Discussion of *The Market's Expectation of AI-Driven Growth*

Boppart, Klenow, Laski and Li

Carla Varona Cervantes

Central Bank of Chile

Labor Markets and the Macroeconomy in a Changing World
November 20, 2026

The question, and a good idea

Q: How much will AI raise productivity growth?

- Published forecasts span an order of magnitude — from under 0.1 pp, through a median around 1.2 pp, to unbounded among the most optimistic.

This paper: do not add a forecast — read the market's. Investors have skin in the game, and P/E prices expected rents.

The answer: AI firms' expected contribution to growth rises 43% $\rightarrow$ 79%, against 33% of sales and 47% of earnings — about 133 bp of TFP growth.

And it comes from bigger steps, not more creative destruction: growth per arrival 0.18 $\rightarrow$ 0.36, with the arrival rate steady at 6% and measured job reallocation steady near 11.3%.

Where the identification comes from

The model. Klette–Kortum quality ladders, with two deviations: idiosyncratic shocks to γ and φ, and an arrival rate that does not scale with firm size so $x = x(\gamma_{jt}, \varphi_{jt})$.

Step 1. Calibrate the model. Two perios: 1976-2019 and 2020-2025

Step 2. Then, firm by firm and year by year:

$$R_{jt} = \frac{x(\gamma_{jt}, \varphi_{jt})^{\nu}}{\varphi_{jt}}$$

$$V_{jt} - \text{PV of existing products} = \text{value of innovation}$$

Two equations, 2 unknown: $(\gamma_{jt}, \varphi_{jt})$.

Step 3. Compute (model implied) expected contribution to growth.

- **The engine.** Size is the accumulation of past realizations but innovation depends only on the current draw — contribution $\neq$ sales share.

This discussion

Then three questions

1. How much weight should the market's forecast carry?
2. A balanced growth path for a diffusion?
3. Does the model reproduce the flat reallocation?

One thought about Chile: whether σ_γ^2 is a technological parameter, or also an institutional one.

And one comment on measurement: where the sample runs out, the forward moments lean on forecasts.

1. How much weight should the market's forecast carry?

The paper's move is to read the forecast off prices rather than add one more.

- Great response for a literature whose estimates span an order of magnitude.

But then 79% **is only as good as the forecaster**.

- A natural question: how good has that forecaster been?

Slide 30 begins to answer this: Expected contribution of the top 7 since 1975:

- Today's reading is the highest in fifty years (about 75%).
- The previous peak was March 2000 (near 67%) — the last technology boom of this shape: firms priced for transformative growth.

What happened *next*? Model predictions vs Realization.

Contribution to aggregate growth is not observable but earnings and margin growth are.

1. Use inferred $(\gamma_{jt}, \varphi_{jt})$ for the top contributors of 1999–2000.
2. Iterate the chain forward from model equations.
3. Compare the implied g_E and $g_{E/S}$ with the realizations.
4. Report the forecast error.

You have already computed those γ_{jt} — it is what slide 28 is built from, and the realizations are in Compustat.

2. A balanced growth path for a diffusion?

Both periods are solved as *stationary* equilibria, yet 2020–25 is six years of ongoing diffusion

- recalibration moves a great deal. VAR(1) dispersion (σ^2_γ: $0.92 \rightarrow 4.90$).

Which equilibrium this is — the one we are in today, or the one we are heading towards?

Combining the contribution shares with the measured growth rates of each period:

Contribution to growth	before	2020–25
AI firms	$0.43 \times 1.2\% \approx 0.5$ pp	$0.79 \times 2.2\% \approx 1.7$ pp
Everyone else	$0.57 \times 1.2\% \approx 0.7$ pp	$0.21 \times 2.2\% \approx 0.5$ pp

The rest of the economy contributes fewer percentage points of growth than before — about 0.5 against 0.7.

Rents, or growth?

Did they get worse at innovating just as AI arrived? Or does the market expect no diffusion at all? **Probably neither.**

- Diffusion takes time and spillovers come later.
- Over a 2020–25 valuation horizon the market is pricing the direct rents of those *building* the technology.

79% might be where the rents land, rather than where the growth is generated — and in this framework those are one object by construction.

- A firm contributes to growth by improving its own lines, and captures the rent from doing so.
- A general purpose technology raises the step size of *other* firms.
- Technology producer prices that gain, the model reads the rent as the producer's own innovation.

A result in the calibration that no slide states

The calibration says something more: part of it comes from labour.

	1976–2019	2020–2025
Average markup	1.13	1.21
Aggregate earnings / sales	0.11	0.20

With labour the only variable input, the cost-weighted markup is the inverse of the labour share in revenue:

$$1/1.13 \approx 88\% \qquad \longrightarrow \qquad 1/1.21 \approx 83\%$$

So AI boom delivers together: faster growth, wider margin dispersion, more market power, and some five points off the labour share. — the paper has it quantified without stating it!

3. Does the model reproduce the flat reallocation?

Growth doubles while job reallocation stays near 11.3% .

- probably one of the most reassuring messages.

But job reallocation is not one of the ten targeted moments, and the model has something to say about it.

In equilibrium, with labour inelastically supplied,

$$\ell_i \;=\; L \cdot \frac{1/\gamma_i}{\int (1/\gamma_j)\, dj}$$

so employment is a *share*: the level of γ washes out, and only its dispersion moves employment across lines — and that is five times larger in 2020–25.

So does the calibrated model reproduce a flat 11.3%, or imply that reallocation should have risen? The one labour market moment here, and untargeted — a natural overidentifying check.

Is σ_γ^2 a technological parameter? The case of Chile

One of the parameters that moves most is the *dispersion* of step sizes: σ_γ^2 rises from 0.92 to 4.90.

Petit (2023) says that dispersion is itself a choice: firing costs make innovation riskier, so firms pick less dispersed, lower-expected-return innovation — conservatively, not merely less.

From the paper we presented yesterday, three things about that cost in Chile:

- **it is large:** 86% of contracts define the worker's functions rigidly
- **it varies:** 7% adaptable in utilities, 21% in business services
- **and it moves with technology:** firms investing more in automation write more flexible contracts, with no such relationship for non-technological investment

What Chile can offer

Not a replication. Chile cannot run your exercise — there are no AI producers here to invert, and no valuations from which to recover γ.

What it can speak to is whether σ_γ^2 is a primitive.

If the dispersion of innovation responds to the cost of reassigning workers, then part of that parameter is the environment firms operate in, rather than the technology available to them.

The natural extension: let the parameters of the γ process depend on measured institutions

- how much of the cross-country difference in AI's contribution is technological

A last comment: One reason the recent reading may not be comparable

Moments 3, 4, 6 and 7 use growth from t to $t+5$. Where $t+5$ runs past the sample, E_{t+5} is the I/B/E/S consensus forecast.

But then both sides of the regression are expectations drawn from the same information set: dependent and independent variables are analyst forecasts.

Moments are not quite the same object in the two columns:

- 1976–2019: link valuations to realized outcomes.
- 2020–25: link valuations partly to expectations of them.

And the paper rests on comparing the two periods.

Measuring the wedge on 1997–2000

Both versions exist for the same firms on the same dates.

- Run regressions with both dependent variables:once *realized* and once *forecast*.
- The gap between them is the wedge.

The answer is not obvious.

- Uniform over-optimism shifts the level, not the slope.
- The moment moves only if the optimism was *correlated with valuation*.

Why this window? Technology boom again.

- Evidence suggests the market was broadly right about the aggregate — productivity growth did accelerate after 1995 — and much less right about *which firms* would deliver it.

Summing up

A clever idea, carefully executed. When a literature's forecasts span an order of magnitude, asking the people with money at stake is the right response.

Four things I would enjoy seeing

1. A way to complete the out-of-sample test that slide 28 begins
2. A note on what the forecast substitution costs, measurable on 1997–2000
3. The rents and the growth separated — and the labour share result drawn out
4. The flat reallocation checked against the model, since it is untargeted

And one thought. σ_γ^2 does a great deal of work here. Whether it is technology or also institutions is a question Chilean contract data can speak to.

Thank you.