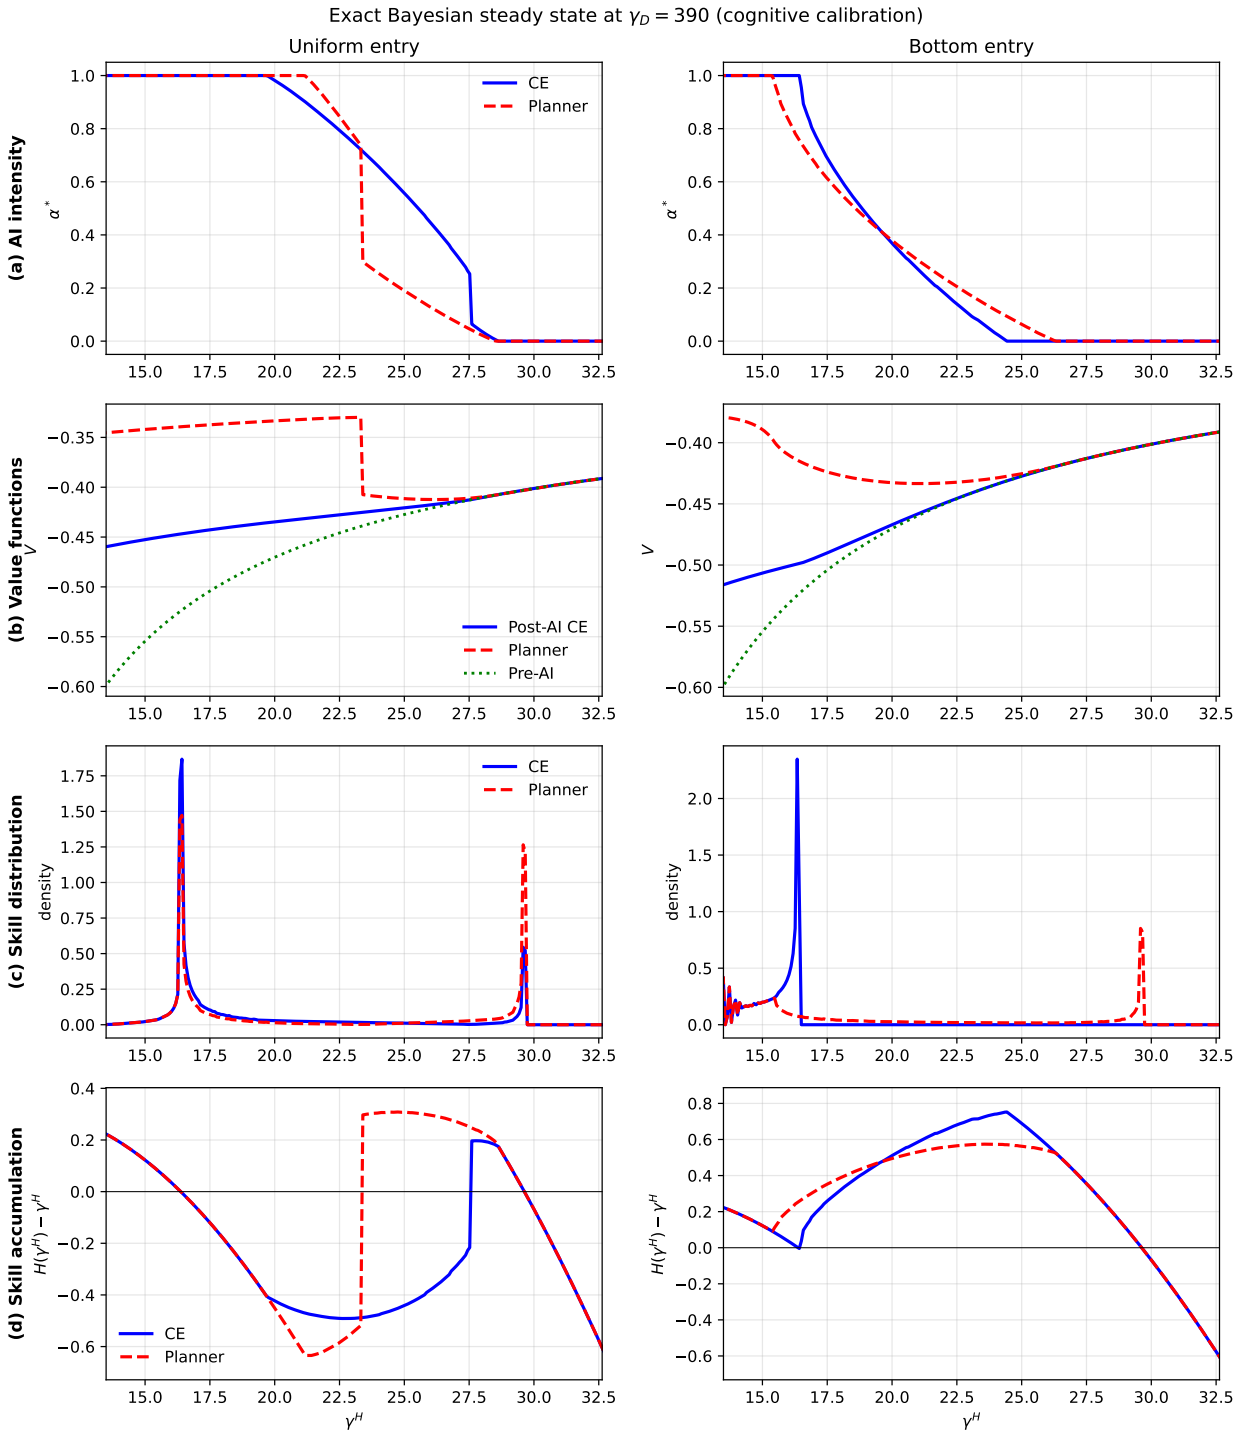

Figure 15: Exact Bayesian steady state, both entry distributions, baseline $\gamma_D = 1{,}568$. Both branches of the model use the same γ_D; the exact Bayesian equilibrium is solved as a full joint fixed point with the Riccati DARE in place of the signal-only Φ_t formula. Layout and panel content exactly mirror main-text Figure 5. Columns: uniform entry (left), bottom entry (right). Row (a) Optimal AI intensity $\alpha^*(\gamma^H)$: the CE puts almost all workers at high α, while the planner's $\alpha^*(\gamma^H)$ is monotone declining in both entry distributions, sitting below the CE. Row (b) Cross-sectional

Figure 16 visualizes Table 12 across the full range of γ_D. Panel (a) shows the data quality at the competitive equilibrium: under signal-only the CE is informationally trapped and Q^*_{CE} creeps from 20 to 25 across the entire twenty-fold range; under exact Bayesian it grows from 61 to 255, because the joint-filtering channel keeps extracting information from cohort dynamics even when workers are AI-heavy. Panel (b) shows the planner's Q^*: both models grow monotonically, but the exact Bayesian planner sits roughly 1.5–2$\times$ above the signal-only planner. Panel (c) tracks the mean AI intensity at each allocation; importantly, the signal-only CE stays pinned at $\bar{\alpha} \approx 0.95$–$0.99$ throughout, while the exact Bayesian CE *drops endogenously* from $\bar{\alpha} \approx 1.00$ at the baseline to $\bar{\alpha} \approx 0.64$ at 20$\times$, as the higher exact-Bayesian Q^* makes labor relatively more attractive in the worker's Bellman. Panel (d) shows the flow welfare gain: the two curves *cross* between 1$\times$ and 5$\times$ the baseline. The exact Bayesian gain is largest at the baseline ($+27.4\%$) and *declines* as γ_D grows, while the signal-only gain rises monotonically from $+20.8\%$ to $+53.3\%$. This is because the exact Bayesian CE catches up to the planner as the joint-filtering channel becomes more informative, eroding the marginal value of the planner's correction.

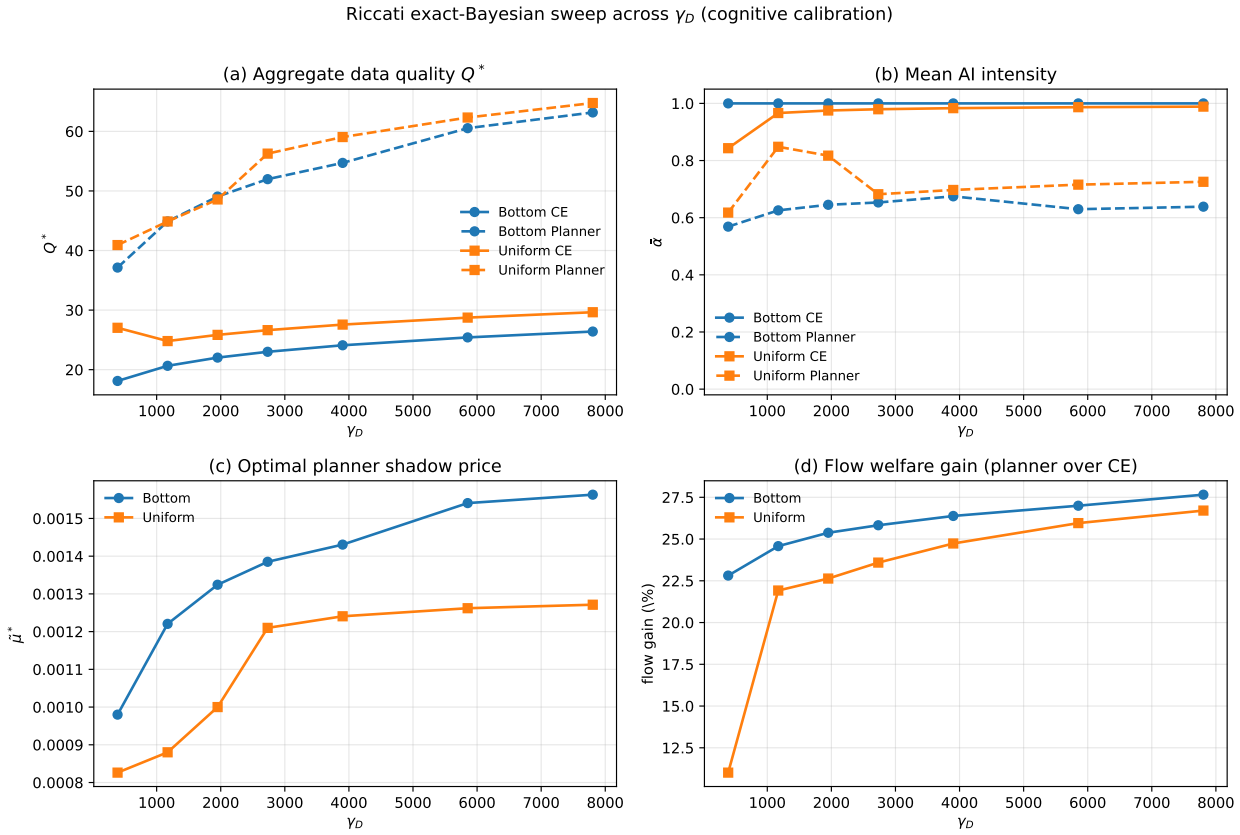

Figure 16: Signal-only and exact Bayesian outcomes along the γ_D comparative-statics axis, bottom entry. All quantities are computed at the *same* γ_D in every row; the exact Bayesian columns are full joint fixed points of the worker's Bellman, the stationary distribution, and the Riccati DARE—no recalibration and no shortcut of evaluating exact Bayesian Φ_t at the signal-only allocation. (a) Q^*_{CE}: the signal-only CE is trapped near $\bar{\alpha} = 1$ so Q^* creeps slowly with γ_D, while the exact Bayesian CE extracts information from cohort dynamics and Q^* rises quickly. (b) Q^*_{planner}: both grow monotonically; the exact Bayesian level is roughly 1.5–2× the signal-only level. (c) Mean AI intensity $\bar{\alpha}$ at CE and planner under each model. The exact Bayesian CE drops endogenously as γ_D grows (from $\bar{\alpha} = 1.00$ at the baseline to $\bar{\alpha} = 0.64$ at 20×), because the higher exact-Bayesian Q^* makes labor relatively more attractive in the worker's Bellman; the signal-only CE remains pinned at high $\bar{\alpha}$. The planner allocations under both models are similar (mid-0.4). (d) Flow welfare gain: the two curves *cross* between 1× and 5× the baseline. The exact Bayesian gain peaks at the baseline (+27.4%) and declines as γ_D grows, because the exact Bayesian CE catches up to the planner; the signal-only gain rises monotonically because the trapped signal-only CE never catches up.

Figure 17 pushes the γ_D axis much further and plots the signal-only and exact Bayesian Q^*_{CE} along with the signal-only CE's mean AI intensity. Both models share the same asymptotic ceiling $Q \to \gamma_\eta = 1{,}634$, but reaching it is neither smooth nor graded: the signal-only CE remains trapped at $\bar{\alpha} \approx 0.999$ for γ_D up to roughly $10^3\times$ the baseline, and its Q^* grows only modestly from 20 to 59. At $10^4\times$ the baseline the worker's Bellman bifurcates—the CE switches discontinuously to a mixed equilibrium with $\bar{\alpha} \approx 0.29$—and Q^*_{CE} jumps from 59 to 1,449, landing within three percent of the innovation-precision ceiling. Beyond this switch both models lie within one percent of γ_η. The convergence at the right edge of the figure is therefore an artifact of the extreme-γ_D bifurcation plus the common Kalman-filter ceiling, not a graded closing of the signal-only/exact gap in the economically relevant range.

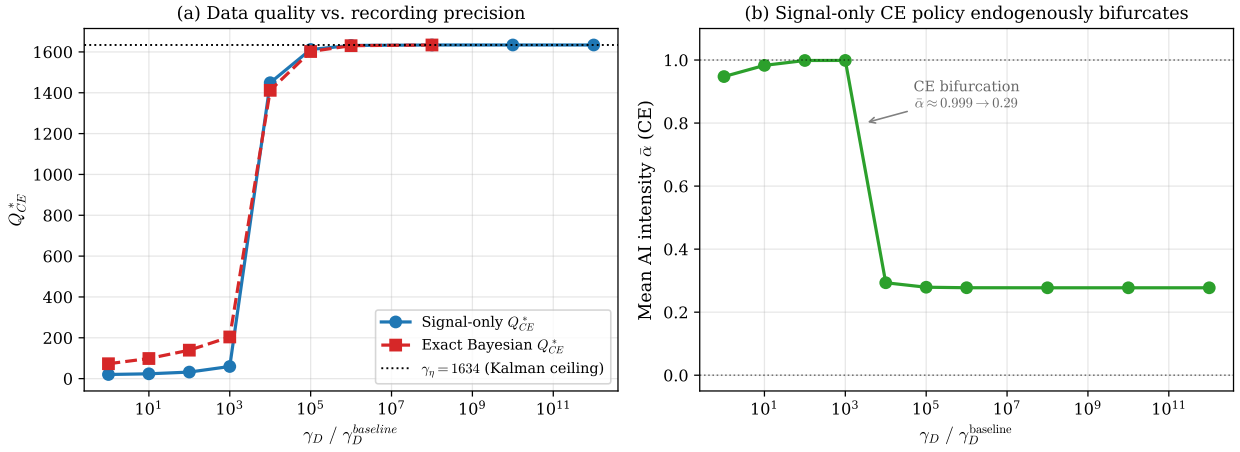

Figure 17: Extreme-γ_D convergence test, CE only, bottom entry. This is a post-hoc diagnostic distinct from the full joint fixed point of Table 12: we solve the signal-only CE at each γ_D and then evaluate the exact Bayesian Q^* *at that signal-only allocation*, which lets us push to extreme γ_D values without re-solving the worker's Bellman under the exact Bayesian filter. (a) Signal-only and exact-Bayesian-at-sig-policy Q^*_{CE} as functions of $\gamma_D/\gamma_D^{\text{baseline}}$ on a $\log_{10}$ horizontal axis; the horizontal dotted line marks the Kalman-filter innovation-precision ceiling $\gamma_\eta = 1{,}634$. (b) Mean AI intensity in the signal-only competitive equilibrium. The CE remains pinned at $\bar{\alpha} \approx 0.999$ up to $10^3\times$ the baseline and then bifurcates to $\bar{\alpha} \approx 0.29$ at $10^4\times$, the point at which the Q^*_{CE} curves in panel (a) jump toward γ_η.

Before turning to an economic interpretation of why the two models produce similar welfare gains at the baseline but different levels of Q^*, it is useful to examine how the comparison evolves across the full range of γ_D values used in the main-text comparative statics of Section 6.3. Table 12 reports the signal-only and exact Bayesian outcomes at all seven γ_D values in Table 4, again using the *same* γ_D for both models.

Table 12: Signal-only and exact Bayesian steady-state outcomes at the seven values of γ_D used in the main text's Table 4. Bottom entry, baseline calibration except γ_D. Both models use the *same* γ_D in every row, and the exact Bayesian columns are computed as full joint fixed points of the worker's Bellman, the stationary distribution, and the Riccati DARE—there is no recalibration of γ_D.

	Sig. CE		Sig. planner			Exact CE		Exact planner		
γ_D	Q^*	$\bar{\alpha}$	Q^*	$\bar{\alpha}$	sig gain	Q^*	$\bar{\alpha}$	Q^*	$\bar{\alpha}$	exact gain
1,568	20.4	0.95	45.1	0.44	+20.4%	60.9	1.00	146.0	0.47	+27.4%
4,704	21.4	0.97	68.0	0.48	+32.7%	77.5	0.99	202.7	0.44	+31.1%
7,840	22.2	0.97	85.9	0.47	+38.5%	110.6	0.89	232.6	0.44	+24.8%
10,977	22.8	0.98	100.8	0.46	+42.3%	144.9	0.81	255.2	0.44	+18.9%
15,681	23.6	0.98	119.0	0.45	+46.2%	180.6	0.74	281.6	0.43	+14.8%
23,522	24.5	0.99	143.1	0.45	+50.4%	223.1	0.68	320.1	0.41	+11.5%
31,362	25.3	0.99	164.7	0.44	+53.3%	255.3	0.64	348.5	0.40	+9.6%

Notes. Bottom entry. The signal-only columns reproduce the main text's Table 4 exactly (same solver, $N_{\gamma^H} = 500$, $N_\alpha = 201$, precise policy via `minimize_scalar`, local Brent planner search around the main-text μ^* for each row). The exact Bayesian columns are computed as a *full* joint fixed point of the worker's Bellman, the stationary distribution, and the Riccati DARE (61), at the *same* γ_D as the signal-only columns; the exact Bayesian planner is solved by local Brent over μ in a neighborhood of the signal-only μ^*. Welfare gains are $(W_{\text{plan}} - W_{\text{CE}})/|W_{\text{CE}}|$ evaluated under each model's own Q^*. The baseline row ($\gamma_D = 1{,}568$) matches the bottom-entry column of Table 11 exactly by construction.

Three patterns emerge. First, and most importantly, the qualitative policy message of the main text is unchanged: under the exact Bayesian model, planner intervention still delivers a large flow welfare gain at every value of γ_D, ranging from +27.4% at the baseline to +9.6% at the top

of the range. Even the smallest value remains firmly in the high single digits, well above any reasonable threshold for "negligible," and the policy prescription—restrict AI use most for early-career workers, allow moderate use for experts—is robust at every γ_D.

Second, the absolute data quality Q^* under exact Bayesian rises much faster with γ_D than under signal-only. Under signal-only, the CE Q^* creeps from 20 to 25 as γ_D grows twenty-fold, because workers are trapped at $\alpha \approx 0.95$–0.99 and the cohort actions carry almost no fresh-signal content. Under exact Bayesian, the CE Q^* grows from 61 to 255, because the joint-filtering channel extracts information about θ_t from cohort action dynamics regardless of how high α is. Importantly, the worker's Bellman responds endogenously to the higher data quality: the exact Bayesian CE $\bar{\alpha}$ drops from 1.00 at the baseline to 0.64 at $20\times$, as workers find labor relatively more attractive when AI's marginal information advantage shrinks. The signal-only model never produces this endogenous switch in the comparative-statics range, because its CE never reaches a high enough Q^* to make labor competitive again.

Third, the welfare gain pattern under exact Bayesian is the *opposite* of the signal-only pattern. The signal-only gain rises monotonically from $+20.8\%$ at the baseline to $+53.3\%$ at $20\times$, because higher γ_D amplifies the planner's correction (every unit of labor-induced novel data is worth more) while leaving the trapped CE unchanged. The exact Bayesian gain peaks at the baseline ($+27.4\%$) and *declines* as γ_D grows, falling to $+9.6\%$ at $20\times$. The mechanism is that the exact Bayesian CE catches up to the planner as γ_D grows: with the joint-filtering channel active, the marginal informational advantage of switching workers from AI to labor diminishes when the CE is already extracting substantial information from cohort dynamics. The two welfare-gain curves cross between $1\times$ and $5\times$ the baseline, with the signal-only model giving the larger gain at high γ_D and the exact Bayesian model giving the larger gain at low γ_D. The two are within a few percentage points of each other at $3\times$ the baseline (sig $+32.7\%$, exact $+31.1\%$).

The two models therefore tell genuinely different stories about how the data-quality externality scales with recording precision. Both agree that the externality is quantitatively important at the calibration baseline—the welfare gains are double-digit in both cases—but they disagree about whether the externality becomes *more* or *less* severe as AI's recording precision improves. The signal-only formula treats the AI system as a pure sampling device, so its CE collapses informationally as workers concentrate on AI; this collapse gives the planner increasing leverage as

γ_D grows. The exact Bayesian formula treats the AI system as an ideal joint Bayesian filter, so its CE remains informationally rich even at high α; this richness erodes the planner's relative advantage as γ_D grows. Section E.5 below develops this contrast and explains which assumption is the appropriate modeling choice for current AI training technology.

E.5 How the two Φ_t formulas differ economically

The difference between the two models has a clean economic origin. Relaxing Assumption 2 endows the AI system with a second information channel—one that extracts information about θ_t from the dynamics of cohort prior means rather than only from fresh private signals. The externality identified in the main text does not disappear under this relaxation—workers still free-ride on the shared dataset in the CE, and the planner still improves outcomes by redirecting them toward labor—but the *level* of data quality in every allocation shifts upward, because the Riccati filter adds a joint-filtering channel that the signal-only reduced form treats as non-codified.

Under Assumption 2, the dataset treats its inference problem as a pure sampling exercise: the only thing it can learn from a recorded action is the fresh private signal ($s_t^{l,i}$ or $s_t^{A,i}$) that the worker drew that period. Everything else in her action—the prior component μ_t^i reflecting her accumulated expertise, the dataset echo μ_t^Ω, and the deterministic weights multiplying each of them—is treated as already codified or otherwise unavailable for further Bayesian updating. When a worker switches to full AI and stops producing fresh labor signals, her recorded action becomes informationally inert from the dataset's perspective, and the dataset's information inflow collapses toward a small residual proportional to $\alpha\,\gamma^a/(\gamma^H + Q + \gamma^a)$, which is tiny when AI adoption is near one. This is exactly why the signal-only CE produces such a low post-AI data quality under bottom entry ($Q^* = 20.2$).

The exact Bayesian treatment lifts this restriction by keeping the cohort prior mean $z_t(\tau)$ in the filter's state vector. Even when $\alpha = 1$ and no fresh labor signal is drawn, the cohort-averaged action still moves one-for-one with θ_t through the weights on $\bar{\mu}_t(\tau)$ and the echo μ_t^Ω—and the Kalman filter, by tracking the joint evolution of θ_t and $\{z_t(\tau)\}$, extracts information about θ_t from the cross-sectional variation of these action dynamics across ages. The AI system runs what is effectively a time-series inference on the cohorts' recorded behavior: it no longer needs fresh labor

signals to keep its estimate of θ_t up to date, because the joint filter can recover the common θ_t component from the way cohort actions comove around their age-specific priors. The accumulated expertise that each worker carries—a deterministic function of past θ's through her Kalman recursion—becomes partially recoverable from the dataset's perspective, and the effective information inflow per period is higher at every allocation. Under bottom entry, this raises the CE Q^* from 20.2 to 60.9 (a factor of 3.0) and the planner Q^* from 42.5 to 146.0 (a factor of 3.4) at the baseline γ_D.

Two economic implications follow. First, the joint-filtering channel lifts Q^* in both the CE and the planner allocation, but the lift is somewhat *larger* at the planner's allocation because the planner's labor-heavy mixture exposes the filter to a richer combination of signal-only and joint-filtering content. As a result, the ratio of planner-Q^* to CE-Q^* rises slightly from 2.10 under signal-only to 2.40 under exact Bayesian at the baseline, and the baseline flow welfare gain is correspondingly slightly larger under exact Bayesian (+27.4% vs. +20.8% under signal-only). Second, as γ_D grows the two models part ways in opposite directions. Under signal-only the CE remains informationally trapped at AI-heavy allocations—its Φ grows slowly with γ_D because the integrand is dominated by the tiny $\gamma^a/(\gamma^H + Q + \gamma^a)$ term—while the planner's labor-heavy allocation captures nearly all of the extra recording precision; the gap widens monotonically and the signal-only gain climbs to +53.3% at 20×. Under exact Bayesian, the CE is *not* informationally trapped because the joint-filtering channel keeps extracting information about θ_t regardless of how high α is; in fact, as γ_D grows the exact Bayesian CE catches up to the planner in Q^* (CE/planner ratio goes from 0.42 at 1× to 0.73 at 20×), and workers endogenously switch toward more labor in response to the higher data quality. The marginal value of the planner's correction therefore *shrinks* with γ_D under exact Bayesian, and the exact Bayesian gain falls from +27.4% at the baseline to +9.6% at 20×. The two welfare-gain curves cross between 1× and 5× the baseline (Table 12, Figure 16).

Which assumption is the right one depends on the sophistication of the real-world data-processing technology. If AI training pipelines treat each recorded action as an independent sample, extracting only the fresh-signal content (as is typical in current large-scale supervised and reinforcement-learning practice), then Assumption 2 is the appropriate modeling choice and the main-text welfare numbers apply directly. If AI systems are ideal joint Bayesian filters that track θ_t and cohort

priors over time and exploit cross-cohort action dynamics, then the exact Bayesian model gives the correct—and generally more conservative—welfare-gain estimate at high γ_D. The two models therefore bracket the quantitative importance of the data externality in the comparative-statics exercise, with the true welfare gain lying somewhere between them depending on how much joint filtering the data processor actually does. Qualitatively, however, the policy prescription is the same in both cases: the CE suffers from a large data-quality externality, planner intervention delivers double-digit flow welfare gains, and the monotone-declining optimal policy—restricting AI use most for early-career workers while allowing moderate use for experts—is robust to the choice of Φ_t formula.

To summarize, the exact Bayesian treatment confirms that the skill trap, the monotone-declining optimal policy, and the data-quality externality are all robust features of the primitive learning environment rather than artifacts of Assumption 2. At the baseline calibration the exact Bayesian flow welfare gain is in fact slightly larger than the signal-only gain (+27.4% vs. +20.8% for bottom entry, +17.4% vs. +6.5% for uniform entry). Across the γ_D comparative-statics range, the two models tell genuinely different stories: the signal-only gain rises monotonically with γ_D as the trapped CE falls further behind the planner, while the exact Bayesian gain peaks at the baseline and declines as the CE catches up to the planner via the joint-filtering channel. The two welfare-gain curves cross between $1\times$ and $5\times$ the baseline. Under either modeling choice, however, planner intervention delivers a double-digit flow welfare gain at every value of γ_D considered, and the policy prescription—restrict AI use most for early-career workers, allow moderate use for experts—is robust to the choice of Φ_t formula. The main-text numerical results are therefore internally consistent with whichever Φ_t formula a reader finds most plausible as a description of real AI training technology.