What holds a job together?

Task bundles, human earnings, and the training of future experts

Sergio Ocampo
ECON 3399 · AI and Robots at Work

A credit analyst does more than calculate repayment risk

A small firm applies for a loan. The analyst forecasts repayment, asks the borrower about its cash flow, and recommends loan terms.

The analyst also builds a relationship with the client to understand their needs and adapt the analysis to their circumstances.

The conversation reveals seasonal sales. That changes the forecast; the forecast raises new questions for the borrower.

> Can AI take over the calculations without losing the context that makes the credit assessment useful?
> The same work also helps a junior learn to judge credit risk.

Weak Bundle, Strong Bundle



Weak Bundle, Strong Bundle:
How AI Redraws Job Boundaries

Luis Garicano* Jin Li† Yanhui Wu‡

March 30, 2026

Abstract

This paper studies how the effect of AI on an occupation depends not just on which tasks AI can perform but also on how costly it is to unbundle those tasks from the job. Much of the discussion of AI and labor markets starts from task exposure: if AI can perform more tasks in an occupation, that occupation should lose employment or earnings. This is incomplete because labor markets price jobs, not tasks. Jobs bundle tasks together, and the effect of AI depends on how costly it is to break the bundle. We build a two-task model in which AI can either assist one task inside a bundled job or supply that task autonomously while a human supplies the residual task. We show that, in weak-bundle occupations, AI automates some tasks and narrows the boundary of the job, activating the standard task-substitution channel once product demand is sufficiently inelastic. In strong-bundle occupations where tasks are not independently reallocable, AI improves performance inside the job, but does not remove the human from the bundle. Thus, bundling provides a force that protects jobs and workers' share of downstream revenue.

1 Introduction

In 2016, Geoffrey Hinton declared that medical schools should stop training radiologists because AI would soon outperform them at reading scans. The remark captured a view that still shapes the discussion: once AI can do a task, the workers who do that task lose.[1]

*London School of Economics.
†University of Hong Kong.
‡University of Hong Kong.

[1]Hinton made the remark at the *Machine Learning and the Market for Intelligence* conference in Toronto, November 2016.

1

How AI redraws job boundaries

Luis Garicano, Jin Li, and Yanhui Wu

Working paper · March 30, 2026

Open the paper

Working paper · First page

What makes credit analysis a strong bundle?

Illustration: calculations and borrower judgment support one loan recommendation.

Source of interdependence	Credit-analysis example
Unpredictable task demands	New facts emerge in the borrower interview; the required analysis changes.
Production spillovers	Seasonal cash flow changes the forecast; the forecast prompts new questions.
Joint outcome measurement	The bank values a sound loan decision, combining numbers with borrower context.

> A bundle is weak when its tasks can be separated with little loss of value.
> A bundle is strong when separating its tasks destroys value.

A credit assessment needs both tasks

$$y = q_1^{\alpha} q_2^{1-\alpha}, \qquad 0 < \alpha < 1$$

Object	Meaning in the credit example
y	Quality-adjusted output: a completed credit assessment.
q_1	Task 1: risk calculations and financial forecasts.
q_2	Task 2: borrower context and judgment about loan terms.
$\alpha, 1-\alpha$	Output elasticities: the contribution of each task.
s_1, s_2	Quantitative and borrower-judgment skills.
$t_1 + t_2 = 1$	Time on tasks 1 and 2 adds up to one workday.

Before task automation: production, wages and entry

Stage	Bundled work (B)
Firm / production	Allocate time $t_1 + t_2 = 1$ across both tasks to maximize returns.
Competitive pay	The wage equals the value of the worker's marginal product.
Worker participation	Enter credit analysis if the best wage offer is at least outside income w_0.

AI-Automation Choice

- AI capability m to perform task 1.
- Two options:
  1. Use AI to assist in task 1 and keep tasks bundled (B).
  2. Automate task 1 and unbundle tasks (U).
- Unbundling destroys fraction c of output.

AI assistance and autonomous risk calculations

Assistance B: $\mu_B(m)$	**Autonomous task 1 U:** $\mu_A(m)$
AI helps with calculations; the analyst checks them.	AI supplies usable calculations independently.
Human time covers both tasks.	Human time goes entirely to task 2.
$q_2 = t_2 s_2$	$q_2 = s_2$
$q_1 = \mu_B(m)^{1-\eta}(t_1 s_1)^{\eta}$	$q_1 = \mu_A(m)$
$y_B = \mu_B(m)^{\alpha(1-\eta)}(t_1 s_1)^{\alpha\eta}(t_2 s_2)^{1-\alpha}$	$y_U = (1-c)\mu_A(m)^{\alpha} s_2^{1-\alpha}$

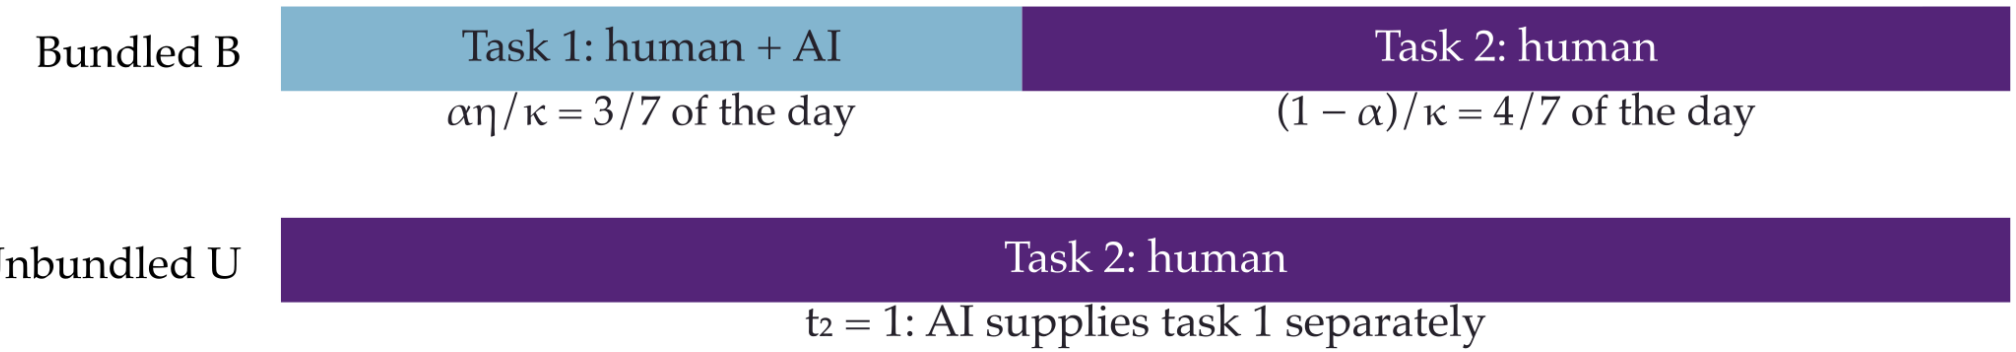
$\eta \in (0,1)$ is the human-input elasticity in task 1.

Unbundling frees time for borrower judgment

Optimal bundled allocation (derivation · production functions)

$$t_1^* = \frac{\alpha\eta}{\kappa}, \qquad t_2^* = \frac{1-\alpha}{\kappa}, \qquad \underbrace{\kappa = \alpha\eta + 1 - \alpha}_{\text{Bundled labor share}}$$

Illustration: $\alpha = ½$, $\eta = ¾$, $\kappa = ⅞$



Task-2 capacity: $\kappa/(1-\alpha) = 7/4 = 1.75$ times the bundled level.
Gain: $\kappa/(1-\alpha) - 1 = \alpha\eta/(1-\alpha) = 3/4 = 75\%$, with skills unchanged.

Real wages depend on revenue shares and output

Real wages w/P are in units of output. For simplicity we set $P = 1$.

$$\frac{w_B}{P} = \kappa y_B, \qquad \frac{w_U}{P} = (1-\alpha) y_U$$

Revenue share $\alpha = \frac{1}{2},\ \eta = \frac{3}{4}$	**Bundled B**	**Unbundled U**
Human input	$\kappa = \alpha\eta + 1 - \alpha = \frac{7}{8} = 87.5\%$	$1 - \alpha = \frac{1}{2} = 50\%$
AI services	$1 - \kappa = \alpha(1-\eta) = \frac{1}{8} = 12.5\%$	$\alpha = \frac{1}{2} = 50\%$

Next, compare κy_B with $(1-\alpha) y_U$: output levels decide between B and U.
Full equilibrium also requires total output Y and its market-clearing price.

Which arrangement supports the higher real wage?

Illustration: $\alpha = \frac{1}{2},\ \eta = \frac{3}{4},\ \kappa = \frac{7}{8},\ y_B = 100,\ P = 1.$

Unbundled output	**Bundled real wage** w_B/P	**Unbundled real wage** w_U/P	**Higher wage offer**
$y_U = 150$	$\kappa y_B = \frac{7}{8}(100) = 87.5$	$(1-\alpha)y_U = \frac{1}{2}(150) = 75$	Bundled B
$y_U = 200$	$\kappa y_B = \frac{7}{8}(100) = 87.5$	$(1-\alpha)y_U = \frac{1}{2}(200) = 100$	Unbundled U

$$\frac{w_U}{P} \geq \frac{w_B}{P} \quad \Longleftrightarrow \quad \frac{y_U}{y_B} \geq \frac{\kappa}{1-\alpha} = 1.75$$

> Unbundling must raise output by at least 75% to offset the smaller human revenue share.

Which arrangement offers the worker a higher wage?

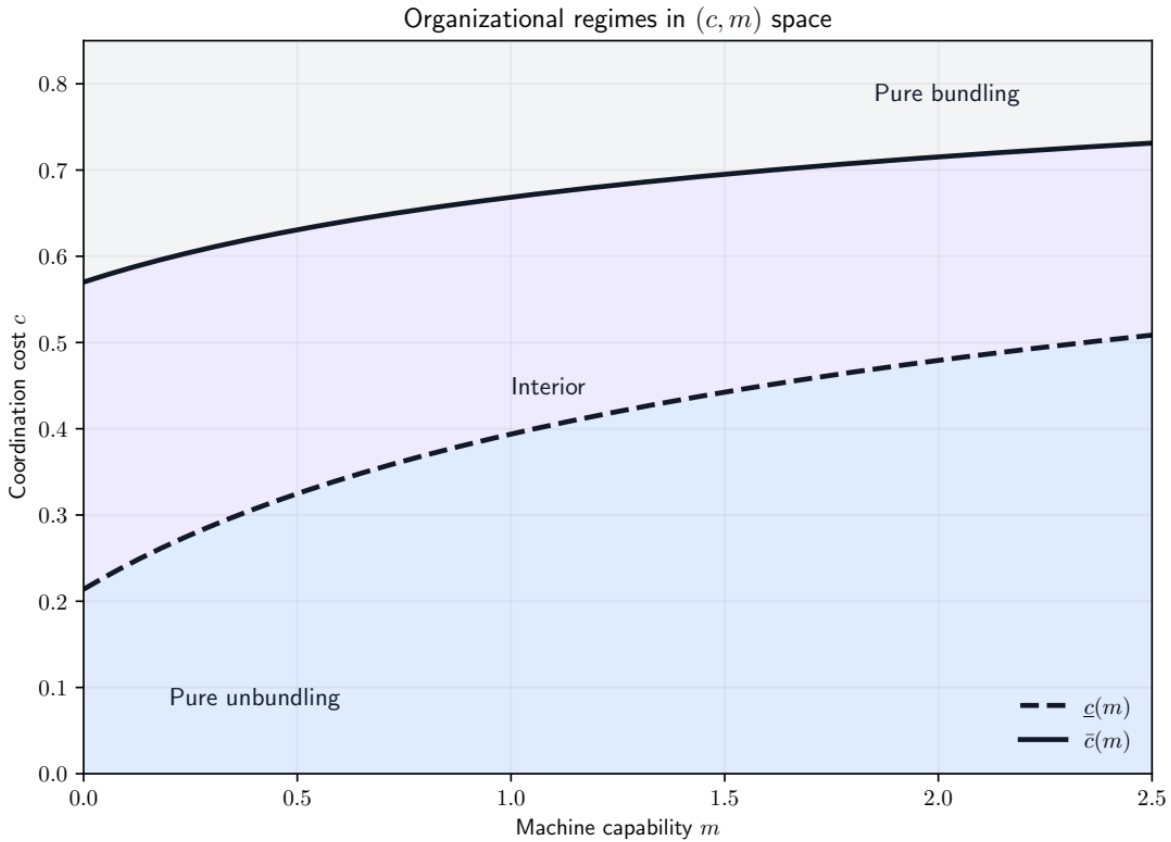
$$\frac{w_U}{w_B} = \frac{(1-\alpha)(1-c)}{\kappa \Omega_B} \Gamma(m)\, s_1^{-\alpha\eta} \quad \longrightarrow \quad \text{Threshold condition for } s_1$$

Ω_B**:** output factor from the optimal time split.
$\Gamma(\mathbf{m}) = \mu_A(\mathbf{m})^{\alpha} \,/\, \mu_B(\mathbf{m})^{\alpha(1-\eta)}$**:** relative autonomous effectiveness.

- Stronger calculation skills make the bundle more valuable.
- Higher c makes separation less attractive.
- Assumption: Γ rises with AI capability m.

The sorting cutoff · P and task-2 skill cancel from the comparison.

When is the bundle worth keeping?



Horizontal m: AI capability.
Vertical c: value lost when separating tasks.

Pure bundling

All workers perform both tasks.

Interior

Some workers bundle and some unbundle;
Depends on s_1!

Pure unbundling

Automate task 1 always.

The worker can also choose another occupation

$$\max\{w_0,\ w_B(P,s,m),\ w_U(P,s_2,m;c)\}$$

- Enter credit-analysis work if its best arrangement pays at least w_0.
- At a given price, w_0/P is outside earnings in units of output.
- Automation depends on s_1; entry into unbundled work depends on s_2!

$$s_2 \geq b_U \equiv \left[\frac{w_0/P}{(1-\alpha)(1-c)\mu_A(m)^\alpha}\right]^{1/(1-\alpha)}$$

Unbundled entry depends on borrower-judgment skill s_2.
Bundled entry still depends on both s_1 and s_2.

Unbundling can open entry into credit analysis

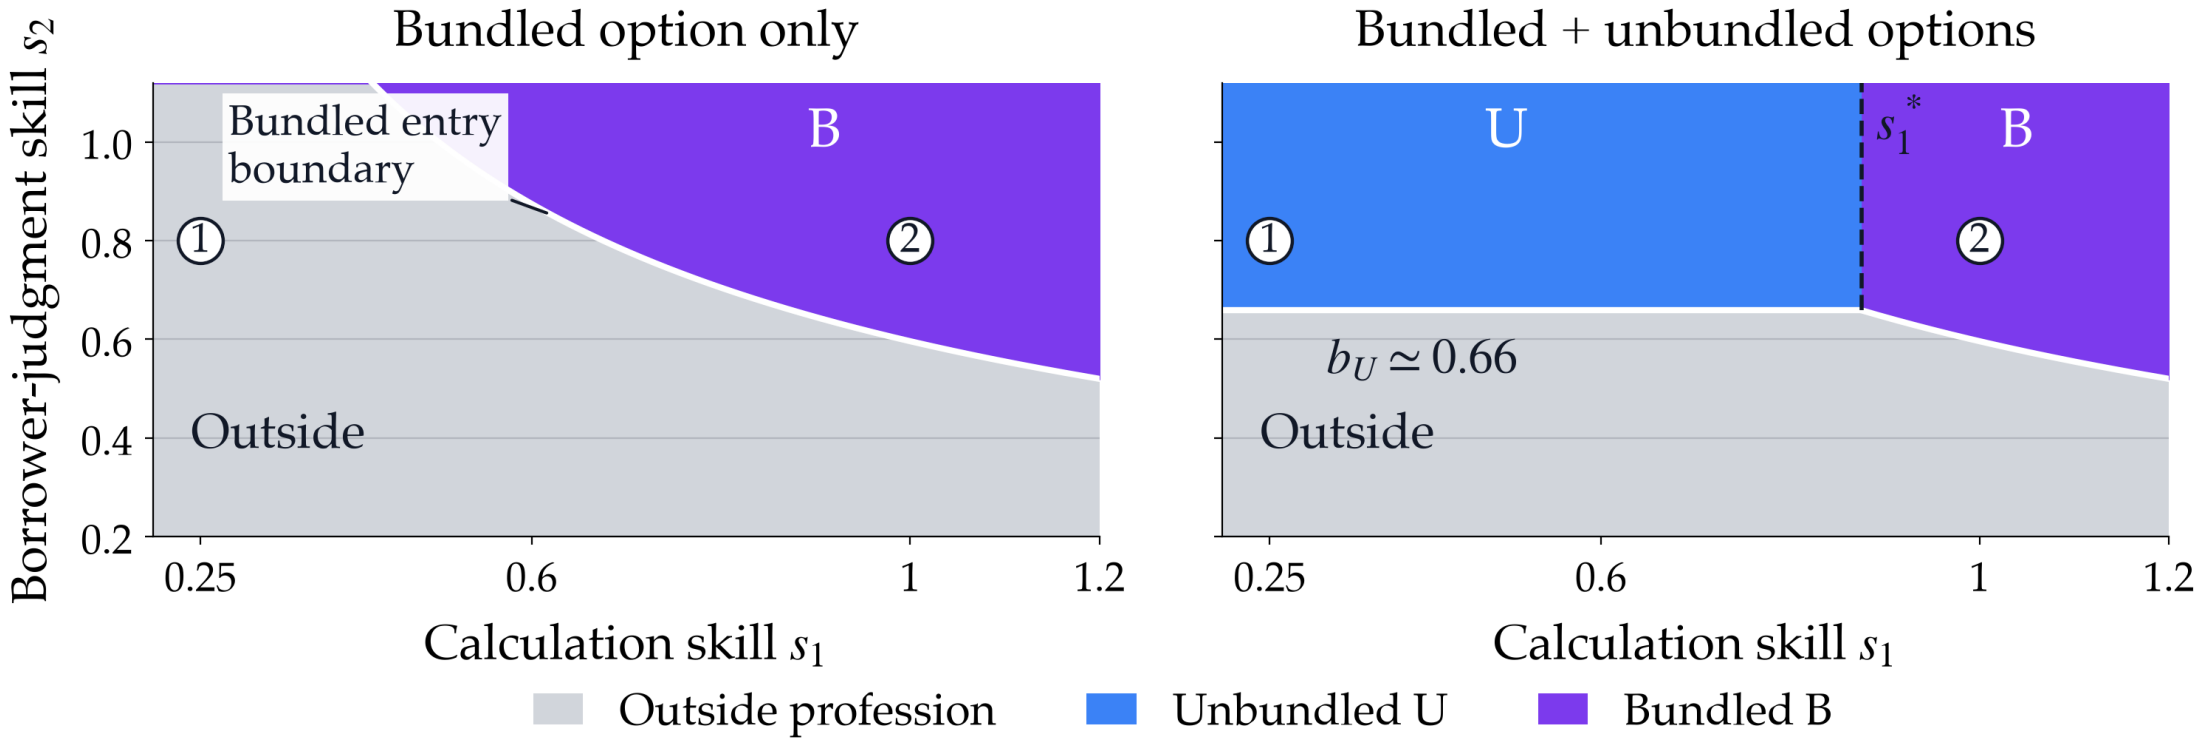


Illustration: $\mathrm{P} = 1$, $\mathrm{w}_0 = 0.4$; $\alpha = \frac{1}{2}$, $\eta = \frac{3}{4}$, $\mathrm{c} = 0.4$, $\mu_{\mathrm{B}} = 1.8$, $\mu_{\mathrm{A}} = 2.7$.

> Worker 1 can enter U even though B pays less than outside work.
> Worker 2 still prefers B at the same borrower-judgment skill.

At a fixed price, better AI expands the residual role

Margin	Why unbundled employment increases
Switching arrangements	The cutoff rises: some B workers switch into U.
Entering the occupation	The U entry threshold falls: more task-2 types qualify.

> To explain displacement, the next step must allow product prices to fall.

- With inelastic demand, output can rise while revenue falls (example).
- Lower prices make entry harder; market clearing determines price and output (figure).

Output can double while employment falls

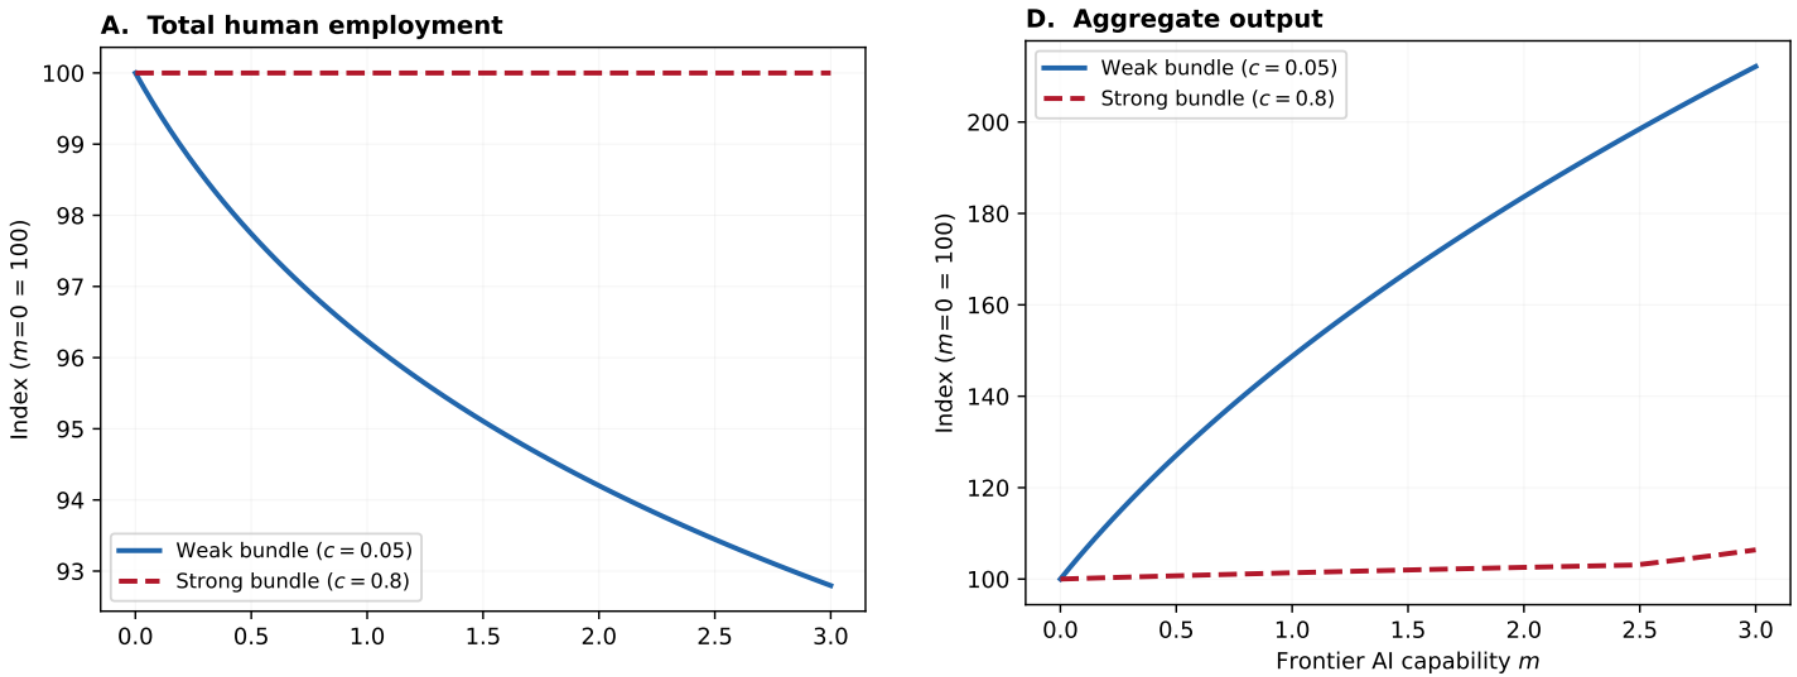


Figure 4, panels A and D: weak bundle (blue) and strong bundle (red); each series starts at 100.

> In this calibration, weak-bundle output more than doubles while employment falls about 7%. The price response offsets the expansion in capacity.

When does the price effect dominate productivity? · Labour income and fixed-worker earnings

Three tiers of the autonomy threshold

Autonomy threshold: AI can deliver finished output well enough to replace paid human work.

Tier 1	→	Tier 2	→	Tier 3
Single verifiable tasks		**Strong task bundles**		**Bundles of relationships**
What sets the threshold?		Why keep tasks together?		Who owns relationships?

- **Tier 1:** coder, junior contract lawyer, wire-service journalist.
- **Tier 2:** technical sales, litigation lawyer, nurse, field reporter.
- **Tier 3:** senior manager, entrepreneur, AI implementation consultant.

Tier 1: Single verifiable tasks

A clean task

- Output and requirements are well defined.
- Quality can be checked.
- AI can meet the required quality at low cost.

Autonomy can still be limited by

- Technology.
- Regulation and accountability.
- Customer preference for a human service.

Examples: a specified program, a standard contract clause, a wire-service summary.

> Once AI can supply the task cheaply, human time no longer limits its supply.
> Competitive pressure can lower the task's price and human earnings.

Tier 2: Strong bundles of tasks

What connects the tasks	Examples beyond credit analysis
Unpredictable needs	Technical sales: client needs emerge during the conversation.
Production spillovers	Litigation: new evidence changes strategy and what to investigate.
Joint quality of the service	Nursing: assessment, explanation and care shape the whole service.

> A task can become automatable before the whole service can be supplied autonomously.
> When splitting the work destroys value, human + AI can preserve the bundle.

Tier 3: Relationships, authority and implementation

What holds work together	Examples in organizations
Decision-making	A manager chooses a reorganization: some teams gain, others lose.
Regulation & accountability	Banking: explain the loan decision and identify who is accountable.
Relationships	Consulting: trust helps client staff share unwritten local knowledge.
Consensus & implementation	AI implementation: secure support from IT, compliance and operations.

> AI implementation requires changing workflows and responsibilities together.
> The scarce work includes securing cooperation and carrying out the decision.

Training in the Age of AI



Training in the Age of AI:
A Theory of Career Viability*

Luis Garicano† Luis Rayo‡

March 3, 2026

Abstract

Across the economy, juniors finance training by performing basic tasks for seniors. AI now automates an increasing share of that work, threatening the exchange. We introduce AI into a dynamic career model with an automation threshold and expert complementarity. The *expertise leverage ratio* R—measuring the output of a fully-trained graduate relative to that of a novice who just outperforms AI—governs the outcome. When R exceeds Euler's e, apprenticeships remain at least as profitable as before AI. Below e, training compresses and advances in AI threaten wholesale career collapse.

1. Introduction

At least since medieval times, economies have relied extensively on apprenticeship-style agreements—broadly defined as work-for-knowledge exchanges—to train new generations of workers. During the initial years of their careers, trainees or "juniors" typically spend much of their time carrying out routine tasks that are valuable to their senior colleagues, but have limited pedagogical value, such as cleaning up spreadsheets and documents in

*We thank Jeff Ely, George Georgiadis, Bengt Holmstrom, Sarit Markovich, Niko Matouschek, Bruno Strulovici, and seminar participants at CEPR, Harvard Kennedy School, and the Eitan Berglas School of Economics at TAU for valuable discussions. We gratefully acknowledge the use of ChatGPT as a drafting aid for early versions of this manuscript. All interpretations, results, and conclusions, as well as any remaining errors are our own. The authors declare no competing financial interests.

†London School of Economics.
‡Kellogg School of Management, Northwestern University.

1

A theory of career viability

Luis Garicano and Luis Rayo

Working paper · March 3, 2026

Open the paper

Working paper · First page

Why would a senior still pay to train a junior?

A senior chooses wages, the knowledge transferred and when the junior graduates.

- **Objective:** maximize discounted junior output minus wages during training.
- **Constraint:** the junior must prefer staying to leaving at every date.
- **Hiring:** the resulting training profit must cover a fixed onboarding cost.

> The senior receives work; the junior receives knowledge.
> AI changes how much useful work can finance that exchange.

Formal contract problem and assumptions

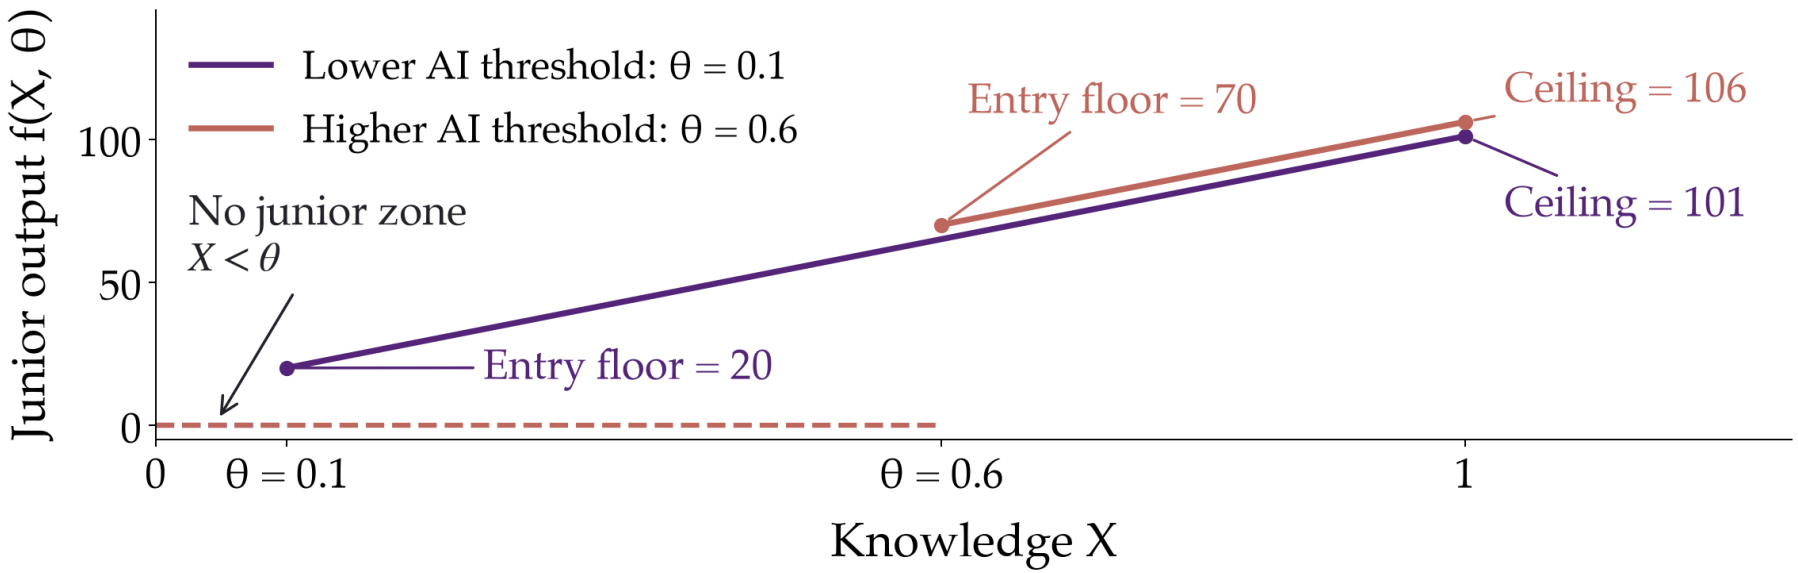
What does the senior choose, and why does the junior stay?

A senior controls a junior's knowledge path $\mathbf{X(t) \in [0, 1]}$.
The junior graduates at time T, when $X(T) = 1$.

Technology	**Meaning**
AI threshold $\theta \in [0, 1)$	Below θ, the junior adds no useful work beyond AI.
Output $f(X, \theta)$ for $X \geq \theta$	More knowledge raises output; better AI weakly raises it.
After leaving	The junior can produce independently at current knowledge, using AI.

> The senior gives enough initial knowledge to clear θ.
> Further training exchanges knowledge for junior work.

AI moves the entry floor and can raise the expert ceiling



Teaching illustration. Dashed segments: no junior zone ($X < \theta$), with no useful junior contribution.

Entry floor (with AI)	**Expert ceiling (with AI)**	**Expertise leverage R**
$f_{\min}(\theta) = f(\theta, \theta)$ Novice who just clears θ.	$f_{\max}(\theta) = f(1, \theta)$ Fully trained graduate.	$R(\theta) = \dfrac{f_{\max}(\theta)}{f_{\min}(\theta)}$ Productive growth from training. 5.05 $\longrightarrow$ 1.51 as θ rises.

The promise of expertise must make the junior willing to stay

Baseline: zero training wages; compensation is future expertise.
r: common discount rate; **T:** graduation date.

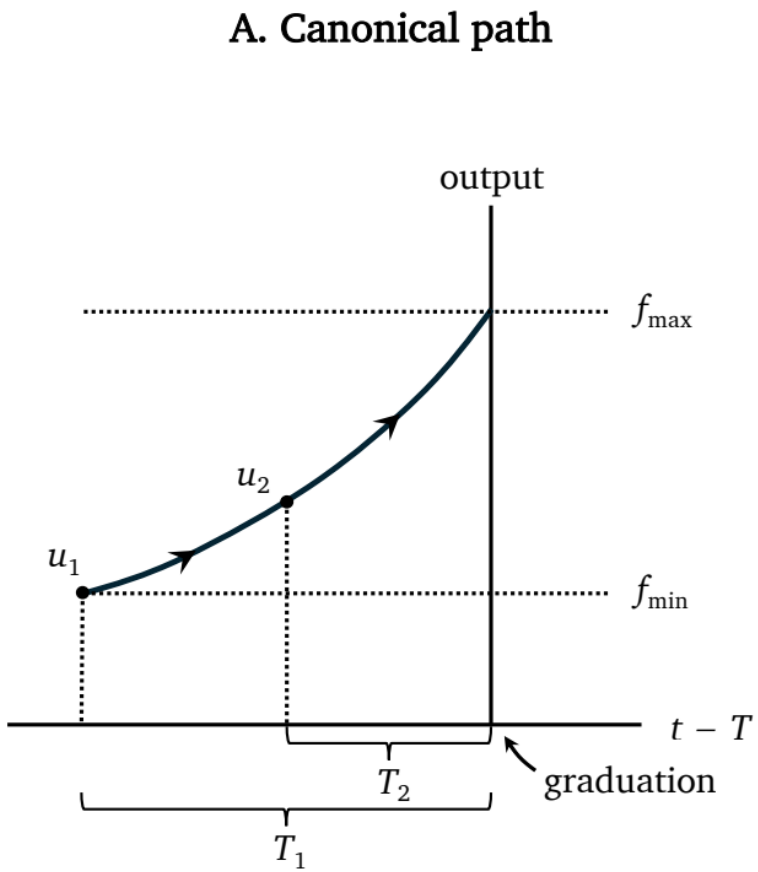
$$\underbrace{\frac{e^{-r(T-t)} f_{\max}}{r}}_{\text{value of staying}} \geq \underbrace{\frac{f(X(t), \theta)}{r}}_{\text{value of leaving now}}$$

- Stay: wait until graduation, then receive expert output forever.
- Leave: produce forever using knowledge already acquired.

> The senior cannot delay training indefinitely.
> The junior must receive enough future expertise to prefer staying.

The no-quit constraint pins down the output path

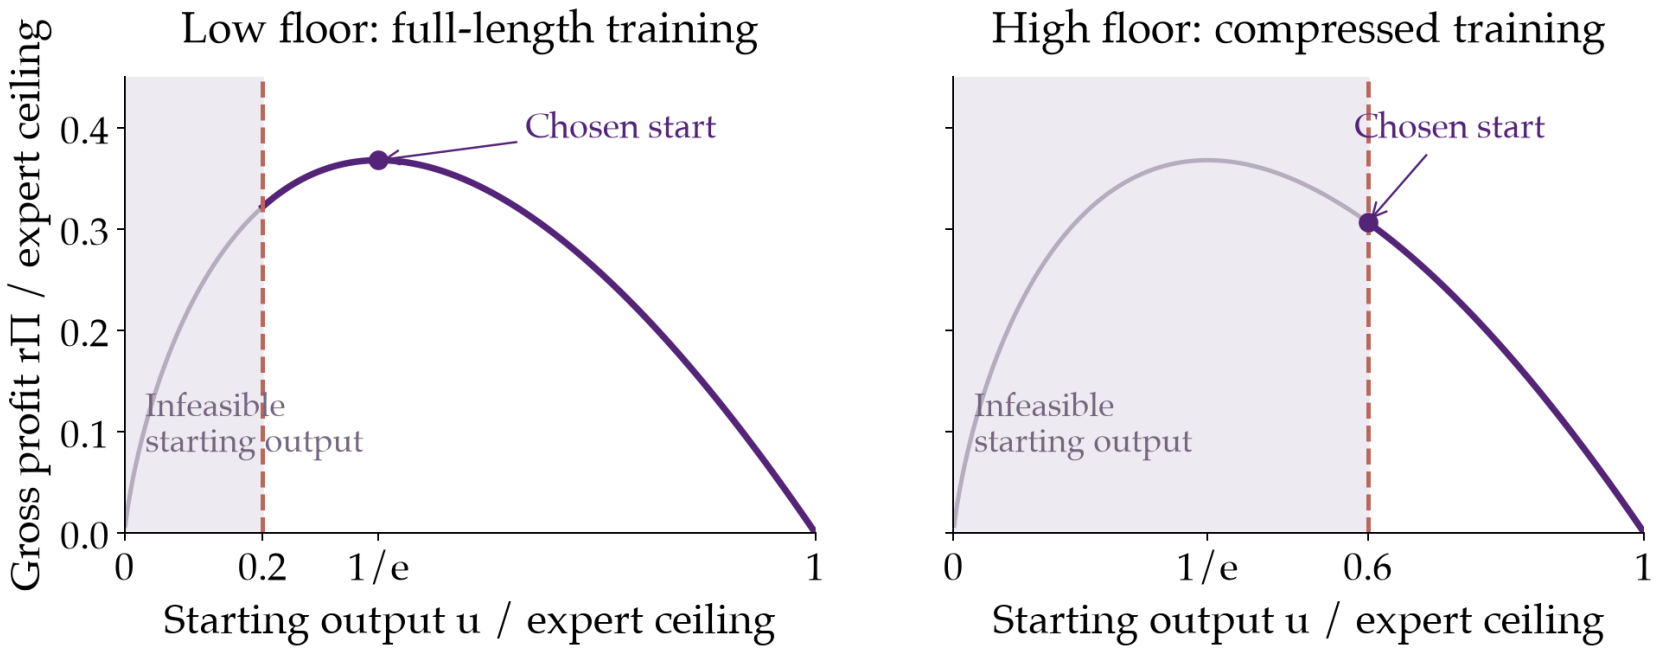
Starting farther up the path shortens the contract



- Horizontal: t - T; zero is graduation.
- Vertical: junior output, up to f_{max}.
- Compare a low start u_1 with a higher start u_2.

> A higher start raises revenue per date.
> It leaves fewer dates before independence.

The AI floor can rule out the profit-maximizing start



u: starting output; Π: training profit; e ≈ 2.718. Gray: starts below the AI floor.

> The ideal start is f_{max} / e; it is feasible when R ≥ e.
> Otherwise, the senior must start at the AI floor.

Profit function and example Optimal contracts and profits Why e?

Two training regimes

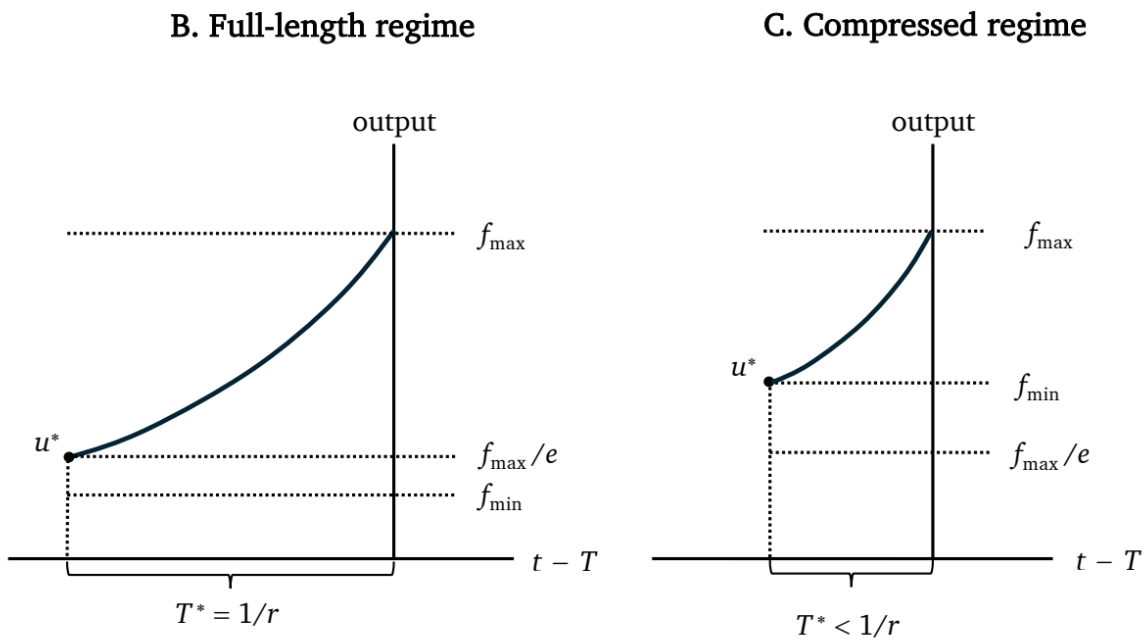
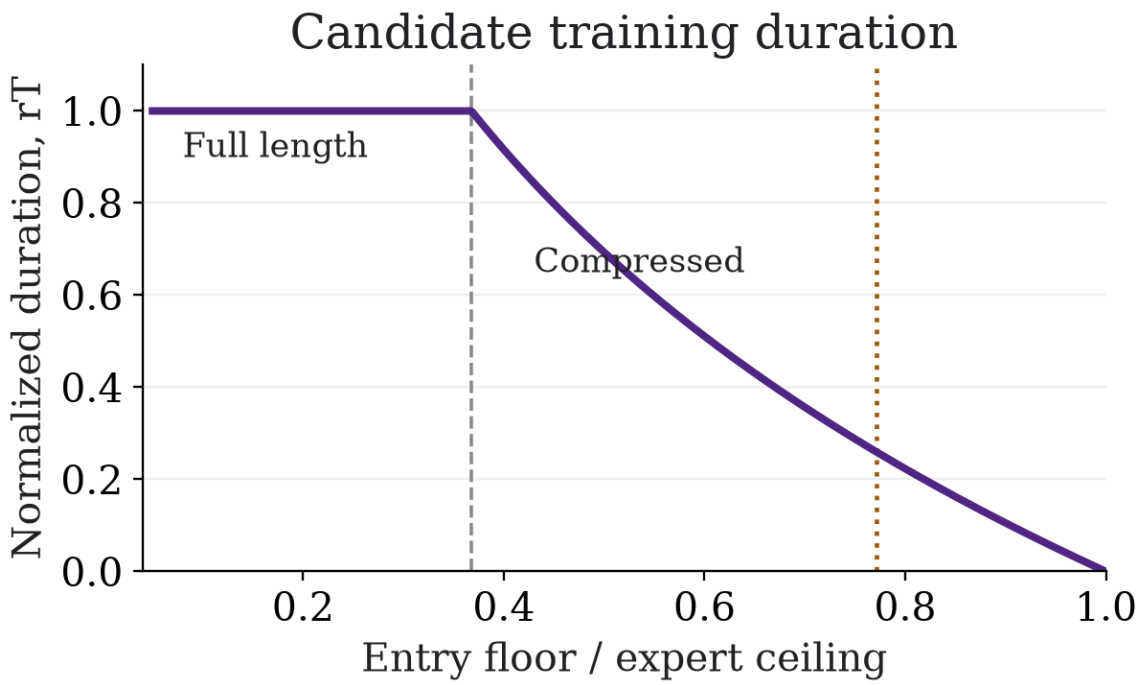


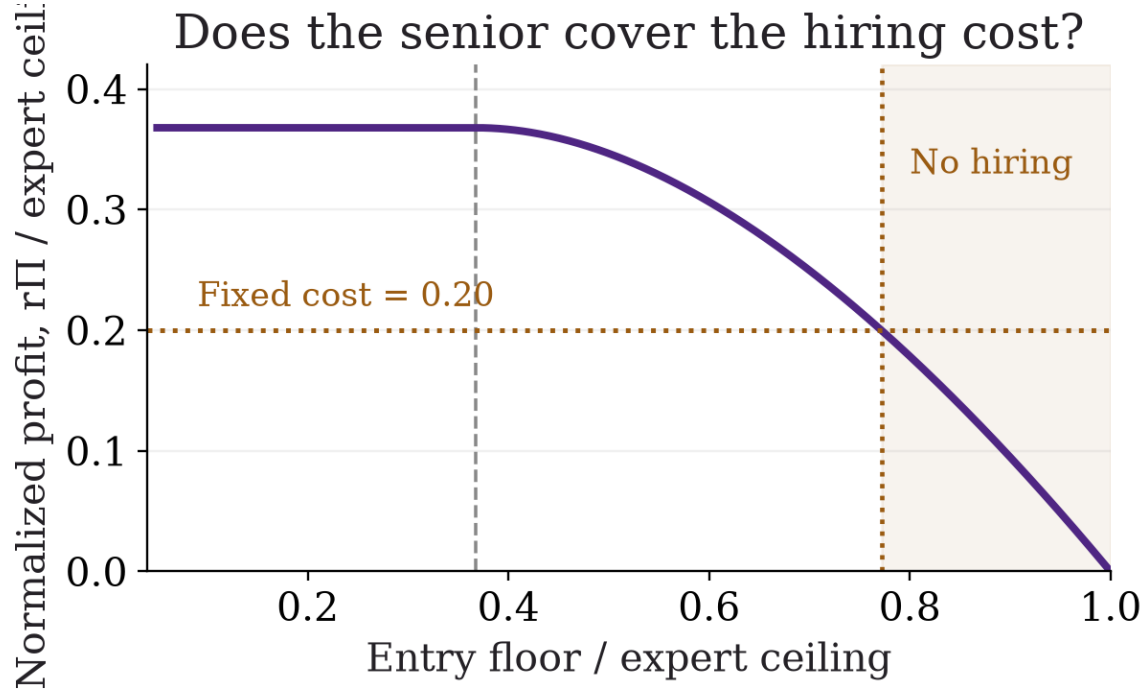
Figure 1B–C: the optimal output path during training, with a full-length or compressed contract.

> A higher entry floor forces juniors to start further up the task ladder.
> Conditional on being hired, they graduate sooner.

Reading the training-path axes

The pipeline can disappear before expert work disappears





Model illustration: fixed expert output. Dotted line: hiring cost F; hire only if $\Pi \geq F$.

> Training first shortens; profits then cross the hiring-cost threshold.
> Experts can remain valuable after the next cohort stops entering.

Profit units and hiring cost Worked hiring examples

What would we expect to observe if the training channel matters?

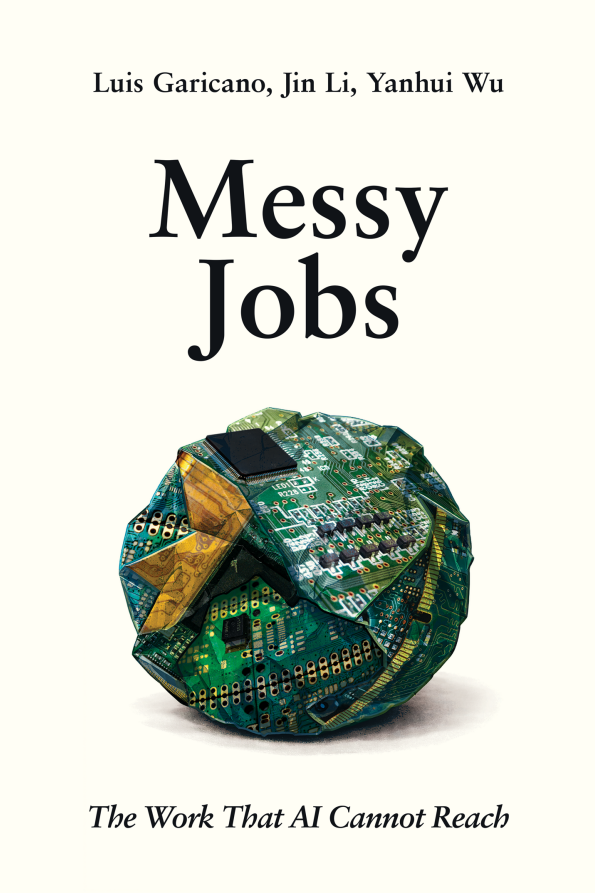
Model change	Observable implication
Entry floor rises	New juniors start on more advanced tasks.
Training compresses	Graduation or promotion is faster among surviving trainees.
Profits fall below fixed costs	Junior entry falls even while senior expertise remains valuable.
Expert complementarity is strong enough	Entry can survive despite loss of routine junior tasks.

> Junior productivity can rise while junior hiring falls.

Job survival and career entry are different economic margins

What we want to understand	Object that matters	Model implication
Does the job retain both tasks?	Coordination loss c vs. autonomous advantage	A strong bundle can preserve the human role.
Do productivity gains raise earnings?	Human revenue share + product demand	More output can coexist with lower earnings.
Will firms train new entrants?	Expertise ratio R + training rent Π^* vs. cost F	Incumbent expertise can remain valuable while entry stops.

Recommended reading: Messy Jobs



The Work that AI Cannot Reach

Luis Garicano, Jin Li & Yanhui Wu

The book connects job design with coordination, trust and responsibility.

Book website and chapters

Interviews with Luis Garicano

- AI and Messy Jobs
  Markus' Academy, episode 164
- Messy Jobs: The Work That AI Cannot Reach
  LSE School of Public Policy

Appendix: Write the two production possibilities

$$y_B^g = \left[\mu_B^{1-\eta}(t_1 s_1)^\eta\right]^\alpha (t_2 s_2)^{1-\alpha}$$

$$t_1 + t_2 = 1, \qquad 0 < \eta < 1$$

$$y_U = (1-c)\mu_A^\alpha s_2^{1-\alpha}$$

Unbundling saves human time on task 1, but pays the coordination cost.

B/U: bundled/unbundled; μ: AI services; s: skills; t: time; η: human share within task 1. Superscript g denotes gross output.

Back to the production setup

Appendix: The worker's share of revenue also changes

$$\kappa = \alpha\eta + 1 - \alpha$$

$$w_B = \kappa P y_B, \qquad w_U = (1-\alpha) P y_U$$

Worker pay: both task contributions when bundled; task 2 when unbundled.

> Higher gross output does not automatically imply higher human earnings.

Appendix: Change capability and the cost of separation

Model illustration at a fixed product price · AI capability m = 1, coordination loss c = 0.2

	Bundled B	Unbundled U
Output	0.600	1.386
Worker earnings	0.525	0.693
Labour revenue share	87.5%	50.0%

Illustrative parameters: $\alpha = 0.5$, $\eta = 0.75$, $s_1 = s_2 = P = 1$, $\mu B = 1 + m$, $\mu A = 1.5 + 1.5m$. Demand, entry and market clearing held fixed. Use the HTML version to vary the inputs.

Appendix: Why the ratio matters for training

As knowledge grows, the junior's outside option improves. Future training must remain valuable enough to retain the junior.

$$T^* = \begin{cases} 1/r, & R \geq e, \\ \log(R)/r, & R < e. \end{cases}$$

Entry floor closer to expert ceiling $\longrightarrow$ shorter optimal training.

Appendix: A shorter training period can make hiring unprofitable

Regime	Training duration	Gross training profit
$R \geq e$	$1/r$	$f_{\max}/(re)$
$1 < R < e$	$\log R/r$	$f_{\min} \log R/r$

Hire a junior only if $\Pi^* \geq F$

R < e means training is compressed.
Career collapse occurs only when the remaining profit falls below the fixed hiring and training cost F.

Appendix: How much expert complementarity preserves training?

In the compressed regime, compare growth in the expert ceiling with growth in the entry floor.

$$\frac{d \log f_{\max}}{d \log f_{\min}} > 1 - \log R$$

At R = 2, the threshold is about 0.31. A 1% increase in the floor requires more than about 0.31% growth in the ceiling for training profits to rise locally.

> A smaller floor–ceiling gap needs stronger expert complementarity to sustain training.

Back to the complementarity result

Optimal time allocation in bundled production

At given skills and AI, maximize output subject to $t_1 + t_2 = 1$.

$$\max_{0<t_1<1} \; \alpha\eta \log t_1 + (1-\alpha)\log(1-t_1)$$

$$\frac{\alpha\eta}{t_1} = \frac{1-\alpha}{1-t_1}$$

$$t_1^* = \frac{\alpha\eta}{\kappa}, \qquad t_2^* = \frac{1-\alpha}{\kappa}, \qquad \kappa \equiv \alpha\eta + 1 - \alpha$$

> Time goes to each human contribution in proportion to its output elasticity.

Back to the main discussion

A task-1 skill cutoff separates the two arrangements

$$s_1^*(m,c) = \left[\frac{(1-\alpha)(1-c)\Gamma(m)}{\kappa\Omega_B}\right]^{1/(\alpha\eta)}$$

Among workers entering the occupation	**Preferred arrangement**
$s_1 < s_1^*$: weaker task-1 skill	AI supplies task 1 autonomously: U
$s_1 > s_1^*$: stronger task-1 skill	Human retains both tasks: B

> Better AI raises the cutoff ⟶ more types prefer U.
> A larger separation loss lowers it ⟶ more types retain B.

Back to the main discussion

Demand determines the price of the extra output

Y: total output of the occupation. **P(Y):** its market price.

$$P(Y) = Y^{-1/\varepsilon}, \qquad \varepsilon < 1 \text{ in the benchmark}$$

ϵ is the price elasticity of demand: a 1% price cut raises demand by less than 1% when $\epsilon < 1$.

> AI raises capacity; selling the extra output requires a lower price.
> If demand responds too little, lower earnings can push workers out.

Back to the main discussion

Choosing an arrangement is different from entering the occupation

Outside earnings are w_0; product price is P*. The ratio w_0/P^* expresses the outside option in units of output.

Figure 2	**How to read it**
Each point (s_1, s_2)	One worker type, with skills in both tasks.
Gray / blue / purple	Outside occupation / unbundled work / bundled work.
b_B / b_U boundaries	Chosen earnings just equal the outside option.
s_1* sorting boundary	Among entrants, which arrangement pays more?

> AI can open the occupation to people with weak task-1 skills.
> Entry still requires enough task-2 skill to beat the outside option.

Back to the main discussion

When does higher productivity push workers out?

For an unbundled worker at fixed skill and c:

$$\frac{d \log w_U}{dm} = \frac{d \log P^*}{dm} + \alpha \frac{d \log \mu_A}{dm}$$

- Direct productivity gain: AI raises the value of the residual human input.
- Market response: extra output lowers the price of the service.
- If the price effect dominates, marginal workers take the outside option.

> In pure unbundling, employment falls if the price effect dominates the productivity gain.

Back to the main discussion

What contract does the senior choose?

Choose wages w(t), knowledge X(t), and graduation time T:

$$\max_{w,X,T} \Pi = \int_0^T e^{-rt}[f(X(t),\theta) - w(t)]dt$$

- Knowledge reaches 1 at graduation; wages cannot be negative.
- At every date, staying must be worth at least as much as leaving.
- Hire only if maximized training profit Π* covers fixed cost F.

> r is the discount rate.
> After graduation, the junior can earn the full expert return independently.

Back to the main discussion

Why does the number e appear?

Choose the starting output that maximizes training profit:

$$\max_{f_{\min} \le u \le f_{\max}} \frac{u}{r} \log \frac{f_{\max}}{u}$$

$$\frac{d\Pi}{du} = \frac{1}{r}\left[\log \frac{f_{\max}}{u} - 1\right] = 0$$

$$u^{\text{ideal}} = \frac{f_{\max}}{e}, \qquad T^{\text{ideal}} = \frac{1}{r}$$

> The ideal start is feasible only if $f_{\min} \le f_{\max}$ / e.
> Equivalently, the expertise leverage ratio must satisfy R ≥ e.

Back to the main discussion

Read the training paths before comparing their lengths

Reading the next figure	Meaning
Horizontal axis: t - T	Time until graduation; zero is the end.
Vertical axis: output	u* at entry; f_{max} at graduation.
Full-length panel	The floor lies below f_{max}/e; duration $= 1/r$.
Compressed panel	The floor lies above f_{max}/e; duration $< 1/r$.

> Compare the floor with the ideal start, holding the expert ceiling fixed.

Back to the main discussion

How much must expert output rise?

In the compressed regime, define the expert-to-entry response as the percentage increase in the ceiling per 1% increase in the floor.

$$\frac{d \log f_{\max}}{d \log f_{\min}} > 1 - \log R \quad \longrightarrow \quad d\Pi^* > 0$$

Current expertise ratio R	Ceiling response needed per 1% floor increase
2.0	More than 0.31%
1.2	More than 0.82%

> A smaller expertise gap requires stronger expert complementarity to preserve training rents.

Derivation and interpretation
Back to the main discussion

Inelastic demand can reduce revenue as output rises

Price per credit assessment: $P(Y) = Y^{-1/\epsilon}$.
$\epsilon = 0.8$: a 1% price cut raises demand by only 0.8%.

Market outcome	Initially	After 10% more output
Output Y	100	110
Price P	100	88.8
Total revenue PY	100	97.6

> Consumers buy more, but the price falls proportionally further.
> More productive work need not generate more income for its suppliers.

Illustration: initial indices = 100; output rises 10%. Demand and employment · Back to the main discussion

Prices and worker participation must agree in equilibrium

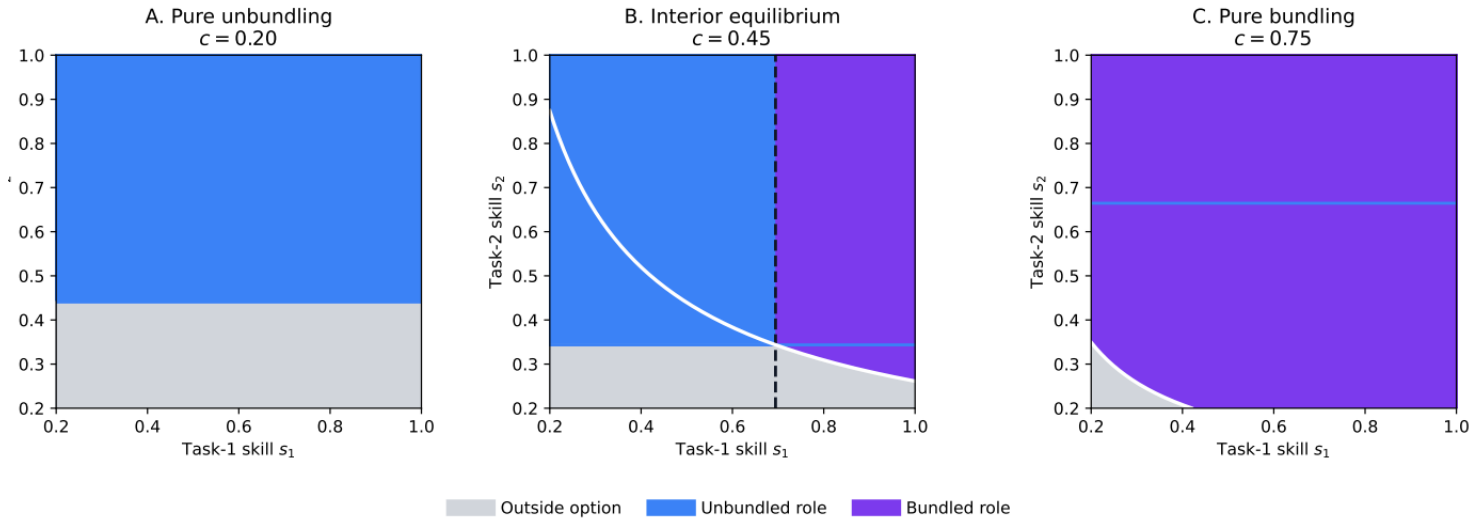
Step	What the model determines
1. Given a candidate price P	Wage offers allocate workers to B, U or outside work.
2. Add their production	Total assessments Y; analysts $N = N_B + N_U$.
3. Apply market demand	The implied price is P(Y) = Y^(-1/ϵ).

$$P^* = P\big(Y(P^*; m, c, w_0)\big)$$

A lower price raises w_0/P and makes entry harder.
The equilibrium balances productive capacity against demand.

Interactive bundle comparison · Back to the main discussion

Unbundling changes who can enter the occupation



Skills s_1 (horizontal), s_2 (vertical); m = 0.8; each panel changes c.
Gray: outside. Blue: U. Purple: B. Vertical s_1*: sorting; b_U / b_B: entry.

> When AI supplies task 1, entry into the residual job depends on task-2 skill.
> Workers previously excluded by weak task-1 skill can qualify.

Reading the entry boundaries · Back to the main discussion

The income loss is not just a change in worker composition

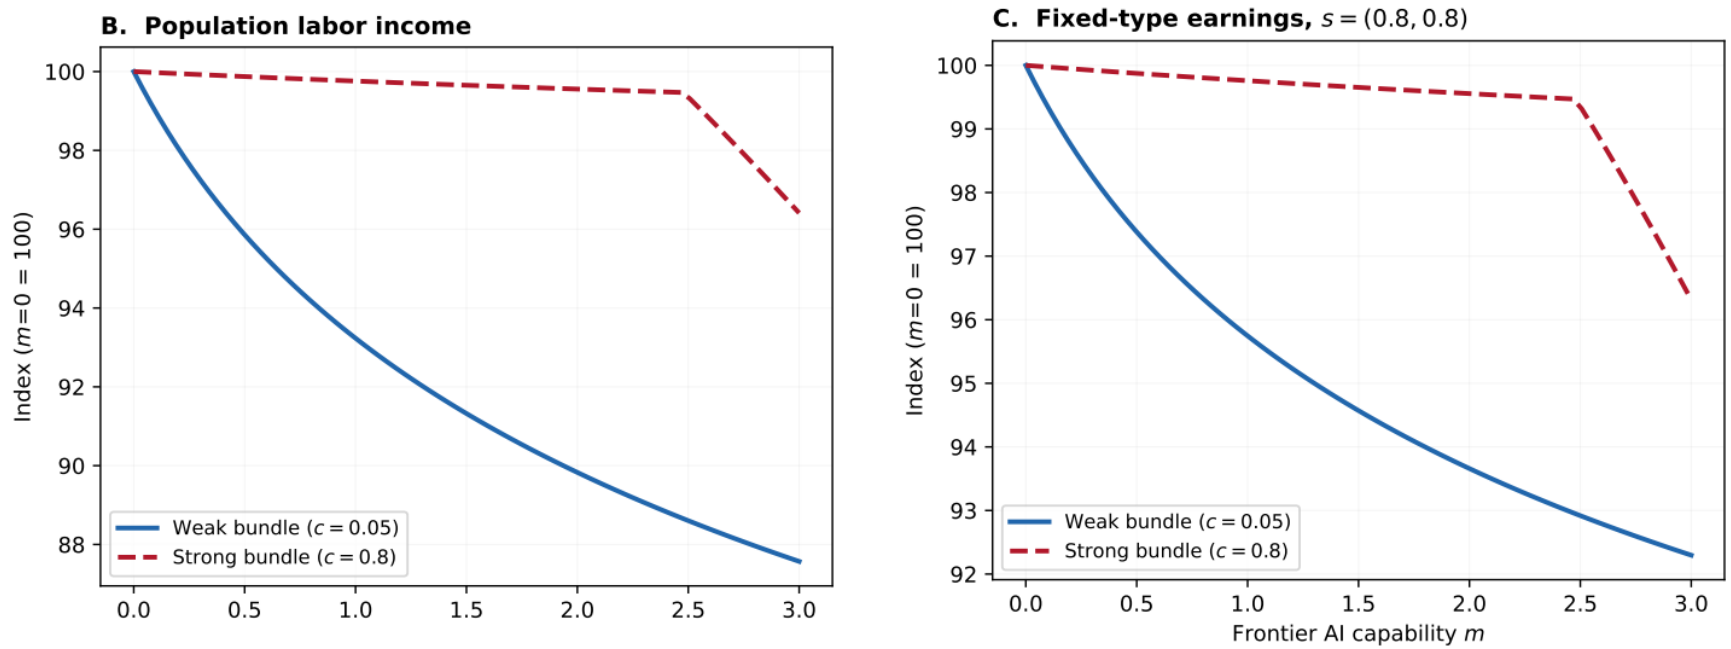


Figure 4, B–C. Labour income from this occupation and earnings of a fixed worker; indices start at 100. · Back to results

Weak-bundle labour income falls about 12%.
Strong bundles delay the drop until capability m ≈ 2.68; keeping the same worker type does not remove the earnings loss.

The no-quit constraint pins down the output path

The senior transfers enough knowledge to make the constraint bind: (Back)

$$f(X(t), \theta) = e^{-r(T-t)} f_{\max}$$

$$f(X(t), \theta) = u e^{rt}, \qquad u \equiv f(X(0), \theta)$$

- Output grows at rate r along the chosen contract.
- Starting output u must be at least the entry floor $f_{\min}$.
- A higher start puts the junior closer to graduation.

> The exponential path follows from incentives.
> The model lets the senior choose how quickly knowledge rises.

The senior trades off a higher output flow against a longer contract

$$T(u) = \frac{1}{r}\log\frac{f_{\max}}{u}$$

$$\Pi(u) = \int_0^{T(u)} e^{-rt} u e^{rt} dt = uT(u) = \frac{u}{r}\log\frac{f_{\max}}{u}$$

Starting output u	Duration T	Gross training profit Π
20	8.05	160.94
36.79	5.00	183.94
60	2.55	153.25

Illustration: expert output = 100; r = 0.2 per time unit; a low enough entry floor. · Back to the profit graph

Two optimal contracts, depending on the floor-to-ceiling gap

Result	**Full length: R $\geq$ e**	**Compressed: 1 < R < e**
Start u*	$f_{\max}/e$	$f_{\min}$
Duration T*	$1/r$	$\log R/r$
Gross profit Π*	$f_{\max}/(re)$	$f_{\min} \log R/r$

For R $\geq$ e, a previously profitable career remains viable.
For R < e, shorter training may not cover the fixed hiring cost.

Back to the profit graph

Hiring compares the value of training with a fixed cost

Π: discounted junior output minus training wages before graduation.
F: fixed hiring and onboarding cost. Hire only if $\Pi \geq \mathbf{F}$.

Next graph / table	**Units**
$x = f_{\min}/f_{\max} = 1/R$	Entry output as a fraction of expert output.
rT	Duration multiplied by the discount rate.
$r\Pi/f_{\max}$; $rF/f_{\max} = 0.20$	Profit and cost in the same normalized units.

> The exercise holds expert output fixed and raises entry output.
> The plotted duration is only a candidate once profit falls below the hiring cost.

Back to career viability

Faster training does not imply more entry

Hold expert output fixed. Express time and profits in normalized units; set the hiring cost to 0.20.

Entry floor / expert ceiling	Candidate duration	Gross profit	Hire?
0.20	1.00	0.368	Yes
0.50	0.69	0.347	Yes
0.80	0.22	0.179	No

> At the high entry floor, an expert remains productive, but the remaining training surplus cannot cover the cost of bringing in a junior.

Back to career viability