

BPEA Fall 2026 Conference
Panel on AI and Policy

The Vanishing Advantage of Specialization

AI, Knowledge Utilization,
and the Boundary of the Firm

Luis Garicano
London School of Economics

September 25, 2026

The vanishing advantage of specialization

The division of labor economizes on the cost of acquiring knowledge.

The reason is that the specialist acquires knowledge once and utilizes it repeatedly, allowing others to rely on her expertise rather than acquiring it by themselves.

Artificial intelligence reduces this fixed cost of acquiring knowledge and therefore also reduces the advantage of specialization.

Specialist's advantage due to utilization

Transaction cost: \$2,000 per case.

Annual cost: \$100,000 before AI; \$10,000 after AI.

A firm faces the problem 10 times a year while the specialist sees it 1000 times.

Cost per case	Before AI	After AI
	\$2,100	\$2,010
Outsourced	$100,000/1,000 + 2,000$	$10,000/1,000 + 2,000$
	\$10,000	\$1,000
In-house	$100,000/10$	$10,000/10$

Even though both sides use the same technology.

Theory

Which force dominates determines whether work moves inside or outside the firm.

Firms and markets

Knowledge is costly to acquire but, once acquired, cheap to reutilize.

Firm

Owns the local knowledge of its core activity.

A firm may undertake few mergers.

Specialist

Sells the same knowledge to many clients, allowing it to be used more intensively.

Cost of crossing the boundary

Translating the firm's language or explaining the local context.

Theory

Order the problem classes the firm faces, indexed by z , from common to rare.

Let:

- $q(z) > 0, q'(z) < 0$ be the number of times the firm faces problem class z
- $n > q(z)$ n be the number of cases of each class across clients dealt with by a specialist
- $k(m) > 0$ be the fixed cost, per period, per problem class, of having reliable knowledge when AI is m
- $\tau > 0$ Per-case cost of asking a specialist for help outside the firm's boundary

We take n to be common across classes.

We assume $k'(m) < 0$, common across classes.

Costs of making and buying

The time spent reviewing each case is a variable cost, assumed equal on both sides.

Producing inside

$$\frac{k(m)}{q(z)}$$

Buying from a competitive specialist

$$\frac{k(m)}{n} + \tau$$

- The firm keeps problem class z inside when

$$\frac{k(m)}{q(z)} \leq \frac{k(m)}{n} + \tau, \quad \text{or equivalently} \quad u(z)k(m) \leq \tau.$$

- Where the specialist's economy of scale due to knowledge utilization is:

$$u(z) \equiv \frac{1}{q(z)} - \frac{1}{n}.$$

The vanishing advantage of specialization

At an interior cutoff $z^*(m)$, the costs of making and buying are equal:

$$u(z^*(m))k(m) = \tau.$$

Proposition 1

A fall in k , holding all else equal, increases the $z^*(m)$ cutoff so that the firm deals inside with a wider range of problems.

$$\frac{dz^*}{dk} = -\frac{u(z^*)}{u'(z^*)k} < 0.$$

The outside specialist deals with a rarer set of problems.

Theory

Corollary 1 (Stigler's theorem)

Hold k and τ fixed. When market expansion allows an outside specialist to pool more cases, n rises.

Work moves out of firms and into specialist suppliers.

Corollary 2 (Minimum internal use)


Firm i makes rather than buys from the outside when

$$q_i(z) \geq \bar{q}(m) \equiv \frac{k(m)n}{k(m) + \tau n}.$$

A fall in k lowers $\bar{q}$ and pushes work inside the firm, while a fall in τ raises $\bar{q}$ and pushes it outside.

Lower knowledge costs and the firm boundary

specialist's cost advantage, $u(z)k$



The outside specialist deals with a rarer set of problems.

Smith versus Coase: what each theory predicts

Observable	Trading gets cheaper	Knowledge gets cheaper
Assignment	Deals inside with fewer problem classes.	Deals inside with more problem classes.
Outside-work mix	Expands to more common problems.	Contracts towards rarer problems.

The relevant question is which cost falls faster in proportional terms.

Holding utilization and other costs fixed.

Professional-service firms can be squeezed in the middle

- Frequent problems move to the firm that owns them.
- Production of general, verifiable knowledge can move up to the model makers.
- The specialist is left with the rare problems for which reliable expertise remains costly to acquire.

Empirics


We are in the very early stages of AI's impact, but we present some preliminary evidence that appears consistent with this mechanism.

Empirical evidence

1. Breadth of AI use
2. Expectations about the firm boundary
3. Freelance market
4. Where specialists work

1. Breadth of AI use

Recent evidence shows workers using AI to cross their occupational boundaries.



More than 800,000 work-related ChatGPT messages from users whose occupations are known.

The first bar divides work messages into generic, within-occupation, and cross-occupation tasks. The second drops generic work and rebases the rest. The source reports no standard errors.

Source: author's transcription of Chin and Richmond (2026).

2. Expectations about the firm boundary

The year-over-year comparisons are restricted to U.S. respondents.

	2024	2025
Reported use of generative AI in legal practice	23%	52%

64%

expect generative AI to reduce reliance on outside providers for routine legal tasks

50%

expect it to reduce the cost of outside providers

These are just expectations from a survey, not purchases.
Association of Corporate Counsel and Everlaw (2025).

3. Freelance market

Freelance writers

−2.0% monthly jobs

−5.2% monthly earnings

After ChatGPT's release, relative to freelancers in less-exposed occupations.

Hui et al. (2024).

Remaining automation-prone posts

+5.7% maximum advertised budgets

+2.2% required skill tags

Among employers posting both before and after ChatGPT.

Demirci et al. (2025).

We do not observe the internal production of the buyer, nor the frequency in each class of problems pre-AI.

4. Where specialists work

99%

of the 337,000 jobs added by the four occupations, were added outside their main specialist industries (335,000).

Function	Growth 2022–25 (%/year)	
	Main specialist industry	Other employers
Legal	1.0	4.0
Accounting	−0.9	1.7
Software developers	−0.7	5.1
Management analysis	1.2	4.4

Other employers are user industries, government, and adjacent specialist suppliers.
Author’s calculations from BLS Occupational Employment and Wage Statistics, May 2022 and May 2025.

4. Where specialists work

Share of each occupation employed in its main specialist industry (percent).



BLS research estimates for 2015–2020 and official estimates for 2021–2025, computed with the same method; the dotted line extends the straight-line trend of 2015–2022.

Software includes quality-assurance analysts and testers. Author’s calculations from the BLS Occupational Employment and Wage Statistics.

Employment growth of each occupation

Function	2015–2019		2019–2022		2022–2025	
	Specialist	Other	Specialist	Other	Specialist	Other
Legal	1.8	2.6	1.3	2.7	1.0	4.0
Accounting	0.2	1.8	3.3	3.3	−0.9	1.7
Software	4.4	6.1	6.9	7.8	−1.1	4.5
Management analysis	4.0	4.8	5.0	4.1	1.2	4.4
Marketing	2.9	7.3	6.8	4.3	1.9	4.2
Design	−1.6	0.9	−5.3	−1.8	0.0	−2.6

Percent per year. Growth rates are compound annual rates between May estimates. “Other” is the number who work for any other employer, including government. Software includes quality-assurance analysts and testers throughout.

Author’s calculations from the BLS Occupational Employment and Wage Statistics.

4. A smaller share and a larger volume can go together

Specialists can lose share and nevertheless gain volume if the number of problems solved expands sufficiently.

In the first eleven months of 2025, in a sample of 184 U.S. law firms:

1.9%

Billable hours grew

7.3%


Worked hourly rates

Thomson Reuters Institute and Georgetown Law (2026).

4. Payroll employment in four specialist industries

Employment in professional, scientific, and technical services firms grew at 0.5% per year from November 2022 to July 2026.

Employment, November 2022 = 100



Monthly, seasonally adjusted, January 2018 to July 2026, with November 2022 set to 100.
The slowdown partly reflects the end of the pandemic hiring boom. Source: BLS Current Employment Statistics, downloaded from FRED.

Policy

Licensing, competition, and training.

Licensing

- Where the solution is equally reliable, and there is not a specific reason to mandate independence (think of external auditors):
 - Rules that constrain the ability of internal professionals to sign their own solutions... can reduce gains and dynamism.

Competition

If the owners of the models, or of the data that train them, have monopoly power, they can charge high fixed access fees or restrict access to their models.

This could keep internal provision expensive and weaken the effect that we are discussing.

Training

Professional service firms run an apprenticeship that produces deep expertise for the whole industry.

Juniors pay for their training with cheap labor on routine cases, and the firm recovers the cost by billing clients for that work.

Routine work supplies the problems that are actually needed for juniors to acquire experience.

Training

AI can affect both: it can remove the routine work that financed the arrangement, and with it the material on which juniors learned by doing.

New institutions may have to emerge to provide a training ladder that routine work used to provide.

The vanishing advantage of specialization

So firms do a larger share of their support work themselves and rely on outsiders for a smaller share.

The specialist's comparative advantage lies in expertise that clients use too rarely to maintain internally and that remains costly to acquire even with AI, including the expertise required to check and verify a solution.

Cheaper knowledge, as we have seen, brings problem solving inside the firm while lower transaction costs lead to buying more solutions outside.

Selected specialist functions

Function	Employment 2025 (thousands)	In main specialist industry (%)		Receipts (\$ billions)
		2022	2025	
Legal	754.5	61.4	59.4	374.6
Accounting	1,449.5	24.2	22.8	208.8
Software developers	1,687.9	34.3	30.5	600.6
Marketing	899.6	4.9	4.6	134.8
Management analysis	898.3	26.6	24.9	435.0
Design	197.8	11.1	11.9	36.8
Six functions	5,887.6	–	–	1,790.6

Receipts cover all establishments in the listed industry, not only the row occupation.

Author's calculations from BLS Occupational Employment and Wage Statistics, May 2022 and May 2025, and the U.S. Census Bureau's 2022 Economic Census, table EC2200BASIC.

Share of each occupation employed in its main specialist industry

Function	2015	2019	2022	2025	Trend 2015–22 pp/year	Change 2022–25 pp	Gap 2025
	percent						
Legal	63.1	62.4	61.4	59.4	−0.23	−2.0	−1.4
Accounting	25.5	24.2	24.2	22.8	−0.08	−1.4	−1.4
Software	36.2	34.7	34.2	30.5	−0.44	−3.6	−1.7
Management analysis	26.7	26.1	26.6	24.9	−0.27	−1.7	+0.1
Marketing	5.3	4.5	4.9	4.6	−0.02	−0.3	−0.2
Design	13.3	12.2	11.1	11.9	−0.40	+0.8	+2.0

“Trend” is the slope of a straight line fitted to 2015–2022. “Change” is 2025 minus 2022. “Gap” is the 2025 value minus the trend line extended to 2025.

Software adds software developers and software quality-assurance analysts and testers. The estimates are subject to sampling error; standard errors for the shares are not reported.

Author’s calculations from the BLS Occupational Employment and Wage Statistics.

Employment growth of each occupation

Function	2015–2019		2019–2022		2022–2025	
	Specialist	Other	Specialist	Other	Specialist	Other
Legal	1.8	2.6	1.3	2.7	1.0	4.0
Accounting	0.2	1.8	3.3	3.3	−0.9	1.7
Software	4.4	6.1	6.9	7.8	−1.1	4.5
Management analysis	4.0	4.8	5.0	4.1	1.2	4.4
Marketing	2.9	7.3	6.8	4.3	1.9	4.2
Design	−1.6	0.9	−5.3	−1.8	0.0	−2.6

Percent per year. Growth rates are compound annual rates between May estimates. “Other” is the number who work for any other employer, including government. Software includes quality-assurance analysts and testers throughout.
 Author’s calculations from the BLS Occupational Employment and Wage Statistics.

4. Empirical evidence and measurement

- The survey observes neither purchases nor problem assignment and excludes the self-employed.
- Each May estimate pools six semiannual panels over three years, so post-2022 changes enter gradually.
- We cannot claim to identify causality here.
- This timing fits both the general drop in knowledge costs with the IT revolution, intensified by AI, and the end of the professional-services boom.

Who purchases professional and business services

User sector	Purchases (\$ billions)	Share (percent)
Professional and business services	581	18.5
Finance, insurance, and real estate	532	17.0
Wholesale and retail trade	510	16.3
Education, health, and social assistance	282	9.0
Government	273	8.7
Manufacturing	270	8.6
Information	225	7.2
Other user sectors	461	14.7
All user sectors	3,134	100.0

Intermediate uses of the commodity “Professional and business services” at producers’ prices, rounded to the nearest billion. Gross purchases, not value added; the commodity is broader than the six functions.

Source: Author’s grouping of the BEA 2017 use table before redefinitions, sector level, 2023 revision.

Firm keeps problem class z inside when

$$\underbrace{\frac{k(m)}{q(z)}}_{\text{producing inside}} \leq \underbrace{\frac{k(m)}{n} + \tau}_{\text{buying from a competitive specialist}}$$

Or equivalently $u(z)k(m) \leq \tau$, $u(z) \equiv \frac{1}{q(z)} - \frac{1}{n}$.

$q(z)$ the number of times the firm faces problem class z

n cases of each class across clients, with $n > q(z)$

$k(m)$ the fixed cost, per period and per problem class, of having reliable knowledge when AI is at level m

τ the per-case cost of asking a specialist for help outside the firm's boundary

What is relevant for the analysis is the part of capability cost that remains fixed, and is divided by q or n .

Minimum internal use

Firm i makes rather than buys from the outside when

$$q_i(z) \geq \bar{q}(m) \equiv \frac{k(m)n}{k(m) + \tau n}.$$

A fall in k lowers $\bar{q}$ and pushes work inside the firm, while a fall in τ raises $\bar{q}$ and pushes it outside.

When both costs fall, z^* rises and work moves inside if, at $u(z^*)k = \tau$,

$$-\frac{k_m}{k} > -\frac{\tau_m}{\tau}.$$

Holding q and n fixed.

Problem-specific knowledge costs

Hold q , n , and τ fixed. Allow the fixed cost of reliable knowledge to depend on the problem class, $k = k(z, m)$.

Suppose first that AI weakly lowers the knowledge cost for every problem class, $k_m(z, m) \leq 0$ for all z .

Any class made inside before the improvement remains inside afterward.

Suppose, in addition, that the specialist's cost advantage rises as problems become rarer:

$$A_z(z, m) = u'(z)k(z, m) + u(z)k_z(z, m) > 0.$$

The firm's choices then continue to have a single common-to-rare cutoff.

A smaller share and a larger volume can go together

Specialists can therefore lose share and nevertheless gain volume if the number of problems solved expands sufficiently.

$$\Delta \log X_z = \Delta \log s_z + \Delta \log Q_z$$

X_z specialist volume

s_z the share handled by outside specialists

Q_z the number of problems of class z solved in the economy