

Taking the Easy Way Out: AI, Self-Control, and Human Capital Formation*

Drew Fudenberg[†] David K. Levine[‡]

May 19, 2026

Abstract

Some researchers are concerned that AI use by junior researchers inhibits acquisition of human capital. We study a dynamic model with a short-term project having an immediate return and a long-term project having a larger delayed return. AI increases success with low human capital, but prevents learning by doing. Increasing available options ordinarily cannot decrease welfare, but if immediate benefits are tempting and self-control is costly, AI reduces welfare if it is not adopted, and can induce the researcher to switch to the short-term project. More surprisingly, the introduction of AI may increase the acquisition of human capital.

*We thank Daron Acemoglu, Joshua Gans, Andrew Koh, and Stephanie Wang for helpful comments and suggestions, and NSF grant SES-2417162 and the Leverhulme Trust for financial support.

[†]Department of Economics, MIT

[‡]Department of Economics, RHUL and Washington University

1 Introduction

An AI agent may give a junior researcher a better chance of producing a useful output. But as Hogg [2026] and Karamanis [2026] emphasize, the value of difficult research work lies not only in the output it generates, but also in the skills and judgment acquired by doing it. We share their concern that by replacing hard research work, AI can weaken the process through which researchers acquire judgment and skill: A researcher who uses AI to generate results and text can produce a passable paper while giving up the learning by doing that would make future work better.

This paper studies the decision of whether to delegate to an AI agent in a simple dynamic model where the researcher can either delegate the task to an AI agent with minimal human input or not use it at all. Each period, the researcher chooses between a short-term project and a long-term one. The short-term project provides an immediate return, while the long-term project only pays off later. In low-human-capital states, working without AI can generate learning by doing and lead to the high-human-capital state, whereas using AI reduces human-capital acquisition. In a standard model without a self-control problem, the availability of AI cannot lower the agent's payoff: They won't use it if it decreases their payoff, and if they don't use AI they are no worse off than if it were not available.

The welfare effect of AI is no longer straightforward when the agent is tempted to forgo future returns for immediate benefits and so has a cost of self control. To model this, we assume that the cost of resisting temptation depends on the best short-run payoff available in the current period, and that this cost may be incurred even if the AI is not used. By making a high-immediate-return option available, AI raises the temptation benchmark. As a result, the availability of AI can worsen welfare and distort project choice.

The issue is not specific to academic research. Many valuable skills are acquired while performing tasks that are initially slow, frustrating, or only imperfectly productive. AI changes the private return to those tasks by making it possible to obtain acceptable output without going through the costly process of skill formation. Thus the central question is not only whether AI raises current productivity, but also whether it changes the incentives to acquire the human capital on which future productivity depends. This makes AI different from a pure productivity shock: it changes both what can be done today and what the worker has reason to learn for tomorrow.

The interaction of AI and self-control gives rise to four distinct effects.

- **Blowing off:** the researcher uses AI to carry out the long-term project in the low-human-capital state and forgoes the accumulation of human capital.
- **Discouragement:** when AI becomes available, the researcher uses AI to carry out the short-term project in the low-human-capital state even though the long-term project would have been chosen if AI were unavailable.
- **Reversal:** the arrival of AI causes the researcher to switch from the long-term project to the short-term project in the low-human-capital state even though AI is not used.
- **Encouragement:** when AI becomes available, the researcher acquires human capital even though, without AI, the researcher would have remained in the low-human-capital state.

The first three effects concern how AI changes behavior within the low human capital state. In blowing off, the researcher continues to choose the long-term project, but does so with AI and therefore gives up the learning by doing that would accompany working without AI. In discouragement, the arrival of AI changes project choice: long-term work would have been chosen without AI, but with AI available the researcher switches to the short-term project. In reversal, AI changes the ranking of the two non-AI tasks. The researcher switches from the long-term project to the short-term project even though AI is not used, because the availability of AI changes the self-control cost of doing long-term projects without it.

Encouragement concerns a different margin. Without AI, the researcher may choose a long-term project and still remain in the low-human-capital state. With AI, the value of remaining in that state changes, and the researcher may instead choose a path that leads to high human capital. Thus the analysis distinguishes effects of AI on current project choice from effects of AI on the continuation value of remaining low skilled.

While our model is highly stylized, the results follow from three key assumptions. The first is that using AI is better than working without it when human capital is low and worse when human capital is high. The second is that human capital accumulation is enhanced by doing it yourself. A third assumption, that there is an increasing marginal cost of self-control, is needed for the reversal result because the independence of irrelevant alternatives is satisfied when the marginal cost of self-control is constant.

Related work. Becker [1962] frames education and training as forward-looking investments that trade off current costs against future returns. Mincer [1962] extends that logic to

on-the-job training and explicitly treats learning from experience as human-capital investment, and Ben-Porath [1967] embeds these choices in a dynamic life-cycle model in which individuals allocate effort and time between current production and skill accumulation. We introduce AI into the human capital literature as a technology that can raise current productivity while simultaneously introducing a self-control problem. Much of the economics of automation studies how new technologies substitute for or complement existing skills, see, for example, Autor, Levy, and Murnane [2003], Autor [2015], Acemoglu and Restrepo [2018], and Agrawal, Gans, and Goldfarb [2018]. Our focus is instead on how technology changes the incentive to acquire those skills in the first place.

Acemoglu, Kong, and Ozdaglar [2026] studies how agentic AI can reduce human learning effort and erode collectively produced knowledge. Ide [2025] studies how automating entry-level tasks can weaken the intergenerational transmission of tacit knowledge. Both papers share our concern that AI can weaken learning, but their mechanisms work through social knowledge transmission and ours is instead an individual self-control problem.

Garicano and Rayo [2017] and Fudenberg and Rayo [2019] study knowledge transfer in apprenticeships; Fudenberg, Georgiadis, and Rayo [2021] adds an explicit learning by doing constraint so that productive work is a vehicle for skill acquisition. These papers all feature contractual frictions between a trainer and an apprentice, while our model features an intrapersonal friction: AI changes the menu of current payoffs, increases temptation, so it can discourage skill-building even in the absence of an external supervisor or contract. Shen and Tamkin [2026] provides complementary evidence from a randomized experiment with software engineers: AI assistance reduced understanding and skills without delivering significant efficiency gains, with the sharpest skill losses among participants who fully delegated tasks to AI rather than remaining cognitively engaged.

Our analysis is based on past theoretical studies of temptation and self-control. Gul and Pesendorfer [2001] provides an axiomatic characterization of a model where self-control costs are proportional to the difference between the maximum currently attainable short-run payoff and the payoff of the chosen action. This paper instead uses the dual-self approach of Fudenberg and Levine [2006], which emphasizes that self-control is often a strictly convex function of foregone short-run utility. With such convex costs, enlarging the opportunity set can change the agent's preferences over previously available options. By contrast, with constant marginal self-control costs, a new tempting alternative shifts the payoff of the

non-AI actions uniformly, and the reversal result disappears.¹

We also add to the literature on skill erosion. Snower [1996] argues that sectors can become stuck in “low-skill, bad job” traps where low productivity and weak training incentives reinforce one another. Similarly, our reversal result shows how AI can shift behavior away from skill-building work and help sustain a low-skill state. Our blowing-off result provides an individual-level analog to this phenomenon: once AI makes it possible to obtain output without learning, the private incentive to build expertise is weakened, potentially sustaining a low-skill equilibrium. This logic extends to the “shadow learning” observed by Beane [2019] in robotic surgery, where trainees must navigate organizational constraints to get the hands-on experience needed to learn. More generally, our model shows how a technology that raises short-run performance can also weaken the process through which expertise is formed.

Recent experimental evidence validates our concern that technologies that boost short-term productivity can weaken the skill-building process on which long-term expertise depends, creating the risks of skill erosion and “blowing off” that we identify. Dell’Acqua et al. [2023] identifies a “jagged technological frontier” where over-reliance on AI leads to systematically flawed output when tasks fall outside the tool’s capabilities. This over-reliance mirrors our discouragement effect, where the immediate return from AI outweighs the incentive to build the human judgment necessary to navigate the frontier. Similarly, Bastani et al. [2025] shows that unrestricted AI access can undermine performance on subsequent exams.²

2 Basic Model

Time is discrete, $t = 1, 2, \dots$, and the researcher discounts with factor $0 \leq \delta < 1$. Human capital can be High (H) or Low (L), and AI may be available (A) or not (N). Although in principle there would be four states, HN, HA, LN and LA , we simplify by assuming that

¹O’Donoghue and Rabin [1999] show that present bias can cause agents to procrastinate on tasks that require bearing an immediate cost for a delayed return. Our model shares that basic tension, but the introduction of AI changes the menu of available actions and thereby changes the cost of self-control, generating effects that do not arise in a present-bias framework.

²The paper also shows that a correctly structured AI tutor can improve exam performance. Our model does not distinguish between different forms of AI, so it cannot capture this effect.

AI is not used when human capital is high and condense HN and HA to H .³ Each period the researcher chooses between two projects. The short-term project yields a current payoff of 1 upon success and no future payoff. The long-term project yields no current payoff, but it generates a next-period project payoff $\Gamma > 1$ upon success. In state LA the researcher can choose whether to perform the task without AI or else let the AI do it. The feasible actions are thus SN, TN, SA , and TA .

Success probabilities are as follows. In the high-human-capital state success is certain. In the low-human-capital state without AI the probability of success is p^N , and in the low-human-capital state with AI the probability of success is p^A , where $0 < p^N < p^A < 1$. If the researcher works on their own in a low-human-capital state, at the end of the period they can choose whether to remain in the low-human-capital state or move to the high-human-capital state next period. If the researcher uses AI in state LA , then next period's state is again LA . In the high-human-capital state the researcher can choose whether to retain their human capital or go back to the low human capital state.

The researcher incurs a quadratic self-control cost from resisting immediate payoff. Let τ denote the highest achievable expected current return in the period. Thus $\tau = p^N$ in state LN , $\tau = p^A$ in state LA , and $\tau = 1$ in state H . If a choice yields current return r_t and next-period project payoff r_{t+1} , then the action value is

$$u_t = r_t - \alpha(\tau - r_t)^2 + \delta r_{t+1},$$

where $\alpha \geq 0$ measures the severity of the self-control problem. Thus the self-control cost is determined by the gap between the current payoff of the chosen action and the highest current payoff available in that state.

Let $G := \delta\Gamma$ be the current value of getting Γ next period.

The resulting action values are shown in Table 1.

³We have shown that the four effects we identify also arise in a more general model that has 4 states, HN, HA, LN and LA and allows for high human capital to decay to low when AI is used (“use it or lose it.”)

State	u^{SN}	u^{TN}	u^{SA}	u^{TA}
H	1	$G - \alpha$	–	–
LN	p^N	$p^N G - \alpha(p^N)^2$	–	–
LA	$p^N - \alpha(p^A - p^N)^2$	$p^N G - \alpha(p^A)^2$	p^A	$p^A G - \alpha(p^A)^2$

Table 1: Payoffs by state and strategy profile.

3 Overview of the Results

Before characterizing the optimal decision, we will explain three key properties of the solution. First, conditional on the state, the value of G , and whether or not AI is used, the choice of which project to choose is purely myopic and does not depend on the discount factor δ . Second, with self-control costs, the high-human-capital state can be worse than the low-human-capital state, even when AI is unavailable. Finally, conditional on G , the discount factor matters only for the choice of using AI in the low human capital state.

The project-choice comparison is myopic because, once the state and the technology are fixed, the future consequences of the short-term and long-term projects are the same. Thus the short-term project is preferred to the corresponding long-term one whenever the easy and long-term projects depends on whether the control cost α is above or below a cutoff value. These cutoffs are shown in Table 2.

Comparison	Condition	Cutoff
$u_{LN}^{SN} > u_{LN}^{TN}$	$\alpha > \bar{\alpha}_L$	$\bar{\alpha}_L = \frac{G - 1}{p^N}$
$u_{LA}^{SA} > u_{LA}^{TA}$	$\alpha > \bar{\alpha}_A$	$\bar{\alpha}_A = \frac{G - 1}{p^A}$
$u_{LA}^{SN} > u_{LA}^{TN}$	$\alpha > \bar{\alpha}_N$	$\bar{\alpha}_N = \frac{G - 1}{2p^A - p^N}$
$u_H^{SN} > u_H^{TN}$	$\alpha > \bar{\alpha}_H$	$\bar{\alpha}_H = G - 1$

Table 2: Static project comparisons and cutoff values.

To see why the high human capital state can be worse even when AI is unavailable,

consider the myopic payoff in the high and low human capital states. The high state raises productivity, but it also raises temptation: in state H the tempting current payoff is 1, while in a low-human-capital state the tempting payoff is p^N without AI and p^A with AI. Let $\hat{u}_H := \max\{1, G - \alpha\}$, $\hat{u}_L := \max\{p^N, p^N G - \alpha(p^N)^2\}$, $\hat{u}_A := \max\{u_{LA}^{SA}, u_{LA}^{TA}\}$, and $\hat{u}_N := \max\{u_{LA}^{SN}, u_{LA}^{TN}\}$. Since temptation lowers utility, \hat{u}_H can be lower than the corresponding low-state payoff, and in this case the researcher may prefer low human capital despite its lower productivity.

Finally, conditional on G , the discount factor matters only for the choice of using AI in the low-human-capital state because it affects only the value of the continuation. In particular, δ changes the value of the no-AI path in the low-human-capital state, since that path may lead to high human capital. Finally, the reason that, conditional on G , the discount factor matters only for the choice of using AI in the low human capital state is that the discount factor enters only through the continuation choice. What it changes is the value of the no-AI path in the low-human-capital state, because that path may lead to high human capital. Thus a higher δ makes the researcher more willing to give up the current AI payoff in order to move toward H, while a lower δ makes AI more attractive.

Having laid out the basic considerations, we now give our main results. We first give the basic properties of the equilibria, then lay out the comparative statics of when each of the four effects can occur.

Proposition 1 (Basic properties).

- (i) *If $\alpha < \min\{\bar{\alpha}_H, \bar{\alpha}_N\}$, only the long-term project is chosen in every state.*
- (ii) *If $\alpha > \bar{\alpha}_L$, only the short-term project is chosen in every state.*
- (iii) *In the low state with AI available, there is a discount factor $\hat{\delta} \in (0, 1]$ such that for $\delta < \hat{\delta}$ AI is used, and for $\delta > \hat{\delta}$ AI is not used. If G is sufficiently large, then $\hat{\delta} < 1$.*

The logic is straightforward and follows the discussion above. Parts (i) and (ii) come from the static project rankings. Part (iii) concerns the low-human-capital state with AI: using AI raises current payoff, while not using AI can lead to H , so as δ rises the value of working without AI rises, and AI use falls whenever the transition to H is valuable enough.

Recall the four effects identified in the Introduction: **blowing off**, where the researcher uses AI to complete the long-term project while in the low-human-capital state, thereby forgoing the accumulation of human capital; **discouragement**, where the researcher uses AI

to complete the short-term project in the low-human-capital state, even though they would have chosen the long-term project had AI been unavailable; **reversal**, where the arrival of AI causes the researcher to switch from the long-term project to the short-term project in the low-human-capital state, even though AI is not actually used; and **encouragement**, where the availability of AI induces the researcher to acquire human capital, even though they would have remained in the low-human-capital state without it. The next table gives a rough sense of when each effect occurs.

Effect	α -range	Discount factor	AI used?
Blowing off	$\alpha < \bar{\alpha}_A$	low δ	Yes
Discouragement	$\bar{\alpha}_A < \alpha < \bar{\alpha}_L$	low δ	Yes
Reversal	intermediate range	high δ	No
Encouragement	intermediate range	high δ	No

Table 3: Summary of the four effects.

The four effects come from two margins. One is project choice within the current state: AI can make short-term work more attractive, or make using AI to do long-term work more attractive than doing long-term work without AI. The other is the continuation value of remaining in the low-human-capital state relative to moving to H . Blowing off and discouragement arise when AI makes behavior more myopic, reversal arises when AI changes the ranking of the two non-AI actions, and encouragement arises when AI changes the value of remaining in the low state relative to acquiring high human capital.

The first two rows are the cases in which AI is used in the low-human-capital state. For blowing off, $\alpha < \bar{\alpha}_A$ implies that TA is the best AI action in state LA ; for sufficiently low δ , the current payoff from using AI dominates the continuation value from working without AI and moving to H , so TA is optimal. Thus the researcher still chooses the long-term project, but delegates it to AI and remains in the low-human-capital state.

For discouragement, $\bar{\alpha}_A < \alpha < \bar{\alpha}_L$ means that TN is chosen in state LN , while SA beats TA among the AI actions in state LA . Let $\hat{u}_N := \max\{u_{LA}^{SN}, u_{LA}^{TN}\}$. If, in addition, $u_{LA}^{SA} > \hat{u}_N$, then for sufficiently low δ the short-term AI action also beats the best no-AI action followed by transition to H . Hence SA is optimal in state LA , even though TN would have been chosen if AI were unavailable. The formal proof of the two AI-use effects is given in Proposition B.1 in Appendix B.

In the last two rows, the availability of AI changes behavior even though AI is not used in the relevant low-human-capital state. In reversal, the self-control cost created by the AI option makes the short-term non-AI action more attractive than the long-term non-AI action in state LA , even though the long-term project would have been chosen in state LN . In encouragement, AI changes the value of remaining low-skilled in a way that can make transition to H more attractive. These effects require intermediate self-control costs and are favored by high δ , because a patient researcher puts more weight on the continuation value associated with working without AI. The next two propositions give the precise conditions for reversal and encouragement.

Proposition 2 (Reversal).

(i) *Reversal occurs when TN is chosen in state LN but SN is optimal in state LA . This requires $\bar{\alpha}_N < \alpha < \bar{\alpha}_L$.*

(ii) *If*

$$(1 - \delta)(p^N - \alpha(p^A - p^N)^2) + \delta \hat{u}_H > \max\{p^A, p^A G - \alpha(p^A)^2\},$$

then reversal occurs.

The key to Proposition 2 is that the availability of AI raises the temptation benchmark enough to reverse the ranking of the two non-AI actions in state LA . The researcher then chooses the short-term project without AI, even though the long-term project would have been chosen if AI were unavailable.

Proposition 3 (Encouragement). *A sufficient condition for encouragement for all δ sufficiently close to 1 is*

$$Gp^N - \alpha(p^N)^2 > 1 > Gp^A - \alpha(p^A)^2.$$

Such parameter values exist if and only if

$$G > \frac{1}{p^N} + \frac{1}{p^A}.$$

Proposition 3 gives the opposite kind of effect. Its condition says that, without AI, remaining in the low-human-capital state is attractive, while with AI the best AI-assisted low-state payoff is not attractive enough to dominate moving to H . For sufficiently high discount factors, the researcher therefore chooses a no-AI path to high human capital. In this case, AI encourages human-capital acquisition by lowering the relative value of staying low-skilled.

4 Discussion of the Four Effects

This section discusses the economic mechanisms behind the four effects.

4.1 Blowing off

Blowing off occurs when the researcher uses AI for the difficult task in the low-human-capital state, thereby forgoing human-capital acquisition. Thus the researcher still wants the payoff from the long-term project, but prefers to obtain it through AI rather than through working without AI.

Economically, the key force is that AI raises current productivity on the difficult task while eliminating the learning by doing that would otherwise come from doing that task unaided. This is the most direct way in which AI can weaken skill formation: the researcher does not switch away from the long-term project, but switches away from the skill-building version of it.

This effect arises when using AI on difficult work is attractive enough both relative to using AI for short-term projects and relative to the best no-AI path that leads to high human capital. In terms of the model, this requires that TA beat both SA and the best non-AI action followed by H . So the economic tradeoff is between higher current performance with AI and the future value of learning. The consequence is that the researcher remains in the low-human-capital state even while continuing to pursue the long-term project. Because any optimal non-AI action in state LA is followed by transition to H , whereas TA keeps the researcher in LA , blowing off replaces a path with learning by doing by one without it. The distortion is not that the researcher stops aiming high, but that they pursue the same difficult objective in a way that prevents skill accumulation.

4.2 Discouragement

Discouragement occurs when without the availability of AI the difficult task is chosen in the low-human-capital state, but when AI is made available there is a switch to the easy task with AI. This differs from blowing off because the availability of AI now changes the project itself, not just the technology used to complete it. Under discouragement, the researcher gives up on difficult work and moves to the short-term project. The economic mechanism is that AI raises the value of immediate-return options, and this makes the short-term project with AI more attractive than the long-term project without AI.

For discouragement to occur, the long-term project must be preferred in state LN , which requires $\alpha < \bar{\alpha}_L$, while in state LA the easy AI action must beat the difficult AI action, which requires $\alpha > \bar{\alpha}_A$. Thus discouragement occurs only for the intermediate range $\bar{\alpha}_A < \alpha < \bar{\alpha}_L$. On that interval, the self-control problem is strong enough that the easy AI action becomes more attractive than the difficult AI action, but not so strong that the short-term project would already have been chosen even without AI. The remaining condition is that SA also beat using the best non-AI action and transitioning to H .

Economically, discouragement is a distortion on the project-choice margin that redirects effort toward the easy task. This lowers eventual human capital only if working without AI in LN would otherwise have led to transition to H . If the no-AI choice in LN would already have been followed by remaining in LN , discouragement changes project choice but not the eventual level of human capital.

4.3 Reversal

Reversal occurs when without the availability of AI the difficult task is chosen in the low-human-capital state, but when AI is made available there is a switch to the easy task without AI, so that AI changes behavior even though it is not used. The logic is similar to the violations of WARP discussed in Fudenberg and Levine [2006]: increasing the highest immediate return increases self-control costs. Here this changes the ranking of the two non-AI actions.

For reversal to occur, the long-term project must be chosen in state LN , so $\alpha < \bar{\alpha}_L$, but conditional on not using AI in state LA , the easy project must beat the difficult one, so $\alpha > \bar{\alpha}_N$. Thus reversal requires the intermediate range $\bar{\alpha}_N < \alpha < \bar{\alpha}_L$. The economic content of this interval is that self-control costs are not large enough to make the researcher avoid the long-term project in the no-AI environment, but become large enough to overturn that ranking once AI is added to the menu.

4.4 Encouragement

Encouragement occurs when no human capital is acquired in the low state without AI, but human capital is acquired in the low state when AI is available. This is the most surprising effect, because it goes against the simple intuition that AI must always weaken incentives to learn. The key point is that AI changes not only the attractiveness of current actions,

but also the value of remaining in the low-human-capital state. Without AI, the researcher may prefer to remain in LN because moving to H raises future temptation too much. With AI, that low-human-capital state becomes less attractive as a place to remain, and this can make transition to H optimal.

Formally, encouragement requires that without AI the agent chooses TN and remains in LN . By Lemma A.2, this means

$$\frac{G}{1 + p^N} < \alpha < \frac{p^N G - 1}{(p^N)^2}.$$

So, absent AI, the long-term project is attractive enough that the researcher does not switch to the short-term project, but the value of remaining in the low state still exceeds the value of moving to H . Encouragement requires that taking the best non-AI action and moving to H is better than having the AI do either action. Thus AI must make the low state unattractive enough that the researcher now prefers skill acquisition. The encouragement effect is welfare-reducing for the researcher: AI induces human-capital acquisition by lowering the value of remaining in the low-human-capital state, not by raising the value of acquiring human capital.

A useful sufficient condition for encouragement for δ close to 1 is

$$Gp^N - \alpha(p^N)^2 > 1 > Gp^A - \alpha(p^A)^2.$$

The left inequality says that without AI the low state is attractive enough that the researcher prefers to remain there. The right inequality says that once AI is available, the AI-assisted low-state option is no longer attractive enough to dominate transition to H . This sufficient condition can hold only if

$$G > \frac{1}{p^N} + \frac{1}{p^A},$$

so encouragement requires the delayed return to difficult work to be large.

5 Conclusion

AI does more than raise productivity. By increasing the best available short-run payoff, it also changes temptation. The model shows that this reduces welfare when behavior does not change, and it can distort project choice in ways that slow or prevent the accumulation of human capital. For applications, what matters is thus not only whether AI makes low-skill workers more productive, but also whether AI reduces the incentive to work on tasks through which expertise is built and maintained.

The model also distinguishes four forms of distortion. Blowing off occurs when the researcher uses AI on the long-term project and thereby avoids the learning by doing that would have raised future productivity. Discouragement occurs when AI induces a switch from non-assisted difficult work to AI-assisted short-term projects. Reversal means that AI changes the ranking of unassisted actions and induces a switch from the long-term project to the short-term project even though AI is never used. Encouragement means that the effect of AI goes in the opposite direction: it induces the acquisition of human capital that would not have occurred in its absence.

Our analysis shows that unrestricted early access to AI can weaken the incentives to acquire human capital. This suggests a role for training rules, staged access, and institutional designs that preserve learning-by-doing while still allowing workers to benefit from AI. The point is not that AI use is undesirable in general. Once the relevant human capital has been acquired, AI can be valuable because as Agrawal, Gans, and Goldfarb [2025] emphasizes, it can complement judgment even as it substitutes for implementation.

A Auxiliary lemmas

Lemma A.1 (Static project rankings). *The cutoffs satisfy*

$$\bar{\alpha}_L > \bar{\alpha}_A > \bar{\alpha}_N, \quad \bar{\alpha}_N \geq \bar{\alpha}_H \iff 2p^A - p^N \leq 1.$$

Proof. Since $p^N < p^A < 2p^A - p^N$, we have $\bar{\alpha}_L > \bar{\alpha}_A > \bar{\alpha}_N$. Also, $\bar{\alpha}_N \geq \bar{\alpha}_H$ is equivalent to $(G - 1)/(2p^A - p^N) \geq G - 1$, which is equivalent to $2p^A - p^N \leq 1$. \square

Lemma A.2 (No-AI values and transition from LN). *Let*

$$\hat{u}_H := \max\{1, G - \alpha\}, \quad \hat{u}_L := \max\{p^N, p^N G - \alpha(p^N)^2\}.$$

The normalized values in the high state and in state LN satisfy

$$\hat{u}_H = (1 - \delta)\hat{u}_H + \delta \max\{W_H, W_{LN}\}, \quad W_{LN} = (1 - \delta)\hat{u}_L + \delta \max\{W_H, W_{LN}\}.$$

Hence $W_H \geq W_{LN}$ if and only if $\hat{u}_H \geq \hat{u}_L$. After a non-AI action in LN , the researcher remains in LN if and only if

$$\hat{u}_L > \hat{u}_H,$$

which is equivalent to

$$\frac{G}{1 + p^N} < \alpha < \frac{p^N G - 1}{(p^N)^2}.$$

Proof. In state H , the researcher first obtains the best one-period payoff $\hat{u}_H = \max\{1, G - \alpha\}$ and then chooses whether to retain high human capital or return to the low-human-capital state. Hence

$$W_H = (1 - \delta)\hat{u}_H + \delta \max\{W_H, W_{LN}\}.$$

Similarly, in state LN the researcher obtains the best no-AI one-period payoff $\hat{u}_L = \max\{p^N, p^N G - \alpha(p^N)^2\}$ and then chooses whether to remain in LN or move to H , so

$$W_{LN} = (1 - \delta)\hat{u}_L + \delta \max\{W_H, W_{LN}\}.$$

Subtracting the two equations gives

$$W_H - W_{LN} = (1 - \delta)(\hat{u}_H - \hat{u}_L).$$

Thus $W_H \geq W_{LN}$ if and only if $\hat{u}_H \geq \hat{u}_L$, and remaining in LN is optimal if and only if $\hat{u}_L > \hat{u}_H$.

Because $p^N < 1$, this inequality is equivalent to

$$\frac{G}{1 + p^N} < \alpha < \frac{p^N G - 1}{(p^N)^2}.$$

□

Lemma A.3. *In state LA , it is never optimal to choose SN or TN and remain in LA next period. Therefore*

$$W_{LA} = \max\{\hat{u}_A, (1 - \delta)\hat{u}_N + \delta\hat{u}_H\}.$$

Proof. In state LA , the feasible action–continuation pairs are (SN, H) , (SN, LA) , (TN, H) , (TN, LA) , (SA, LA) , and (TA, LA) .

Since $p^A > p^N$, if the continuation state is fixed at LA , the AI version of each project strictly dominates the corresponding non-AI version. Hence no optimal policy chooses SN or TN and remains in LA , so $\hat{u}_A > \hat{u}_N$. Since any optimal no-AI action in state LA must be followed by transition to H , the best no-AI value is $(1 - \delta)\hat{u}_N + \delta W_H$. If instead an AI action is optimal, then $W_{LA} = (1 - \delta)\hat{u}_A + \delta W_{LA}$, so $W_{LA} = \hat{u}_A$. Combining the two cases gives $W_{LA} = \max\{\hat{u}_A, (1 - \delta)\hat{u}_N + \delta\hat{u}_H\}$. \square

Lemma A.4 (AI use in the low-human-capital state). *Suppose AI is available in state LA . An AI action is strictly optimal in state LA if and only if $\hat{u}_A > (1 - \delta)\hat{u}_N + \delta\hat{u}_H$. Moreover, if $\hat{u}_A \geq \hat{u}_H$, an AI action is optimal for every δ , and if $\hat{u}_N < \hat{u}_A < \hat{u}_H$, there is a unique cutoff*

$$\hat{\delta} = \frac{\hat{u}_A - \hat{u}_N}{\hat{u}_H - \hat{u}_N} \in (0, 1)$$

such that an AI action is strictly optimal for $\delta < \hat{\delta}$, a no-AI action followed by transition to H is strictly optimal for $\delta > \hat{\delta}$, and the researcher is indifferent at $\delta = \hat{\delta}$.

Proof. If $\hat{u}_A > \hat{u}_H$ then AI is used with low human capital regardless of the discount factor. Otherwise, define $w_A = \hat{u}_A$ to be the optimal average present value conditional on using AI with low human capital. Similarly, let

$$w_N = \hat{u}_N^A + \delta(\hat{u}_H - \hat{u}_N^A)$$

be the value conditional on doing it yourself and acquiring human capital. Consequently, AI is strictly worse if

$$\delta > \frac{\hat{u}_A - \hat{u}_N^A}{\hat{u}_H - \hat{u}_N^A} = \hat{\delta}.$$

For G large, $\hat{u}_A < \hat{u}_H$ and the long-term project is always optimal. Hence, as $G \rightarrow \infty$,

$$\hat{\delta} \rightarrow \frac{p_A - p_N}{1 - p_N} < 1.$$

\square

Lemma A.5 (Sufficient condition for encouragement near $\delta = 1$). *Suppose*

$$Gp^N - \alpha(p^N)^2 > 1 > Gp^A - \alpha(p^A)^2.$$

Then, for all δ sufficiently close to 1, encouragement occurs.

Proof. Let $f(x) := Gx - \alpha x^2$. The hypothesis says $f(p^N) > 1 > f(p^A)$. Since $p^N < p^A < 1$ and f is concave, $f(p^A) < f(p^N)$ implies that f is decreasing on $[p^A, 1]$, so $f(1) < f(p^A) < 1$.

Thus in state H the best one-period payoff is $\hat{u}_H = 1$, while in state LN the best one-period payoff is $\hat{u}_L = f(p^N) > 1$. By Lemma A.2, without AI the researcher chooses TN and remains in LN .

Now consider state LA . The best short-run payoff using AI $\hat{u}_A = \max\{p^A, f(p^A)\} < 1$. By Lemma A.2, since $\hat{u}_L > \hat{u}_H$, the value of reaching state H is $W_H = (1 - \delta)\hat{u}_H + \delta\hat{u}_L$. Thus $W_H \rightarrow \hat{u}_L = f(p^N) > 1$ as $\delta \uparrow 1$. Therefore the value of the best no-AI choice followed by transition to H also converges to $f(p^N) > 1$. Since $\hat{u}_A < 1$, for all sufficiently high δ the best no-AI choice followed by transition to H is better than using AI. Thus, with AI available, the researcher acquires human capital, whereas without AI the researcher remains in LN . \square

B Proofs of the propositions

Proof of Proposition 1. Parts (i) and (ii) follow immediately from Lemma A.1 and the ordering $\bar{\alpha}_L > \bar{\alpha}_A > \bar{\alpha}_N$, together with $\bar{\alpha}_L > \bar{\alpha}_H$. Part (iii) follows from Lemma A.4. \square

Proposition B.1 (AI-use effects). *Fix G .*

- (i) *If $\alpha < \bar{\alpha}_A$, then TA is the best AI action in state LA , and blowing off occurs for all sufficiently low δ .*
- (ii) *If $\bar{\alpha}_A < \alpha < \bar{\alpha}_L$ and $u_{LA}^{SA} > \hat{u}_N$, then TN is chosen in state LN , while SA is optimal in state LA for all sufficiently low δ .*

Proof. Part (i): By Lemma A.3, blowing off occurs if TA is optimal in state LA : in that case the researcher chooses the long-term project using AI and remains in state LA . The condition $TA > SA$ is $\alpha < \bar{\alpha}_A$. We now show that this same inequality also implies that TA beats both no-AI actions in current payoff. The comparison with TN is immediate, since $u_{LA}^{TA} - u_{LA}^{TN} = (p^A - p^N)G > 0$. For SN , TA beats SN by direct algebra under $\alpha < \bar{\alpha}_A$. Hence if $\alpha < \bar{\alpha}_A$, then TA is the best current-period action in state LA , so $\hat{u}_A = u_{LA}^{TA}$ and $\hat{u}_A > \hat{u}_N$.

By Lemma A.4, if $\hat{u}_A \geq W_H$, an AI action is optimal for every δ . If instead $\hat{u}_A < W_H$, then $\hat{u}_N < \hat{u}_A < W_H$, so there is a cutoff $\hat{\delta} \in (0, 1)$ such that an AI action is strictly optimal for $\delta < \hat{\delta}$. Since under $\alpha < \bar{\alpha}_A$ the optimal AI action is TA , blowing off occurs for all sufficiently low δ .

Part (ii): Discouragement means that TN is chosen in state LN but SA is optimal in state LA . When $\alpha < \bar{\alpha}_L$ TN is chosen in state LN , while $\alpha > \bar{\alpha}_A$ implies that SA beats TA among the AI actions in state LA . Thus on the interval $\bar{\alpha}_A < \alpha < \bar{\alpha}_L$, the best AI action in state LA is SA , so $\hat{u}_A = u_{LA}^{SA}$.

The condition $u_{LA}^{SA} > \hat{u}_N$ implies $\hat{u}_A > \hat{u}_N$. By Lemma A.4, SA is optimal in state LA for all sufficiently low δ . Hence AI induces a switch from TN in state LN to SA in state LA , which is discouragement. \square

Proof of Proposition 2. (i) Because TN is chosen in state LN if and only if $\alpha < \bar{\alpha}_L$, and SN beats TN if and only if $\alpha > \bar{\alpha}_N$, reversal requires $\bar{\alpha}_N < \alpha < \bar{\alpha}_L$.

(ii) On the interval $\bar{\alpha}_N < \alpha < \bar{\alpha}_L$, the best no-AI action in state LA is SN . By Lemma A.3, any optimal no-AI action in state LA is followed by transition to H , so the value of choosing SN in state LA is $(1 - \delta)(p^N - \alpha(p^A - p^N)^2) + \delta\hat{u}_H$. If this value exceeds both AI payoffs, p^A and $p^A G - \alpha(p^A)^2$, then SN is optimal in state LA . Since TN is chosen in state LN , reversal occurs. \square

Proof of Proposition 3. The sufficient condition follows from Lemma A.5. It remains to show when such parameter values exist. The inequalities $Gp^N - \alpha(p^N)^2 > 1 > Gp^A - \alpha(p^A)^2$ are equivalent to $\frac{Gp^A - 1}{(p^A)^2} < \alpha < \frac{Gp^N - 1}{(p^N)^2}$.

Finally, this interval is nonempty if and only if $\frac{Gp^A - 1}{(p^A)^2} < \frac{Gp^N - 1}{(p^N)^2}$, that is, if and only if $G > \frac{p^N + p^A}{p^N p^A} = \frac{1}{p^N} + \frac{1}{p^A}$. \square

References

- Acemoglu, D., D. Kong, and A. Ozdaglar (2026). *AI, Human Cognition and Knowledge Collapse*. Working Paper. National Bureau of Economic Research.
- Acemoglu, D. and P. Restrepo (2018). “The race between man and machine: Implications of technology for growth, factor shares, and employment”. *American economic review*, 108, 1488–1542.

- Agrawal, A., J. Gans, and A. Goldfarb (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston, MA: Harvard Business Review Press.
- Agrawal, A. K., J. S. Gans, and A. Goldfarb (2025). *The Economics of bicycles for the mind*. Tech. rep. National Bureau of Economic Research.
- Autor, D. H. (2015). “Why Are There Still So Many Jobs? The History and Future of Workplace Automation”. *Journal of Economic Perspectives*, 29, 3–30.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration”. *The Quarterly Journal of Economics*, 118, 1279–1333.
- Bastani, H. et al. (2025). “Generative AI without guardrails can harm learning: Evidence from high school mathematics”. *Proceedings of the National Academy of Sciences*, 122, e2422633122.
- Beane, M. (2019). “Shadow learning: Building robotic surgical skill when approved means fail”. *Administrative Science Quarterly*, 64, 87–123.
- Becker, G. S. (Oct. 1962). “Investment in Human Capital: A Theoretical Analysis”. *Journal of Political Economy*, 70. Part 2: Investment in Human Beings, 9–49.
- Ben-Porath, Y. (Aug. 1967). “The Production of Human Capital and the Life Cycle of Earnings”. *Journal of Political Economy*, 75, 352–365.
- Dell’Acqua, F. et al. (2023). “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality”. *Harvard Business School Technology & Operations Management Unit Working Paper*.
- Fudenberg, D., G. Georgiadis, and L. Rayo (Oct. 2021). “Working to Learn”. *Journal of Economic Theory*, 197, 105347.
- Fudenberg, D. and D. K. Levine (2006). “A dual-self model of impulse control”. *American economic review*, 96, 1449–1476.
- Fudenberg, D. and L. Rayo (Nov. 2019). “Training and Effort Dynamics in Apprenticeship”. *American Economic Review*, 109, 3780–3812.
- Garicano, L. and L. Rayo (Sept. 2017). “Relational Knowledge Transfers”. *American Economic Review*, 107, 2695–2730.

- Gul, F. and W. Pesendorfer (2001). “Temptation and self-control”. *Econometrica*, 69, 1403–1435.
- Hogg, D. W. (2026). “Why do we do astrophysics?” *arXiv preprint arXiv:2602.10181*.
- Ide, E. (2025). *Automation, AI, and the Intergenerational Transmission of Knowledge*. Discussion Paper. CEPR.
- Karamanis, M. (2026). *The Machines Are Fine*. URL: <https://ergosphere.blog/posts/the-machines-are-fine/> (visited on 04/26/2026).
- Mincer, J. (Oct. 1962). “On-the-Job Training: Costs, Returns, and Some Implications”. *Journal of Political Economy*, 70. Part 2: Investment in Human Beings, 50–79.
- O’Donoghue, T. and M. Rabin (1999). “Doing it now or later”. *American economic review*, 89, 103–124.
- Shen, J. H. and A. Tamkin (2026). “How AI impacts skill formation”. *arXiv preprint arXiv:2601.20245*.
- Snower, D. (1996). “The low-skill, bad-job trap”. *Acquiring skills: Market failures, their symptoms and policy responses*, 111–24.