

Transformative AI, existential risk, and real interest rates

Trevor Chow¹ Basil Halperin² J. Zachary Mazlish³

¹Stanford, ²University of Virginia, ³Oxford

December 2025

AI researchers expect economic growth to accelerate

AGI could create about \$100 trillion in new wealth – roughly doubling the world's GDP – that actually sounds quite plausible to me.

— Sam Altman, CEO of OpenAI



AI researchers expect economic growth to accelerate

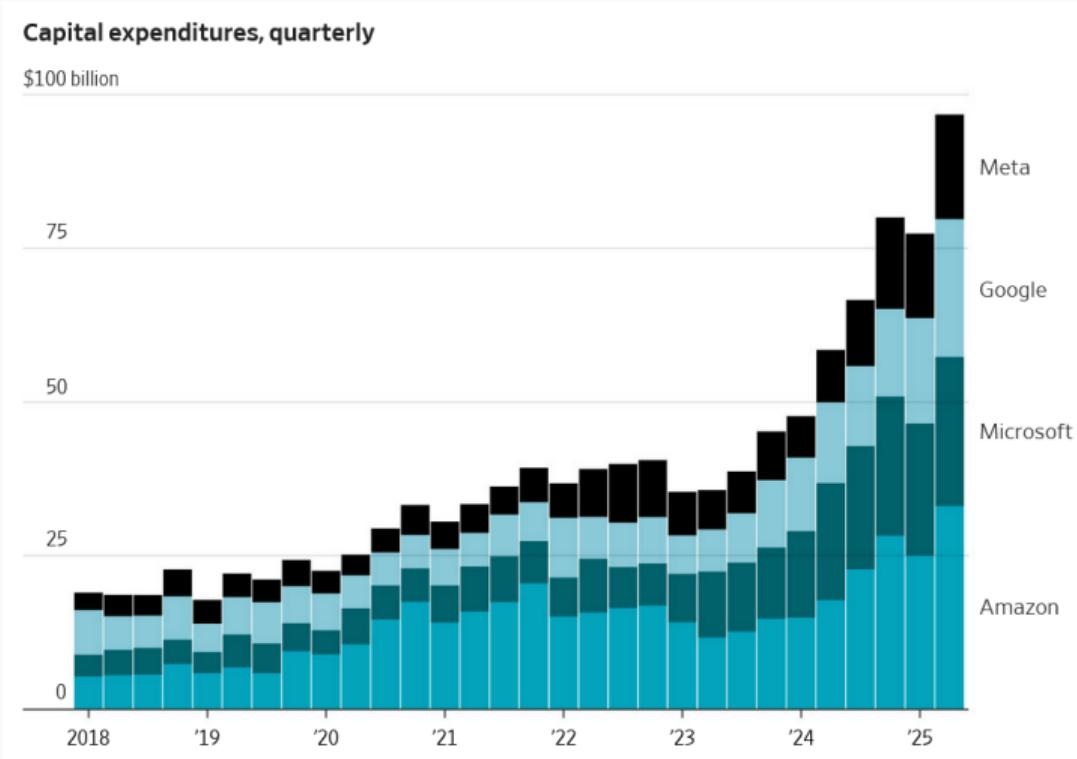
We are likely to see **periods of 4 years in the future that experience more change than the last 200.**

— Paul Christiano, head of AISI & inventor of RLHF



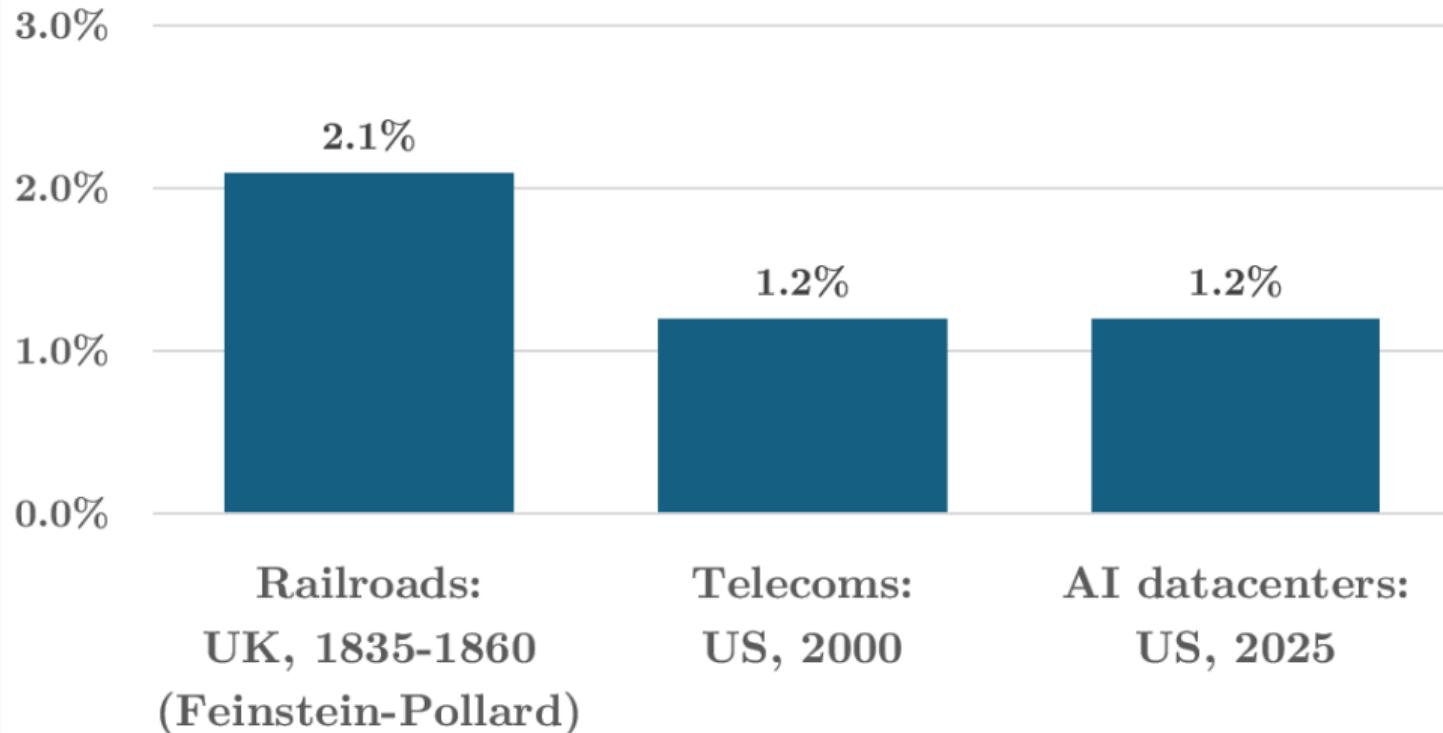
Investment is booming

Hyperscaler capex is 1% of GDP:



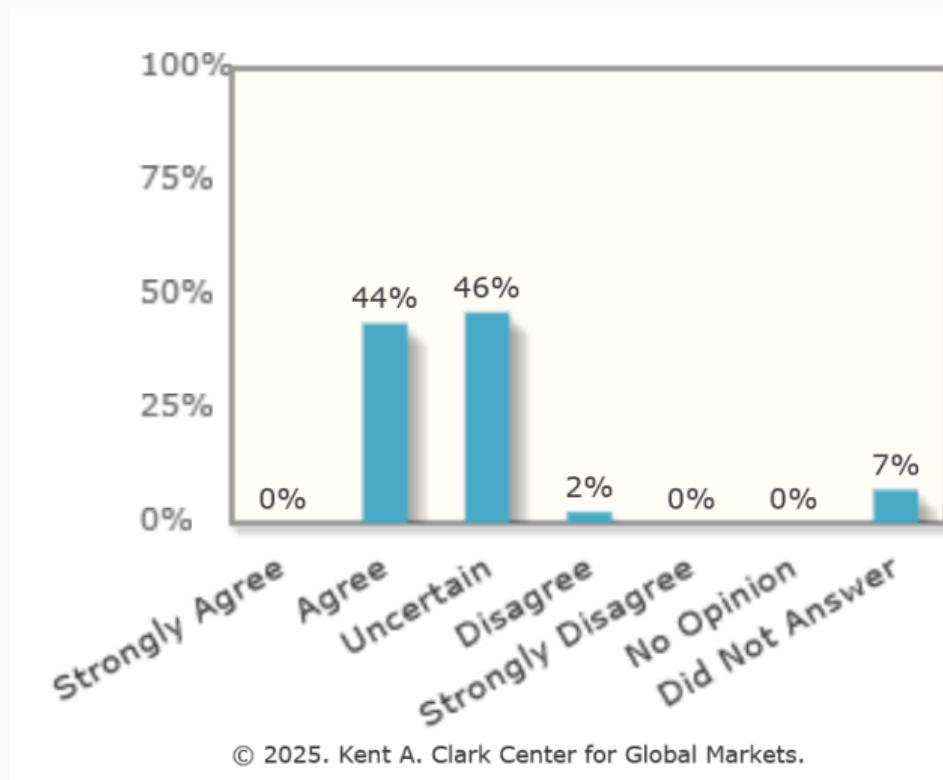
Investment is booming

Capex as % of GDP, by episode



Economists are, on average, measured about AI impacts

AI will cause a “substantial increase” in GDP per capita growth over next 20 years:



Econ prior: stable 2% GDP growth

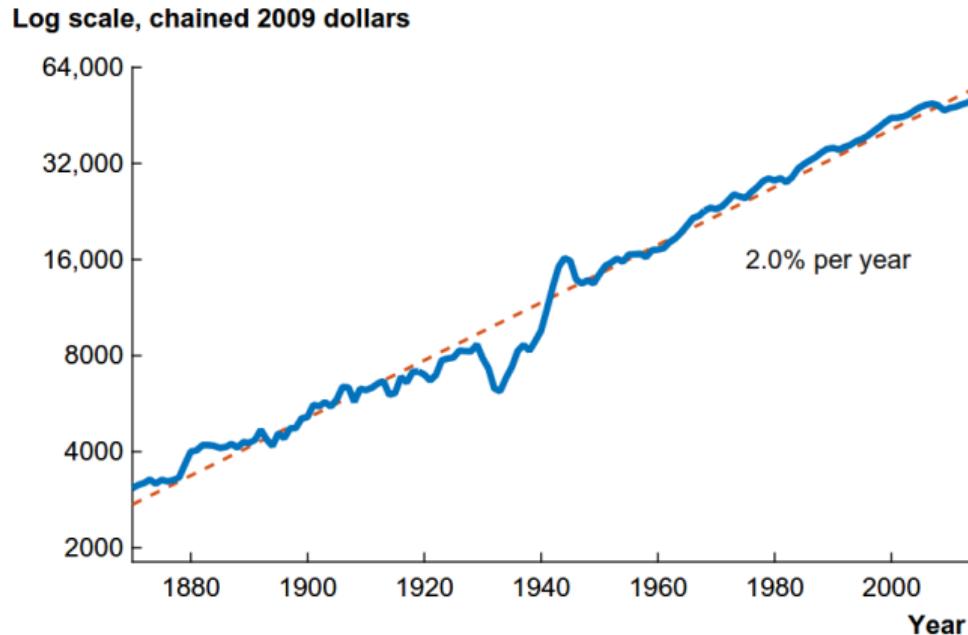
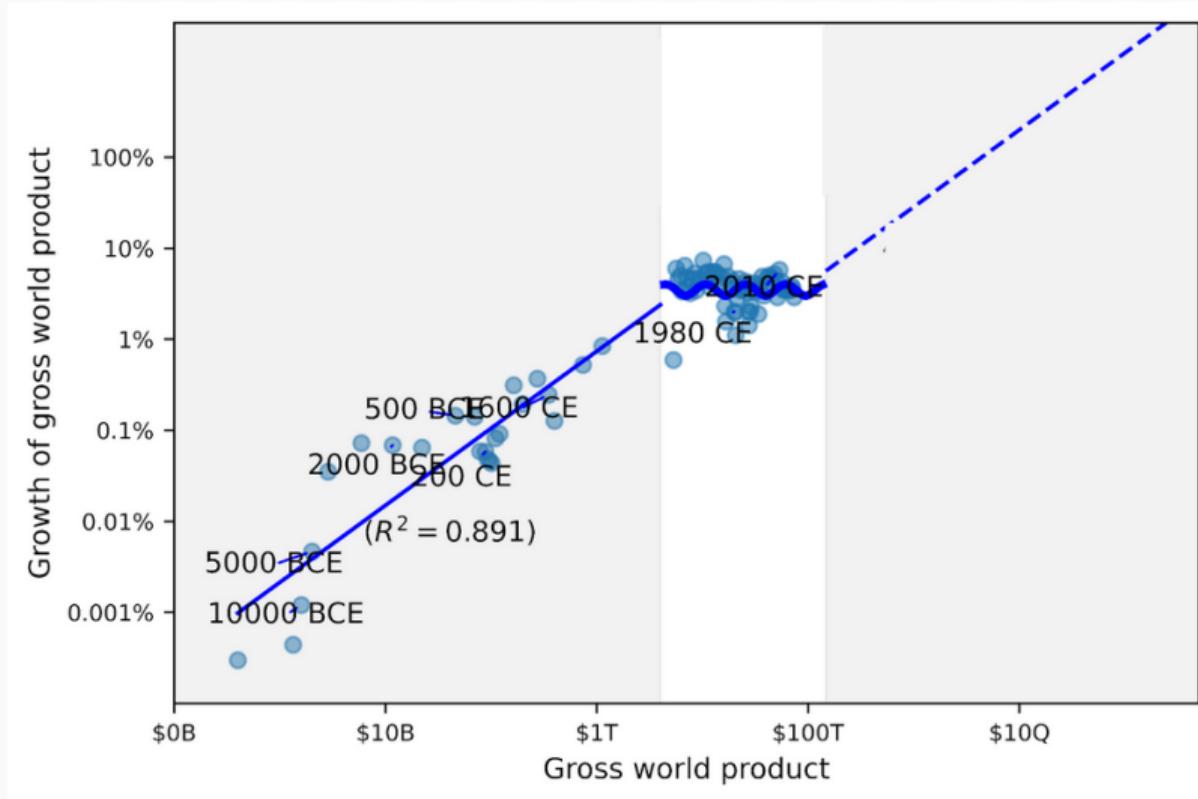


Fig. 1 GDP per person in the United States. Source: Data for 1929–2014 are from the U.S. Bureau of Economic Analysis, NIPA table 7.1. Data before 1929 are spliced from Maddison, A. 2008. Statistics on world population, GDP and per capita GDP, 1–2006 AD. Downloaded on December 4, 2008 from <http://www.ggdc.net/maddison/>.

Alternative prior: macrohistory shows hyperbolic growth

[Roodman 2020; Kremer 1993]



Theory: growth explosion is *possible*

Once machines can produce ideas, the limits to growth may no longer hold, and **growth rates could speed up, potentially even leading to a so-called 'singularity'** with infinite consumption

— Jones (2024)

$$\underbrace{\frac{\dot{A}_t}{A_t}}_{\text{TFP growth}} = \underbrace{A_t^{-\beta}}_{\text{ideas getting harder to find}} \cdot L_t^\alpha K_t^\beta$$

How can we forecast the development of transformative AI?

How can we forecast the development of transformative AI?

1. Speculation

- ▶ “AI 2027”
- ▶ Elon Musk: “AI will superset the intelligence of maybe all humans by 2027”

How can we forecast the development of transformative AI?

1. Speculation

- ▶ “AI 2027”
- ▶ Elon Musk: “AI will superset the intelligence of maybe all humans by 2027”

2. Surveys of AI researchers

▶ more

- ▶ Grace et al (2024): 2047

How can we forecast the development of transformative AI?

1. Speculation

- ▶ “AI 2027”
- ▶ Elon Musk: “AI will superset the intelligence of maybe all humans by 2027”

2. Surveys of AI researchers

▶ more

- ▶ Grace et al (2024): 2047

3. Models / trend extrapolation:

▶ more

- ▶ Kurzweil (2005): 2045
- ▶ Cotra (2022): 2040
- ▶ Davidson (2023): 2043

How can we forecast the development of transformative AI?

1. Speculation

- ▶ “AI 2027”
- ▶ Elon Musk: “AI will superset the intelligence of maybe all humans by 2027”

2. Surveys of AI researchers

▶ more

- ▶ Grace et al (2024): 2047

3. Models / trend extrapolation:

▶ more

- ▶ Kurzweil (2005): 2045
- ▶ Cotra (2022): 2040
- ▶ Davidson (2023): 2043

4. Asset prices

- ▶ Prices aggregate dispersed wisdom (Hayek 1945)...
- ▶ ...financial market prices especially so (Fama, etc)

▶ examples

This paper: a framework for contrasting *market expectations* of transformative AI vs. researcher beliefs

This paper: a framework for contrasting *market expectations* of transformative AI vs. researcher beliefs

Central point: short timelines for transformative AI would *increase real interest rates*

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

Want to recognize **double-edged sword** nature (Jones 2023):

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

Want to recognize double-edged sword nature (Jones 2023):

1. May rapidly accelerate growth...

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

Want to recognize double-edged sword nature (Jones 2023):

1. May rapidly accelerate growth...
2. ...may pose an “existential risk”

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

Want to recognize double-edged sword nature (Jones 2023):

1. May rapidly accelerate growth...
2. ...may pose an “existential risk”

Definition (Aligned transformative AI)

Technology that causes growth in global GDP in excess of 30% per year.

- ▶ Follows Davidson (2023): a 10× increase in growth

“Transformative AI”: defining the scenario under consideration

Definition (Transformative AI)

“Artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution”

Want to recognize double-edged sword nature (Jones 2023):

1. May rapidly accelerate growth...
2. ...may pose an “existential risk”

Definition (Aligned transformative AI)

Technology that causes growth in global GDP in excess of 30% per year.

- ▶ Follows Davidson (2023): a 10× increase in growth

Definition (Unaligned AI)

Technology that causes human extinction.

Theory

Empirics

Discussion

Appendix

The basic logic

$$1 = \beta \mathbb{E}_t \left[\frac{U'(C_{t+1})}{U'(C_t)} \right] (1 + r_t)$$

Real interest rates are determined by the supply and demand for savings

Ramsey rule:

$$r = \rho + \frac{1}{\sigma} g$$

- ▶ r : real interest rate
- ▶ ρ : time discounting
- ▶ g : growth rate
- ▶ $\sigma > 0$: elasticity of
intertemporal substitution

Real interest rates are determined by the supply and demand for savings

Ramsey rule:

$$r = \rho + \frac{1}{\sigma}g$$

- ▶ r : real interest rate
- ▶ ρ : time discounting
- ▶ g : growth rate
- ▶ $\sigma > 0$: elasticity of intertemporal substitution

1. Time discounting and mortality risk

- “Intrinsic preference for the present”
 - +
- Probability of death
- Intuition: no reason to save if dead

Real interest rates are determined by the supply and demand for savings

Ramsey rule:

$$r = \rho + \frac{1}{\sigma} g$$

- ▶ r : real interest rate
- ▶ ρ : time discounting
- ▶ g : growth rate
- ▶ $\sigma > 0$: elasticity of intertemporal substitution

1. Time discounting and mortality risk

- “Intrinsic preference for the present”
 - +
- Probability of death
- Intuition: no reason to save if dead

2. Economic growth

- Intuition: consumption smoothing (“no reason to save if going to be rich”)

Real interest rates and transformative AI

Ramsey rule:

Aligned transformative AI: $g = 30\%$

$$r = \rho + \frac{1}{\sigma}g$$

- ▶ r : real interest rate
- ▶ ρ : time discounting
- ▶ g : growth rate
- ▶ $\sigma > 0$: elasticity of
intertemporal substitution

Real interest rates and transformative AI

Ramsey rule:

$$r = \rho + \frac{1}{\sigma}g$$

- ▶ r : real interest rate
- ▶ ρ : time discounting
- ▶ g : growth rate
- ▶ $\sigma > 0$: elasticity of intertemporal substitution

Aligned transformative AI: $g = 30\%$

Example calibration: $\rho = 1\%$, $\sigma = 1$, $g = 1\%$

Then:

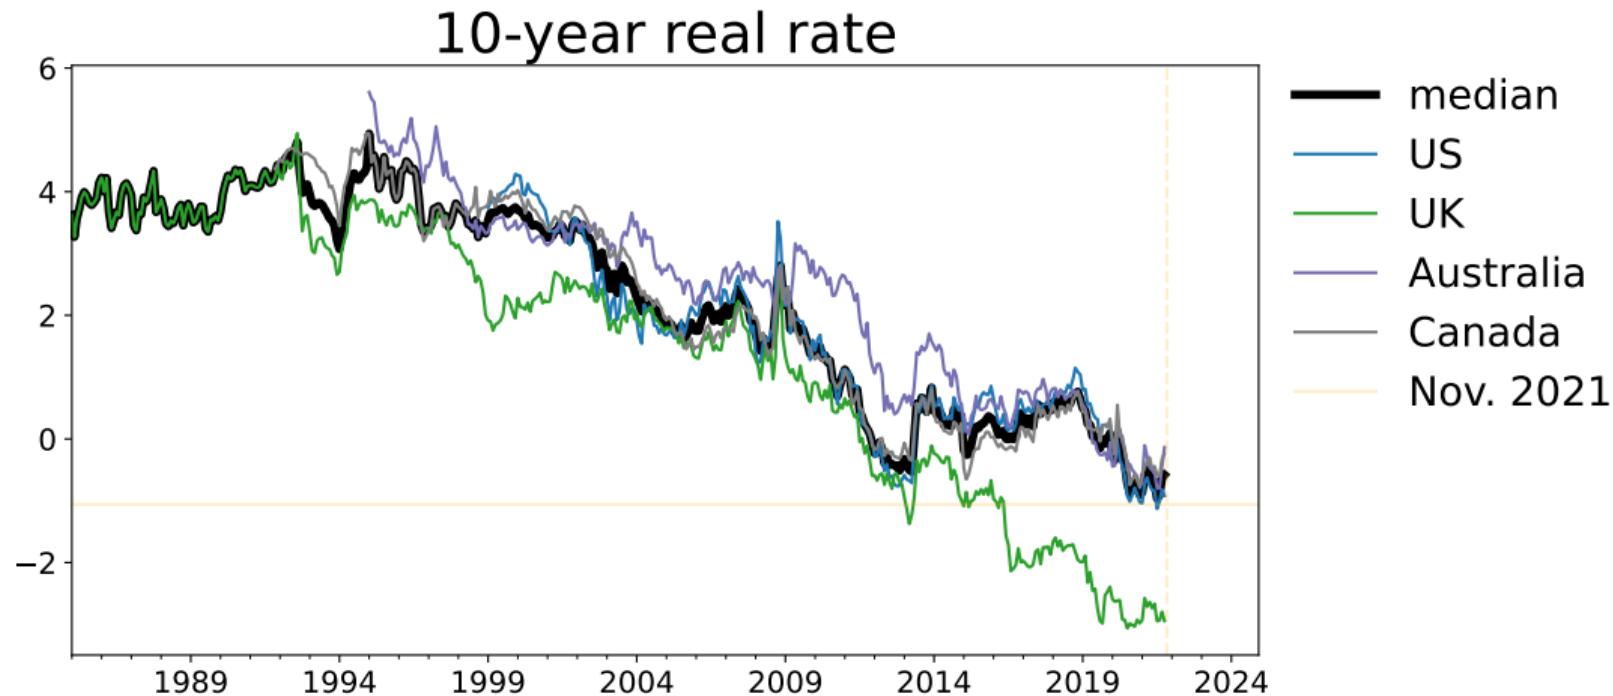
$$r = 2\%$$

vs.

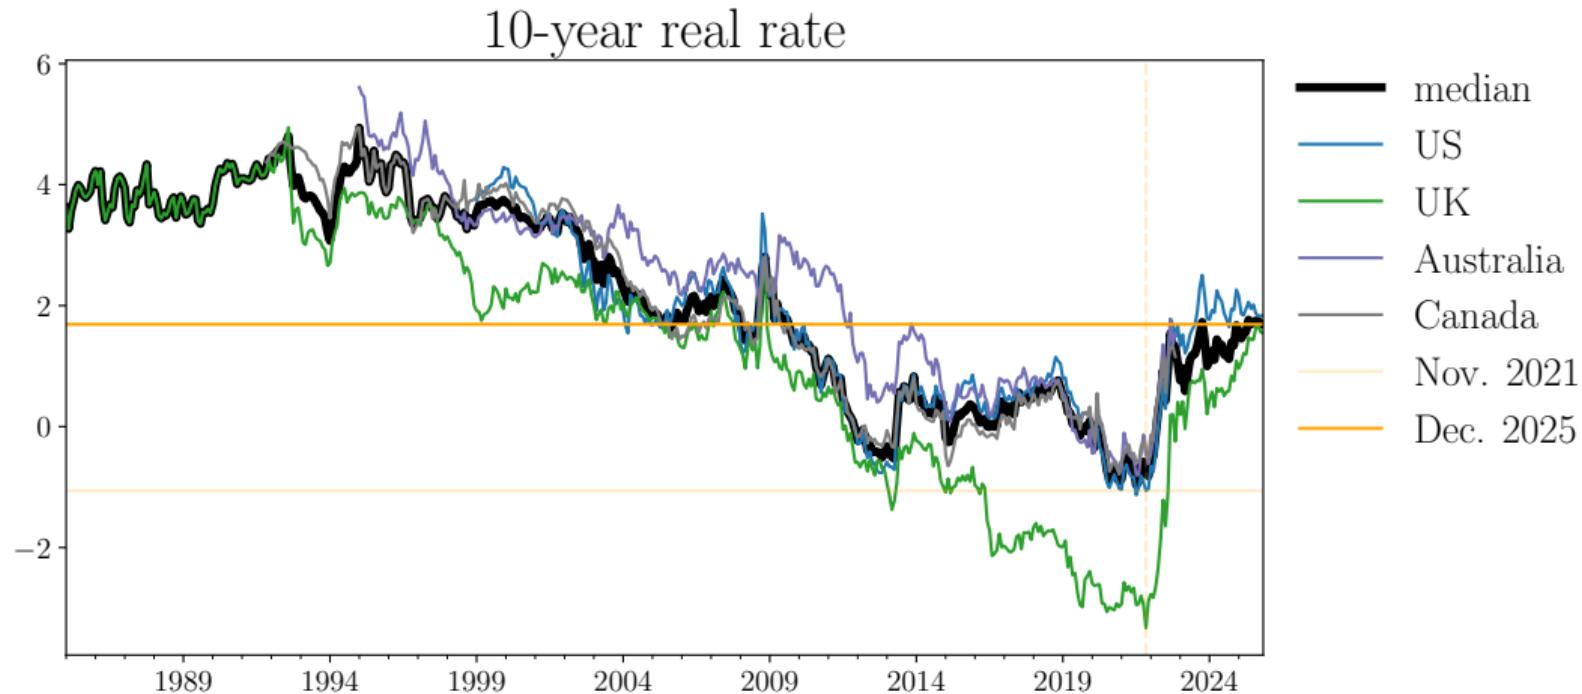
$$r = 31\% !$$

Recent history of real rates: at normal levels, though have risen

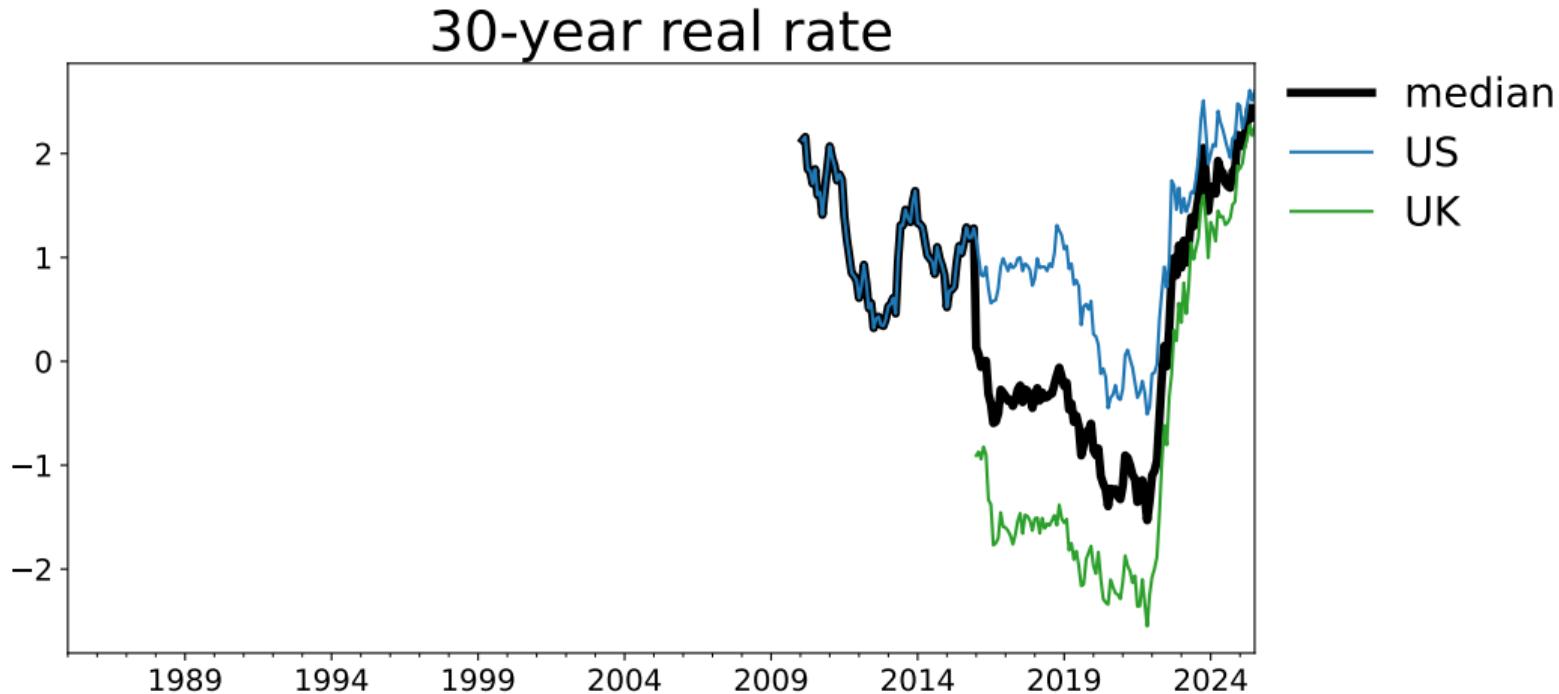
Recent history of real rates: at normal levels, though have risen



Recent history of real rates: at normal levels, though have risen



Recent history of real rates: at normal levels, though have risen



"I'm usually on the more frugal, conservative finance side. My whole life's been budgeting, and we're like, 'Nope, we're going to take that extra, very expensive vacation this year," said Sharon Korinek.

Korinek used to do chief financial officer-type work for private companies but hasn't returned to an office job because she doesn't see the point of it. Her husband, Anton, agrees. He's an economist at the University of Virginia who researches artificial intelligence and has aired his views at [Marketplace Tech](#).

The Korineks can fund that extra vacation partly because they're not saving for one very-big-ticket item in the future.

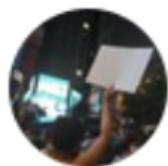
"I've had so many arguments about this. I'm like, 'Our kids aren't going to college,'" Sharon said. "Most people look at us like we're crazy."

Sharon and Anton have two kids, ages 8 and 6. The thought of opening a 529 college savings plan feels absurd to them, considering that they envision a future in which AI will be smarter than most humans.



Zachary Anglin @zachanglin · Sep 26, 2022

Have you?



Kelsey Piper
@KelseyTuoc

taken out large (>1M) long-term loans? yes

9:12 PM · Sep 26, 2022



Cate Hall 

@catehall

...

In 2017 I was convinced AI timelines were <5 years so I cashed in my 401k and blew away the money and let me tell you this particular form of intellectual consistency is Not Recommended

Benefits

For employees

- ✓ Health, dental, and vision insurance
- ✓ Mental healthcare support and services
- ✓ Commuter benefits
- ✓ 401(k) with generous matching
- ✓ Domestic conference budget for each employee

Theory

Empirics

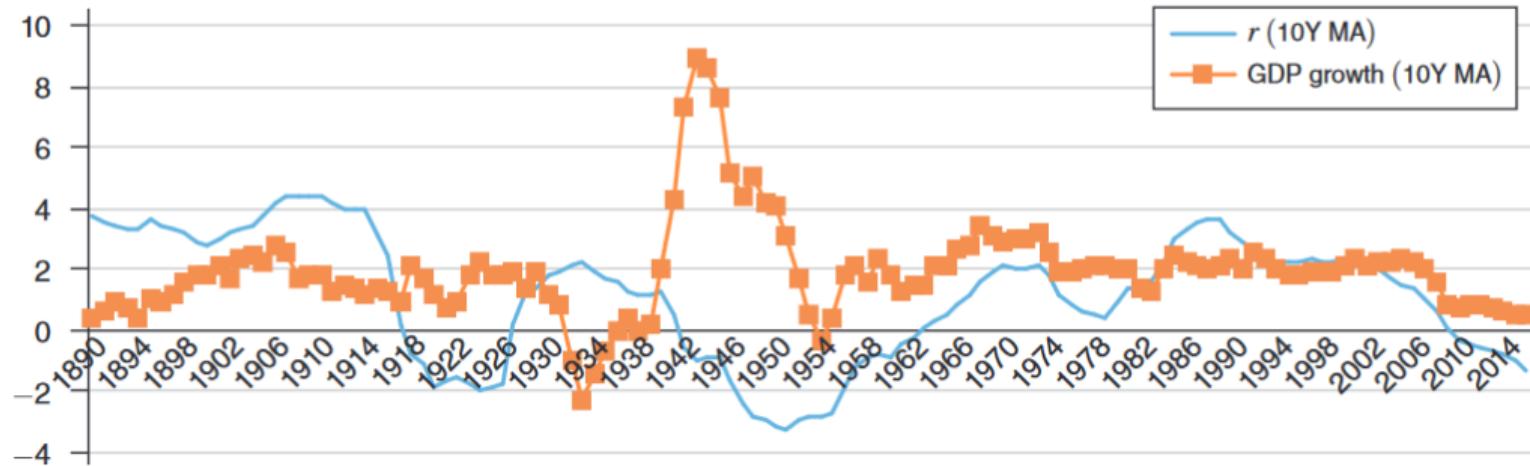
Discussion

Appendix

Existing literature: weak or no relationship between r and g [Rogoff-Rossi-Schmelzing 2024; Lunsford-West 2019; Hamilton et al 2016; Borio et al 2022; ...; Hall 1978]

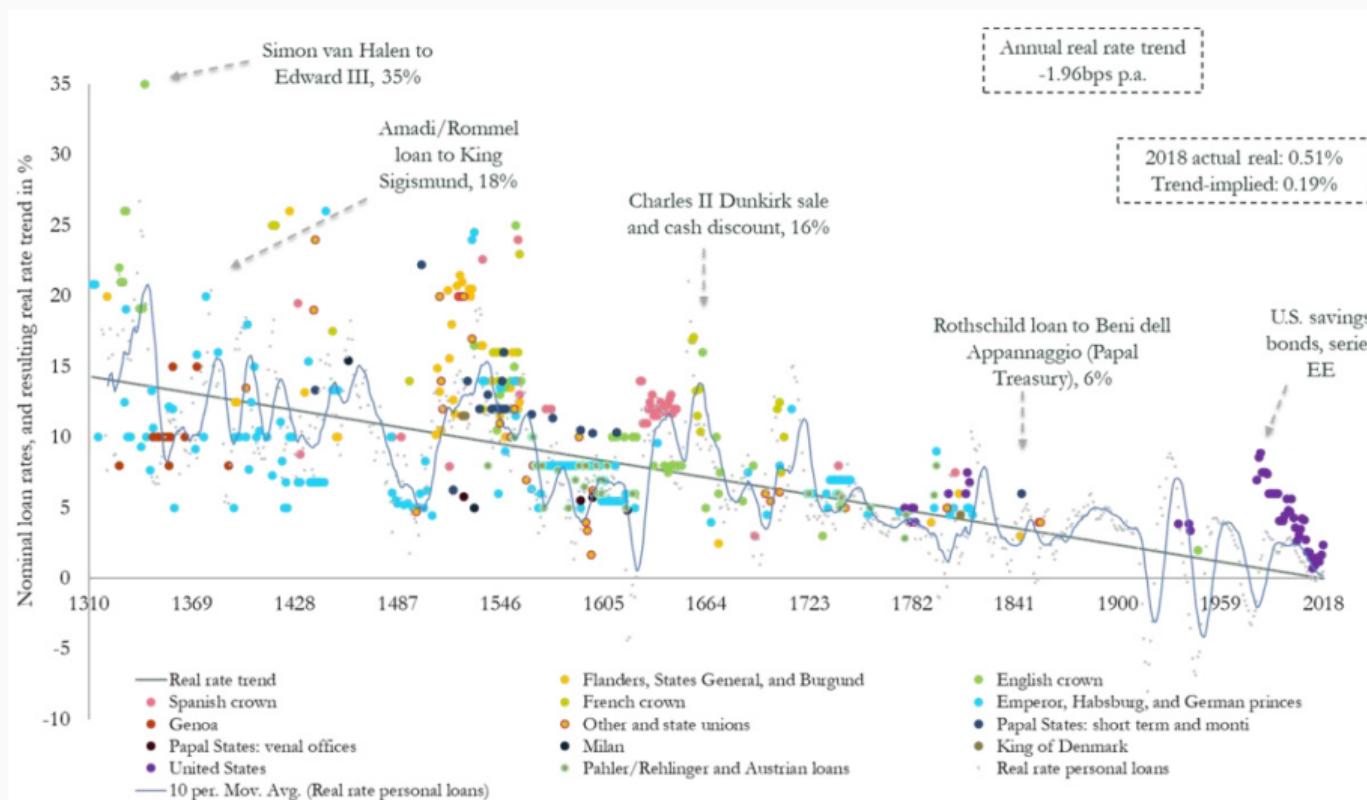
Existing literature: weak or no relationship between r and g [Rogoff-Rossi-Schmelzing 2024; Lunsford-West 2019; Hamilton et al 2016; Borio et al 2022; ...; Hall 1978]

Panel A. Real rate r_t , per capita GDP growth, ten-year moving averages



Existing literature: weak or no relationship between r and g [Rogoff-Rossi-Schmelzing

2024; Lunsford-West 2019; Hamilton et al 2016; Borio et al 2022; ...; Hall 1978]



Challenge 1: credible measurement of inflation expectations

► TIPS data

Goal: measure *long-term* real interest rates

$$r_t = i_t - \mathbb{E}_t \pi_{t+1}$$

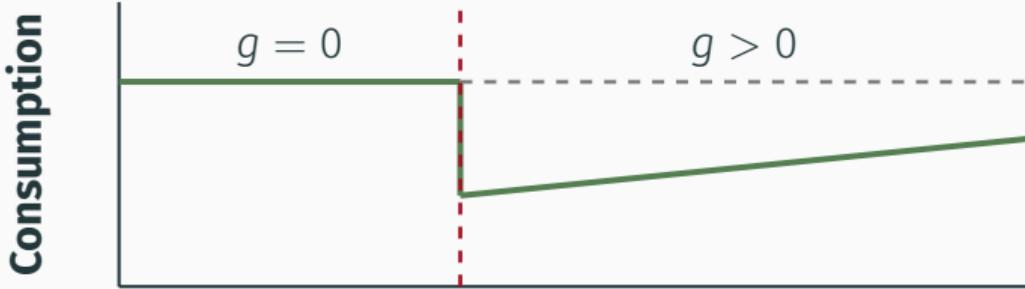
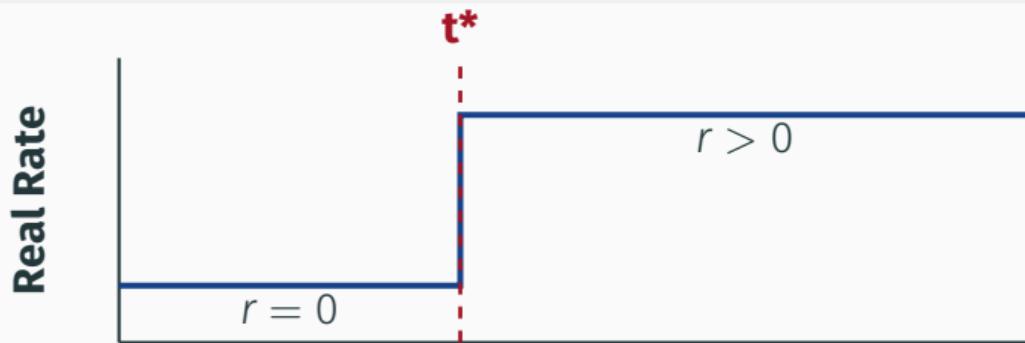
Goal: measure *long-term* real interest rates

$$r_t = i_t - \mathbb{E}_t \pi_{t+1}$$

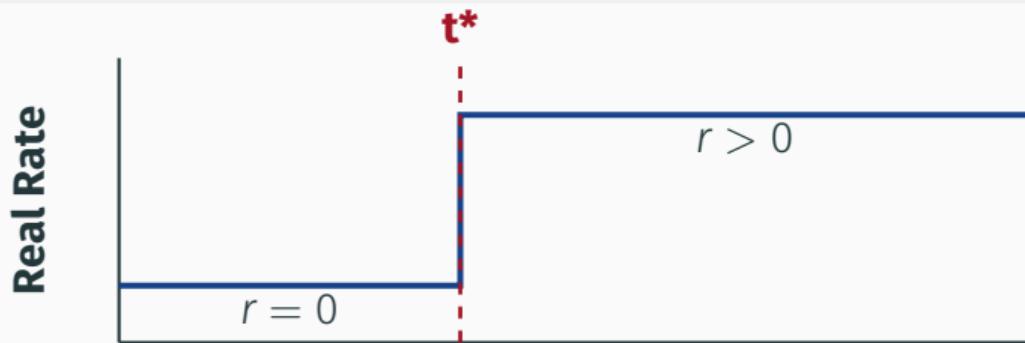
Inflation expectations: cross-country survey data

- 59 countries + up to 35 years of data (variation!)
- 1-to-10-year horizon
- Source: Consensus Economics

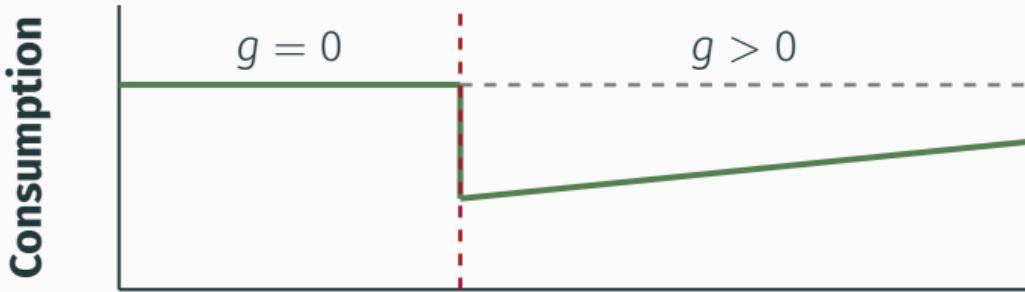
Challenge 2: distinguishing the short run and long run



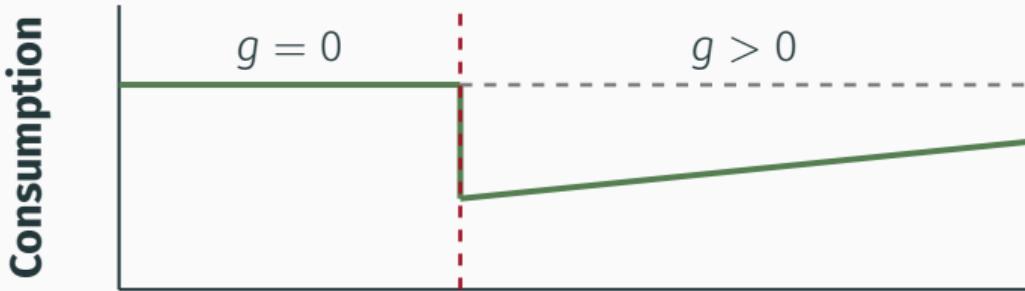
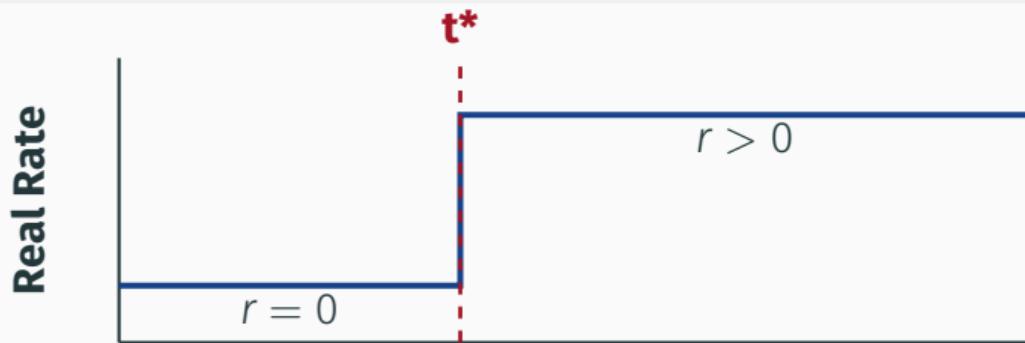
Challenge 2: distinguishing the short run and long run



⇒ Use *long-term* real interest rates



Challenge 2: distinguishing the short run and long run



⇒ Use *long-term* real interest rates

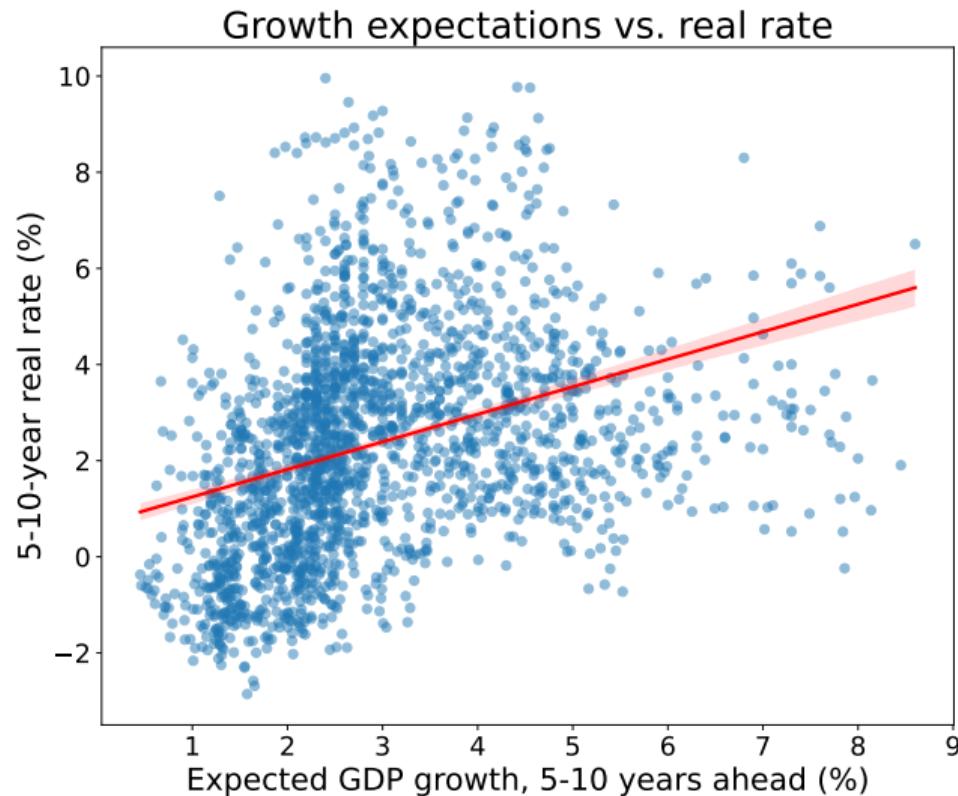
► Main analysis: 5y5y

Expected growth vs. ex ante real rates

▶ more

Expected growth vs. ex ante real rates

▶ more



Expected growth vs. ex ante real rates

▶ more

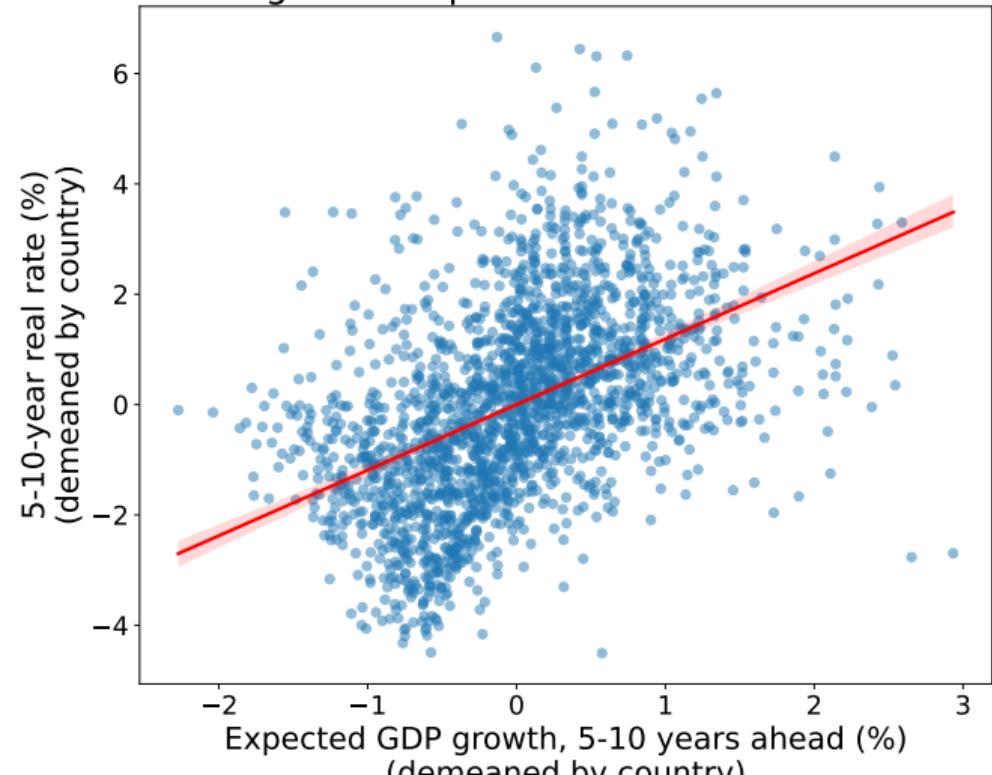
Table 1: Expected growth vs. real rate

	<i>Dependent variable: 5-10-year real rate</i>			
	(1)	(2)	(3)	(4)
5-10-year GDP growth forecast	0.71*** (0.12)	1.15*** (0.19)	1.80*** (0.21)	1.36*** (0.26)
SD(5-10-year GDP growth forecast)			-0.19 (0.47)	-0.36* (0.21)
5-year GDP growth forecast			-1.09*** (0.23)	-0.59** (0.28)
CDS spread			0.262** (0.112)	0.120** (0.059)
Country FE	No	Yes	No	Yes
Observations	3080	3080	2113	2113
Overall R^2	0.16	0.10	0.39	0.34
Within R^2		0.23		0.22

Expected growth vs. ex ante real rates

▶ more

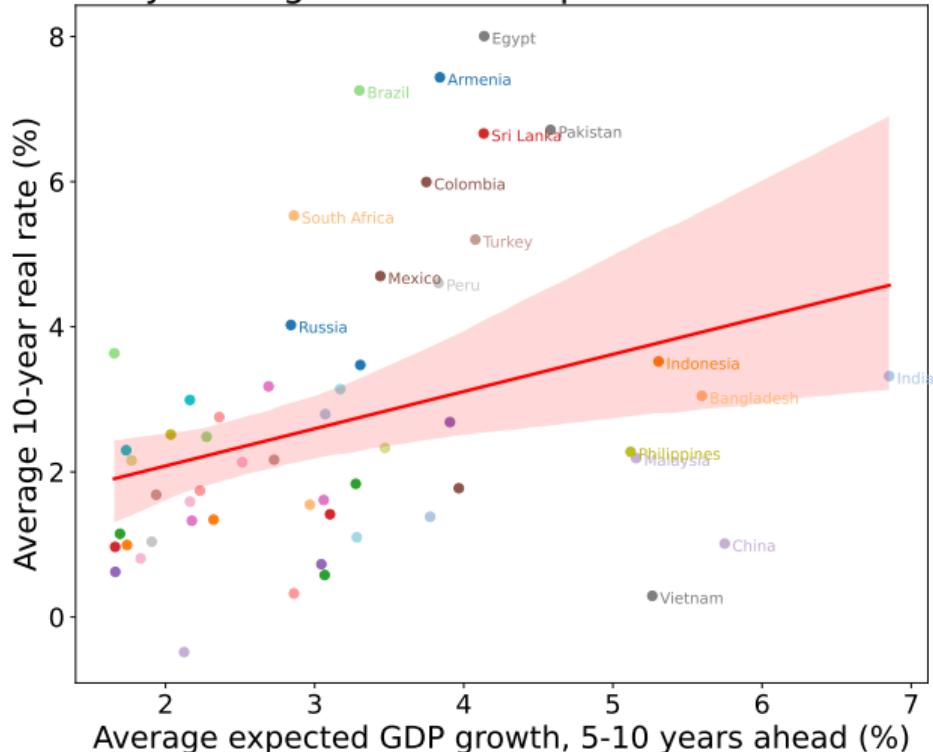
Demeaned growth expectations vs. demeaned real rate



Expected growth vs. ex ante real rates

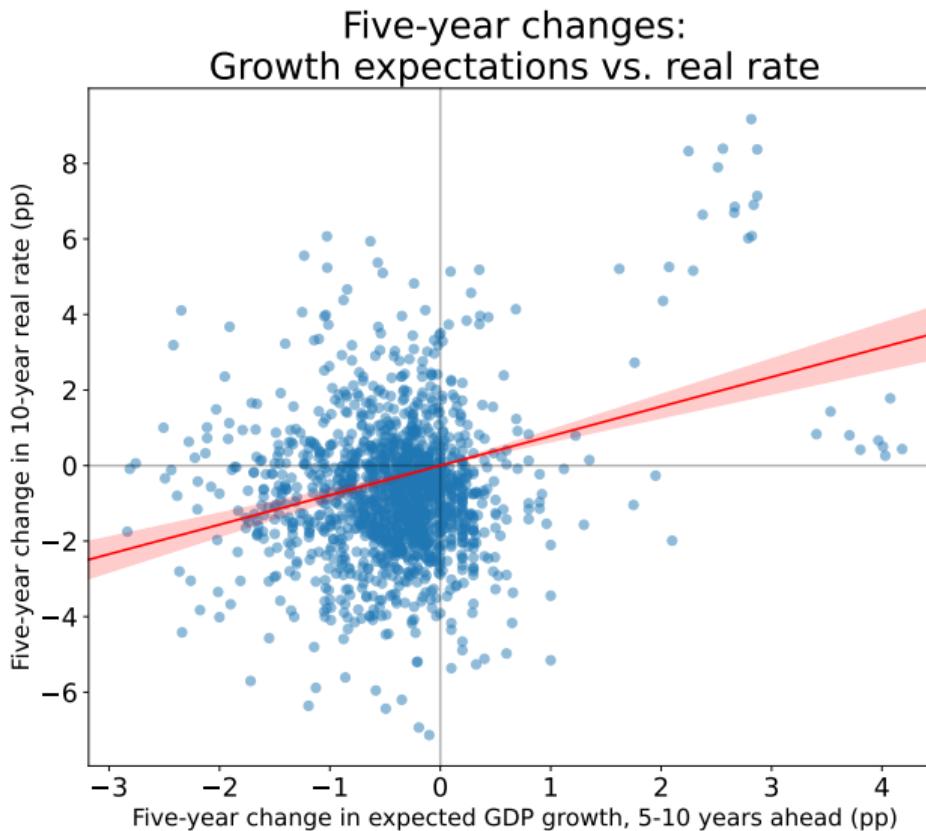
▶ more

Country averages: Growth expectations vs. real rate



Expected growth vs. ex ante real rates

▶ more



Taking the results extremely seriously

Taking the results extremely seriously

Baseline regression result
(no controls, no FE):

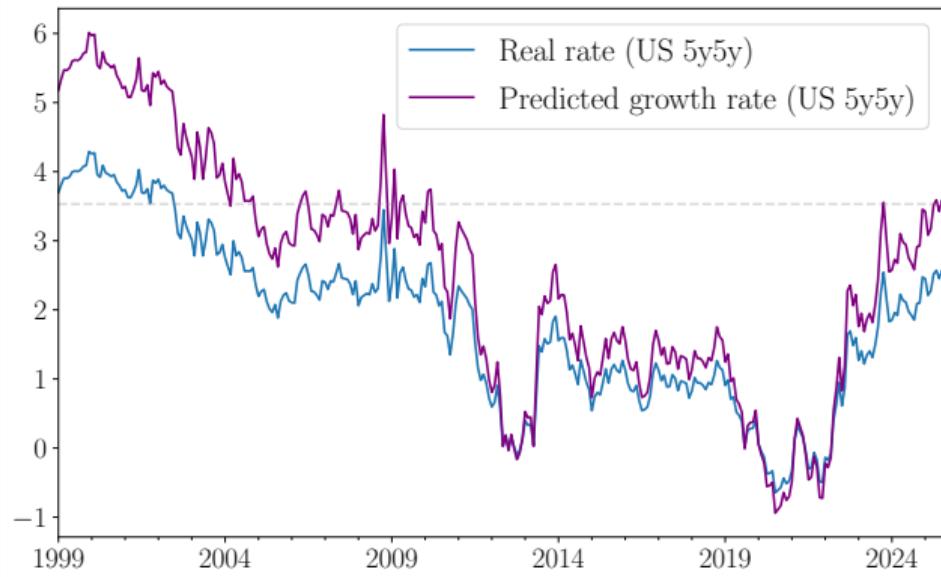
$$g_{5y5y} = \frac{1}{0.71} \cdot r_{5y5y} - 0.03$$

US 5y5y real rate as of Nov 30:

2.33%

Projected 5y5y growth rate:

= 3.25%



Theory

Empirics

Discussion

Appendix

Best arguments against

1. *Non-transformative* AI is consistent even with low rates

Best arguments against

1. Non-*transformative* AI is consistent
even with low rates

- ▶ Precautionary savings against
automation

Best arguments against

1. *Non-transformative* AI is consistent even with low rates

- ▶ Precautionary savings against automation
- ▶ TAI won't affect growth

Best arguments against

1. Non-*transformative* AI is consistent even with low rates

- ▶ Precautionary savings against automation
- ▶ TAI won't affect growth

2. Marginal utility may remain high post-singularity

- ▶ New products [Trammell (2024)]
- ▶ Habit formation

Best arguments against

1. Non-transformative AI is consistent even with low rates

- ▶ Precautionary savings against automation
- ▶ TAI won't affect growth

2. Marginal utility may remain high post-singularity

- ▶ New products [Trammell (2024)]
- ▶ Habit formation

Bad arguments:

- (i) “Want to invest more – to have a shot at controlling the lightcone”.
Maresca 2025: you want to save more, but this still pushes up r
- (ii) “High expected returns”: *movement along supply curve vs. shift in supply curve*
 - Also distinguish between high risk-free rate (discussion here) versus high *risk premium* (not discussed here)

Real rates and mortality risk

Challenge: measuring existential risk over time

Challenge: near-existential disasters could lower real rate!

- ▶ E.g. nuclear war \Rightarrow precautionary saving \Rightarrow lower real rate

Real rates and mortality risk

Challenge: measuring existential risk over time

Challenge: near-existential disasters could lower real rate!

- ▶ E.g. nuclear war \Rightarrow precautionary saving \Rightarrow lower real rate

Individual mortality risk and individual savings behavior

- ▶ \Rightarrow Test mechanism
- 1. New AIDS therapy rollout \Rightarrow more savings, more education
- 2. Information treatment \Rightarrow more investment in agriculture, livestock
- 3. Huntington's disease testing \Rightarrow 30pp less likely to go to college

Real rates and mortality risk

Challenge: measuring existential risk over time

Challenge: near-existential disasters could lower real rate!

- ▶ E.g. nuclear war \Rightarrow precautionary saving \Rightarrow lower real rate

Individual mortality risk and individual savings behavior

- ▶ \Rightarrow Test mechanism
- 1. New AIDS therapy rollout \Rightarrow more savings, more education
- 2. Information treatment \Rightarrow more investment in agriculture, livestock
- 3. Huntington's disease testing \Rightarrow 30pp less likely to go to college
- 4. Cross-sectionally: pessimistic survival beliefs are correlated with a lower savings rate

Real rates and mortality risk

Challenge: measuring existential risk over time

Challenge: near-existential disasters could lower real rate!

- ▶ E.g. nuclear war \Rightarrow precautionary saving \Rightarrow lower real rate

Individual mortality risk and individual savings behavior

- ▶ \Rightarrow Test mechanism
- 1. New AIDS therapy rollout \Rightarrow more savings, more education
- 2. Information treatment \Rightarrow more investment in agriculture, livestock
- 3. Huntington's disease testing \Rightarrow 30pp less likely to go to college
- 4. Cross-sectionally: pessimistic survival beliefs are correlated with a lower savings rate
- 5. Cold War evidence (sort of)

Other asset prices: stocks

Stocks are harder to use to forecast timelines:

1. Only reflect aligned case (future profits)

Other asset prices: stocks

Stocks are harder to use to forecast timelines:

1. Only reflect aligned case (future profits)
2. Stocks only reflect public companies (private companies, nonexistent companies)

Other asset prices: stocks

Stocks are harder to use to forecast timelines:

1. Only reflect aligned case (future profits)
2. Stocks only reflect public companies (private companies, nonexistent companies)
3. Capped profits / nationalization

Other asset prices: stocks

Stocks are harder to use to forecast timelines:

1. Only reflect aligned case (future profits)
2. Stocks only reflect public companies (private companies, nonexistent companies)
3. Capped profits / nationalization
4. TAI could lower stock prices: depends if $\sigma > 1$

$$P = \frac{D}{r - g}$$

Other asset prices: stocks

Stocks are harder to use to forecast timelines:

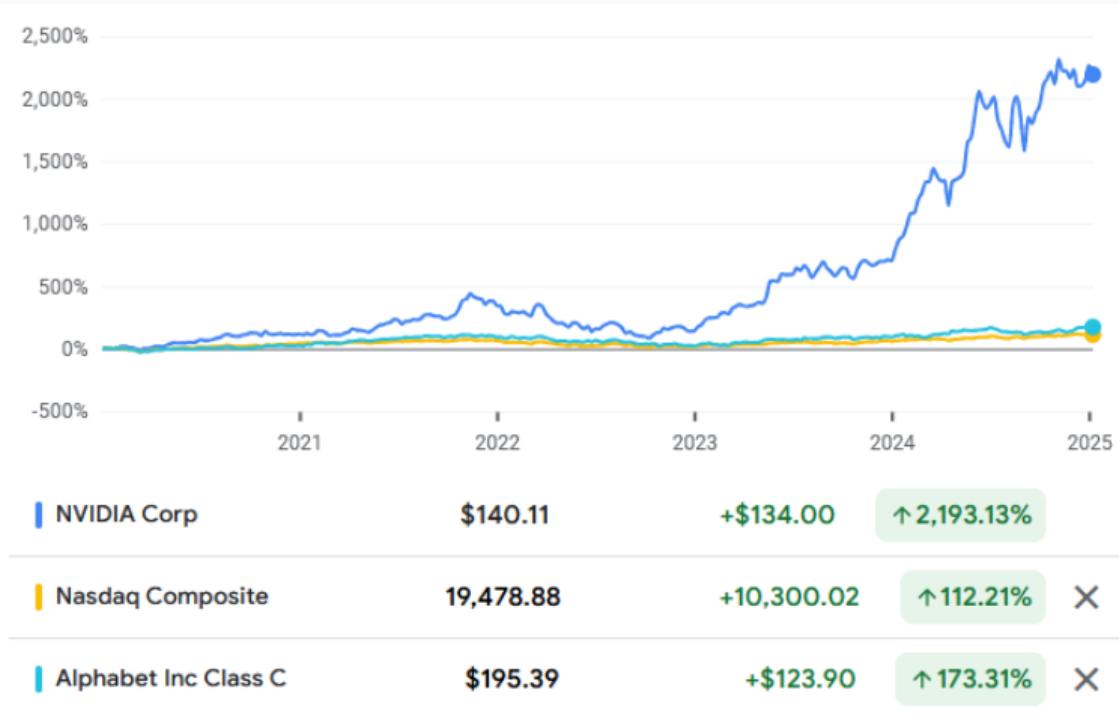
1. Only reflect aligned case (future profits)
2. Stocks only reflect public companies (private companies, nonexistent companies)
3. Capped profits / nationalization
4. TAI could lower stock prices: depends if $\sigma \gtrless 1$

$$\begin{aligned} P &= \frac{D}{r - g} \\ &= \frac{D}{(\rho + \frac{1}{\sigma} \cdot g) - g} \\ &= \frac{D}{\rho + (\frac{1}{\sigma} - 1) \cdot g} \end{aligned}$$

r affects all assets

Other asset prices: stocks

Given our estimate the $\sigma < 1$, **rising AI stock valuations** are consistent with **market expecting AI to be profitable, but not transformative**



Conclusion: Two possibilities

1. Markets are efficient information aggregators

2. Markets are wrong

- ▶ Trade opportunity? (“Get rich or die trying”)
- ▶ Opportunity to borrow cheaply? (“Impatient philanthropy”)

Conclusion: Two possibilities

1. Markets are efficient information aggregators

2. Markets are wrong

- ▶ Trade opportunity? (“Get rich or die trying”)
- ▶ Opportunity to borrow cheaply? (“Impatient philanthropy”)

Contribution:

1. **AI safety:** *outside view* evidence on AI timelines
2. **Mainline economics:** fundamental question about determinants of real interest rates

Theory

Empirics

Discussion

Appendix

1. Surveys

▶ back

1. Surveys

▶ back

- ▶ Grace et al. (2018) survey of AI researchers: “when unaided machines can accomplish every task better and more cheaply than human workers”
 - Median: **2061**
- ▶ Grace et al. (2022):
 - Median: **2058**

- ▶ Grace et al. (2018) survey of AI researchers: “when unaided machines can accomplish every task better and more cheaply than human workers”
 - Median: **2061**
- ▶ Grace et al. (2022):
 - Median: **2058**
- ▶ Grace et al. (2024):
 - Median: **2047**

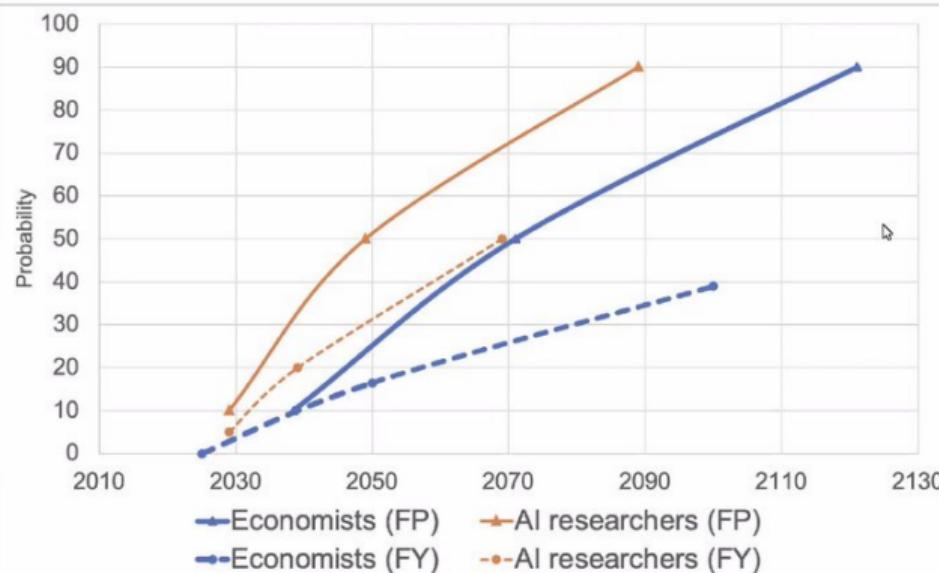
- ▶ Grace et al. (2018) survey of AI researchers: “when unaided machines can accomplish every task better and more cheaply than human workers”
 - Median: **2061**
- ▶ Grace et al. (2022):
 - Median: **2058**
- ▶ Grace et al. (2024):
 - Median: **2047**
- ▶ Metaculus forecasting platform
 - Median: **2031**

1. Surveys

▶ back

Korinek et al. (2022) survey of economists: **2070-2130+**

HLMI PREDICTIONS: ECONOMISTS & AI RESEARCHERS



Economists ascribe a median 10% probability to HLMI never being developed.

FP = fixed probability (shown probability, asked for years)

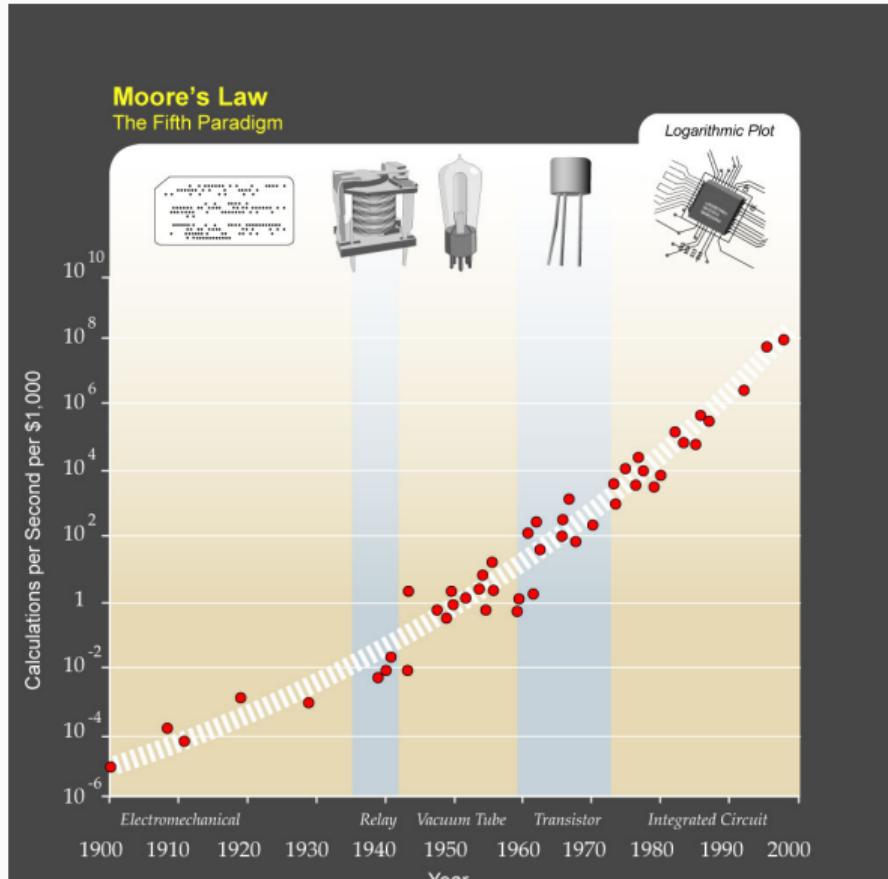
FY = fixed year (shown year, asked for likelihood)

2. Models / trend extrapolation

▶ back

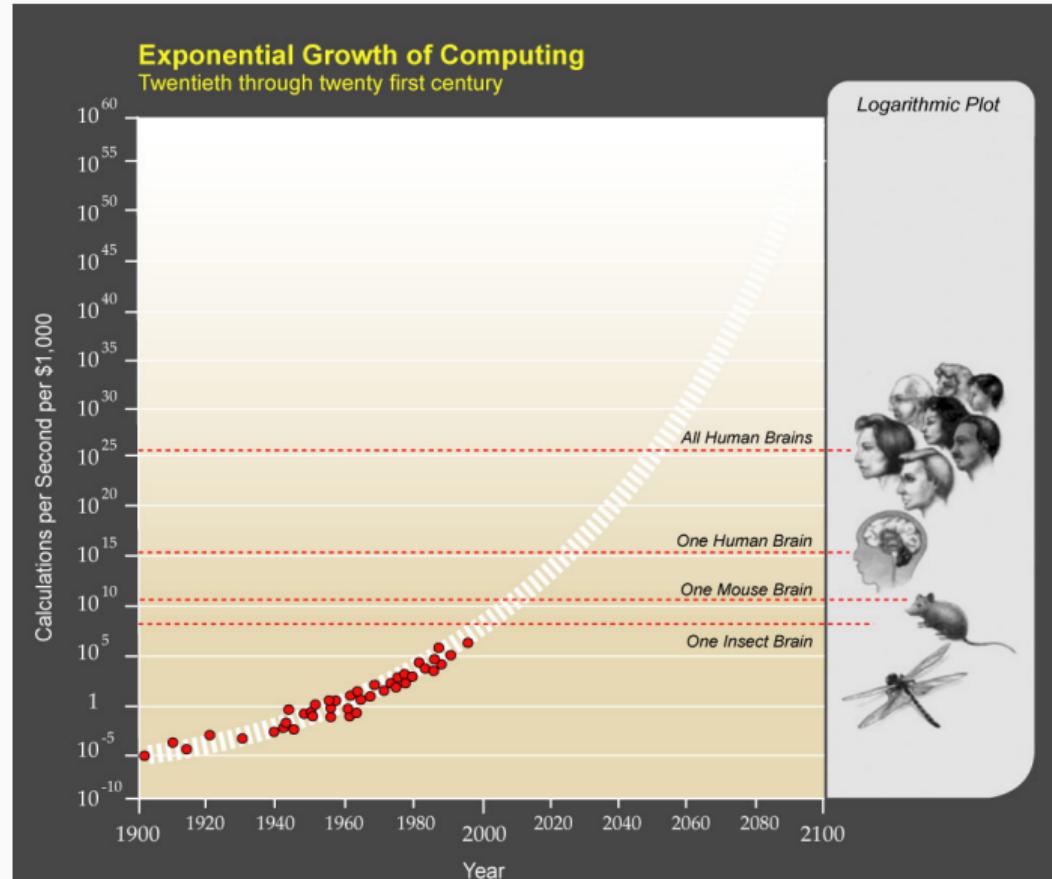
2. Models / trend extrapolation

▶ back



2. Models / trend extrapolation

▶ back



2. Models / trend extrapolation

▶ back

Deep and steep

Computing power used in training AI systems

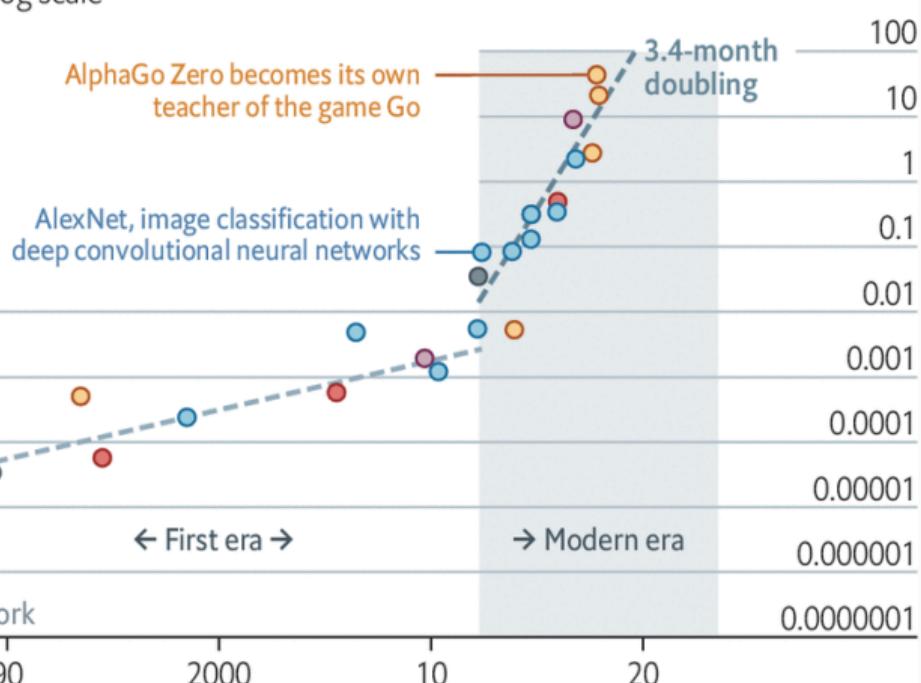
Days spent calculating at one petaflop per second*, log scale

By fundamentals

- Language
- Games
- Speech
- Other
- Vision

Two-year doubling
(Moore's Law)

Perceptron, a simple artificial neural network



Compute-centric forecasting models:

1. Estimate # of computations brain per second

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”
- ▶ Performance = compute \times \$ \times algorithms

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”
- ▶ Performance = compute \times $\$$ \times algorithms
 - Temporarily high growth

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”
- ▶ Performance = compute \times \$ \times **algorithms**
 - Temporarily high growth
 - Made up [cf Nostalgebraist, 2022]

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”
- ▶ Performance = compute \times \$ \times algorithms
 - Temporarily high growth
 - Made up [cf Nostalgebraist, 2022]
- ▶ Follow up work by Davidson (2022) and Epoch

Compute-centric forecasting models:

1. Estimate # of computations brain per second
2. Project forward trends in power of computers (e.g. Moore's Law)
3. \Rightarrow project date when computers can do as many calculations per second as brain

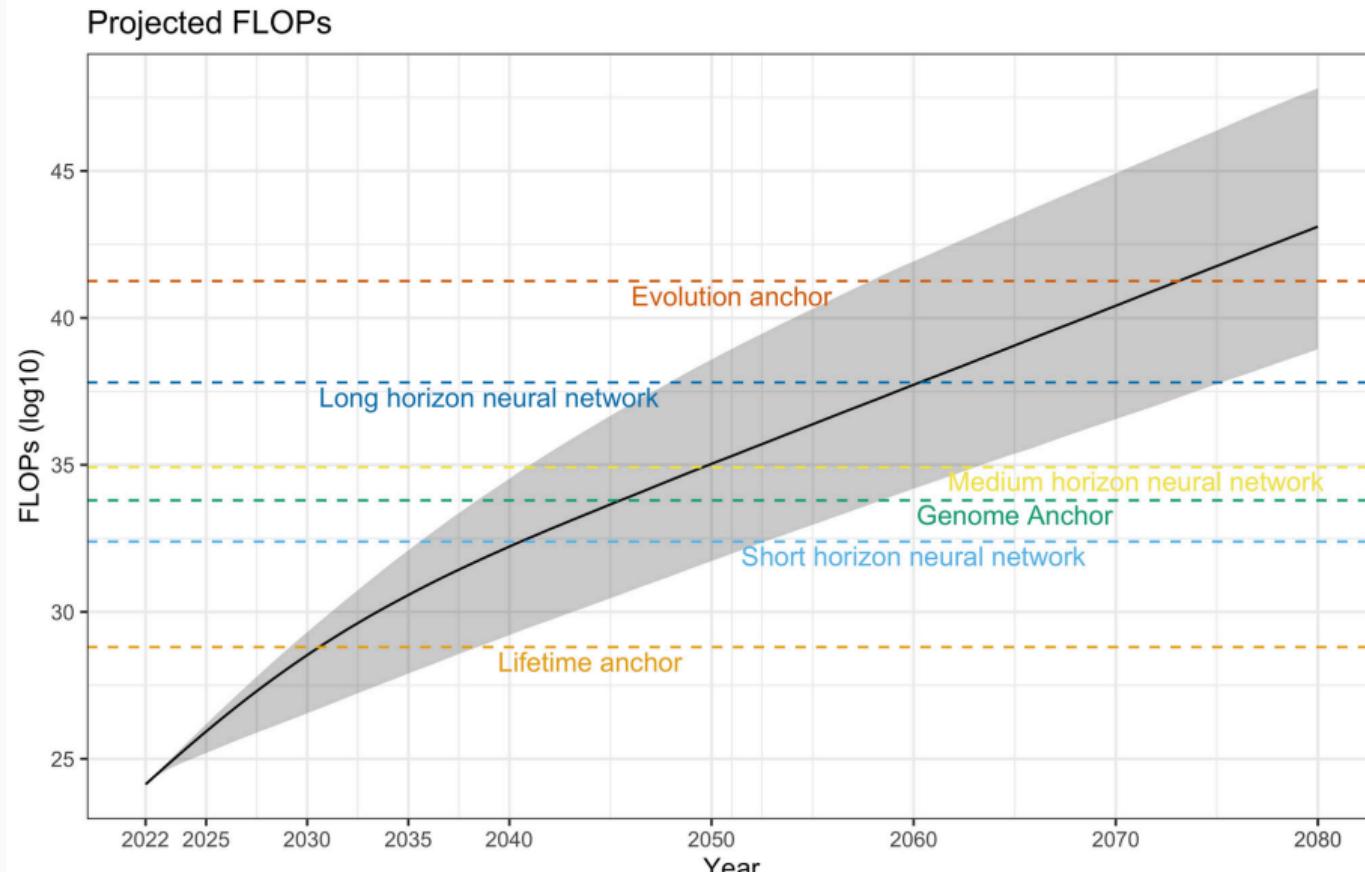
- ▶ **Cotra (2020):** 2050
- ▶ **Cotra (2022):** 2040
- ▶ **Davidson (2023):** 2043
- ▶ **Epoch (2022):** 2046

Seminal paper: Cotra (2020)

- ▶ “Biological anchors” or “bio anchors”
- ▶ Performance = compute \times \$ \times algorithms
 - Temporarily high growth
 - Made up [cf Nostalgebraist, 2022]
- ▶ Follow up work by Davidson (2022) and Epoch

Bio anchors (Epoch version)

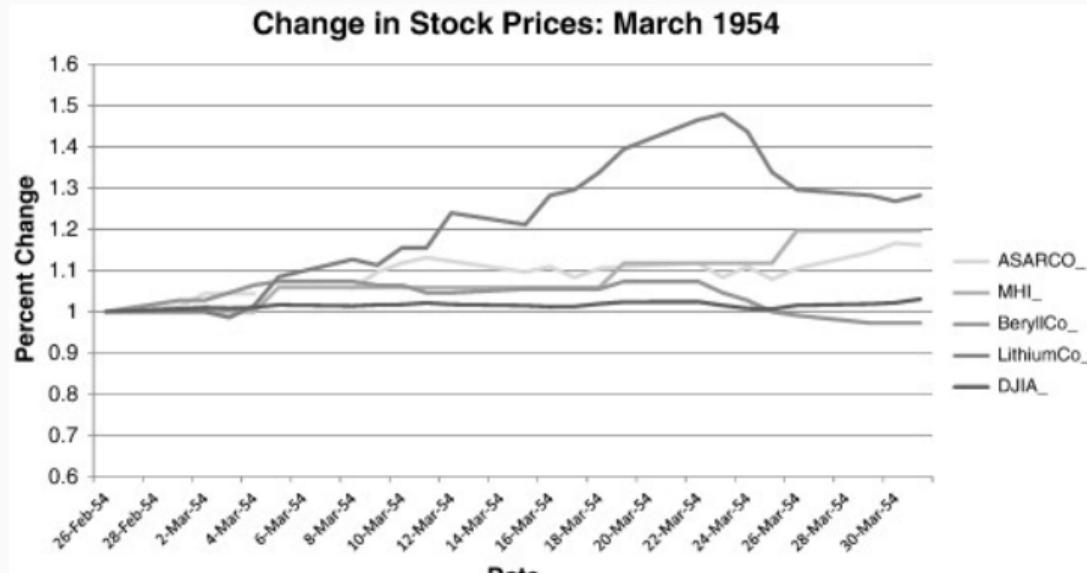
▶ back



Prices aggregate dispersed information (Hayek 1945); financial market prices especially so (Fama, etc)

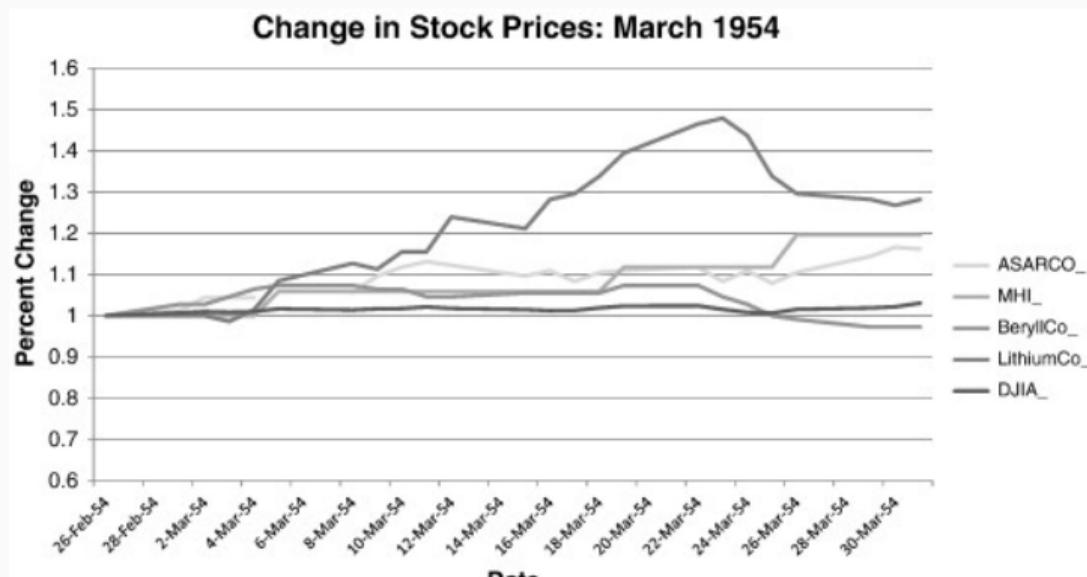
Prices aggregate dispersed information (Hayek 1945); financial market prices especially so (Fama, etc)

- ▶ Alchian and the hydrogen bomb

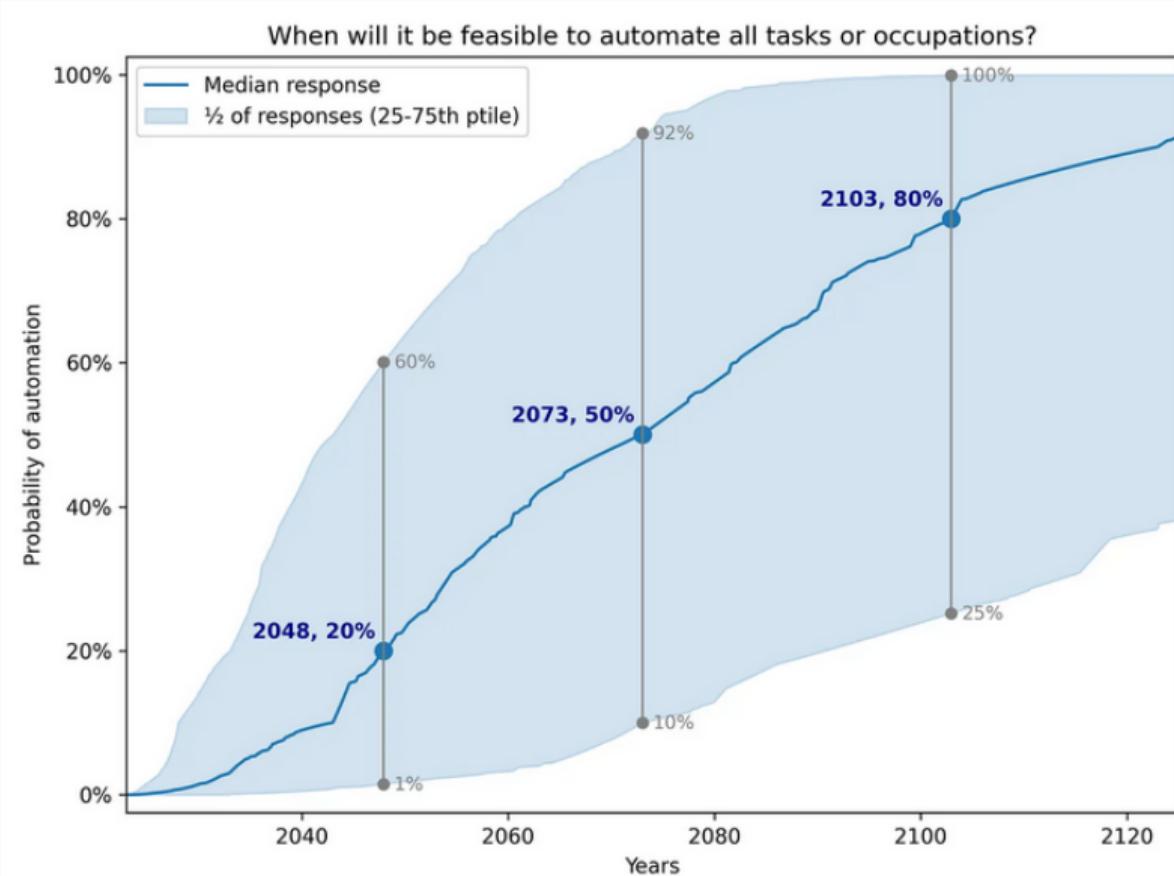


Prices aggregate dispersed information (Hayek 1945); financial market prices especially so (Fama, etc)

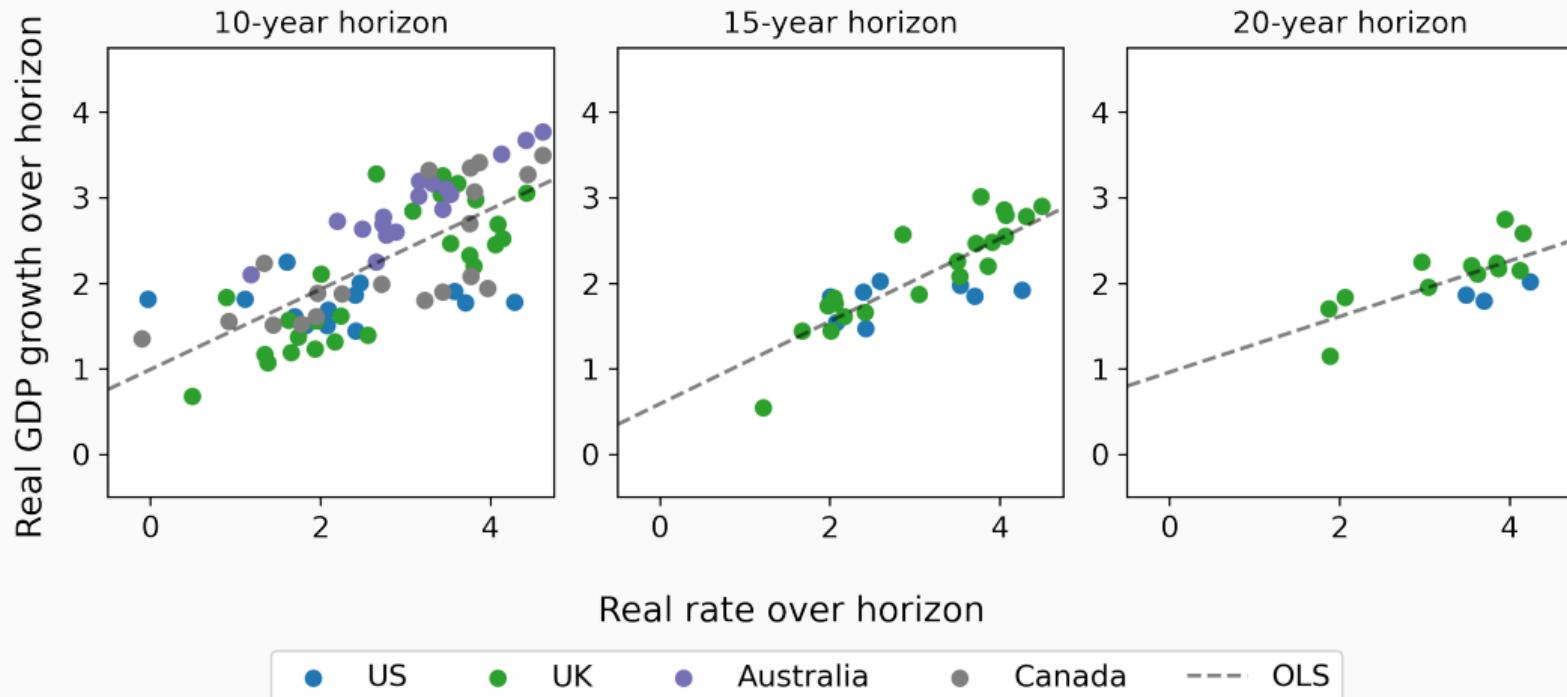
- ▶ Alchian and the hydrogen bomb
- ▶ Space shuttle Columbia disaster; election markets; inflation breakevens; ...



Weinstein-Raun (2024): Grace et al (2024) reanalysis

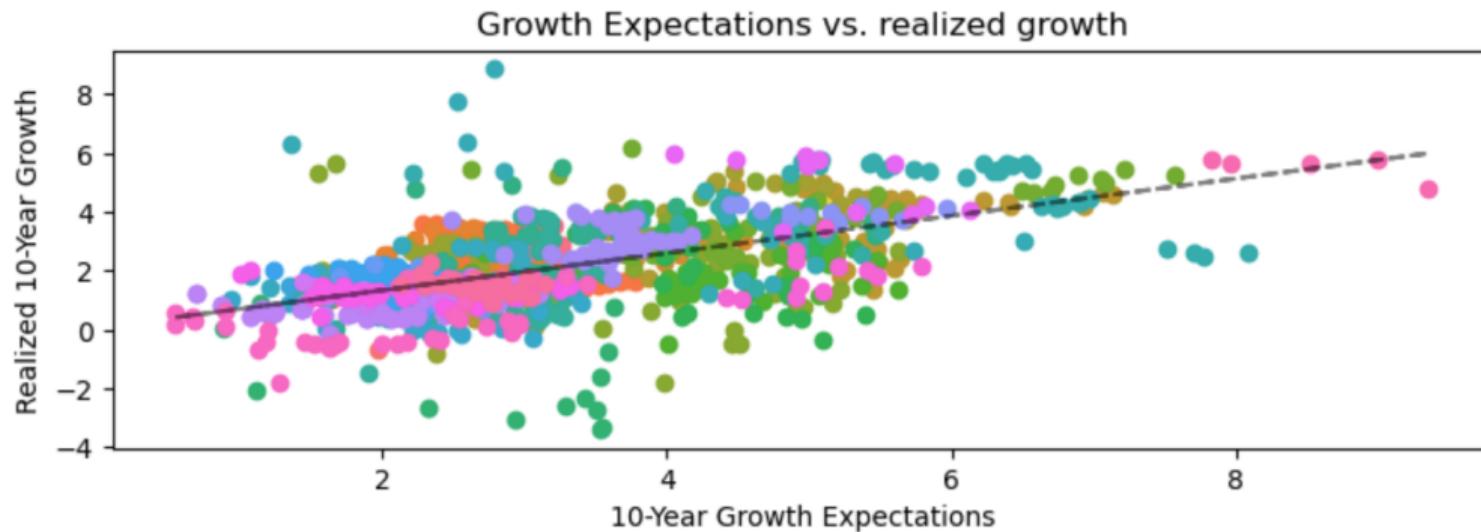


Real rate vs. future real GDP growth



Expected growth vs. realized growth

▶ back

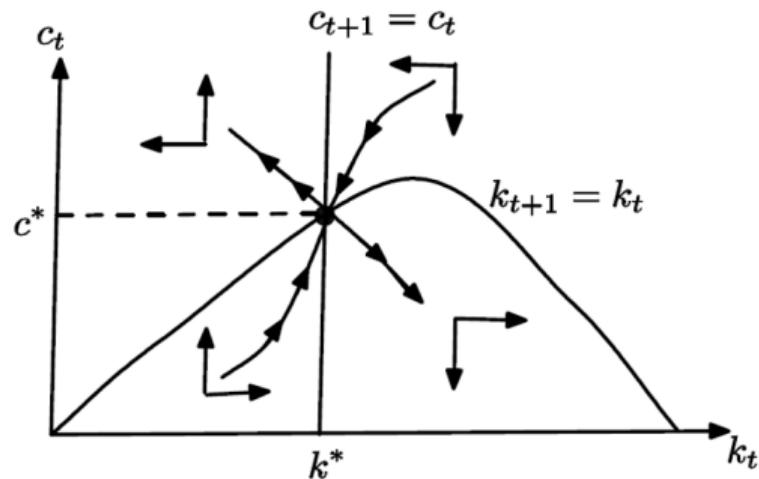


Robustness

- 1. Preferences:** recursive utility (Flynn, Schmidt, and Toda 2023) and under internal habit formation (Bhamra and Uppal 2014; Hamilton et al. 2016; Dennis 2009)
- 2. Incomplete markets:** Werning (2015)
- 3. Belief heterogeneity:** Buraschi and Whelan (2022), Molavi et al (2025)

Phase diagram analysis

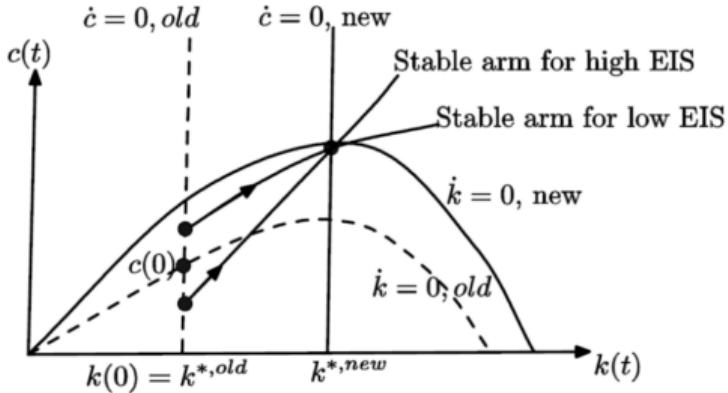
NCG:



► Neoclassical growth

Phase diagram analysis

Permanent productivity shock today:

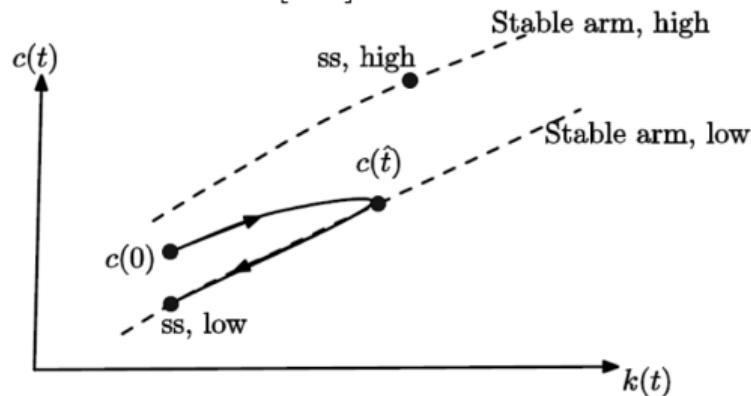


- ▶ Consumption can jump up, jump down, or be constant; but then increases to new SS
- ▶ Depends on EIS i.e. income vs. substitution effects: if $EIS > 1$, then more willing to trade consumption tomorrow for consumption today, i.e. substitution outweighs income effect and jumps down; vice-versa if $EIS < 1$

Phase diagram analysis

Transitory TFP shock:

$A' > A$ for $t \in [0, \hat{t}]$, then revert to A



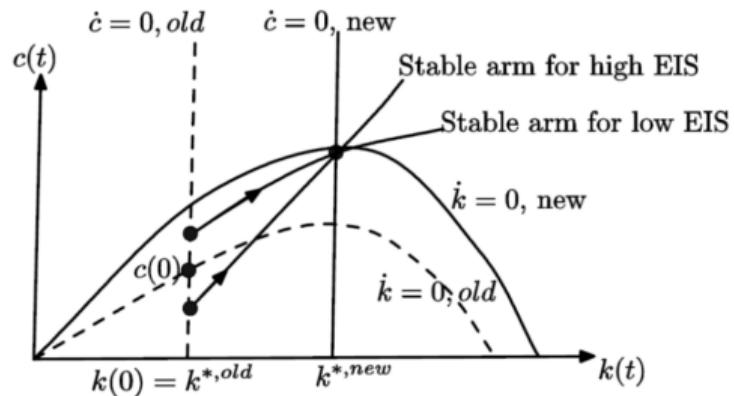
- ▶ Between $[0, \hat{t}]$, follow the dynamics of the new system
- ▶ At time \hat{t} , need to land on original stable arm
- ▶ Consumption jumps up; then could increase, decrease, or stay constant; until \hat{t} when it decreases back to original SS. Thus, a cusp in the time-path for c

Phase diagram analysis

News about future productivity shock:

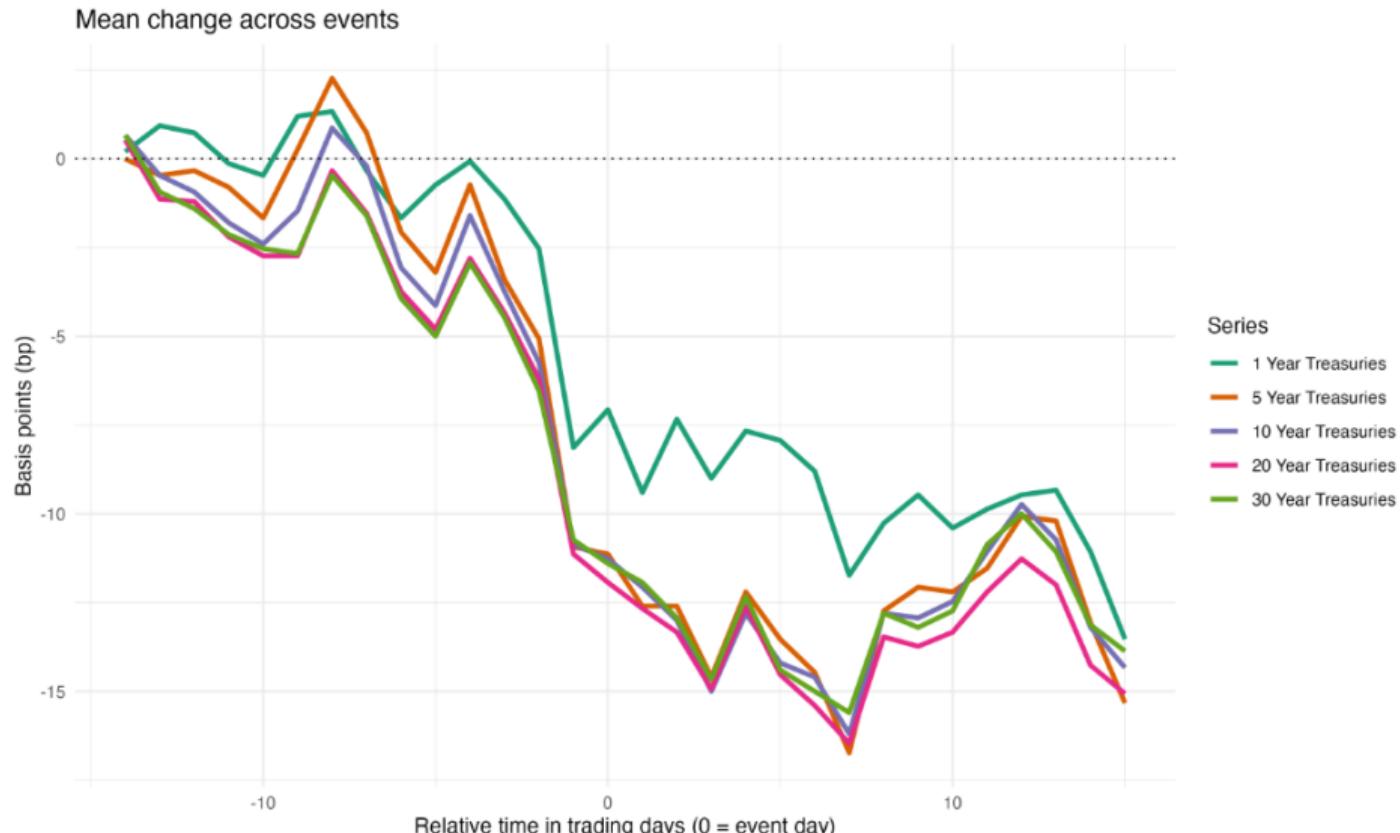
$A' > A$ at $t = \hat{t}$; but news received at $t = 0$

Permanent productivity shock today:



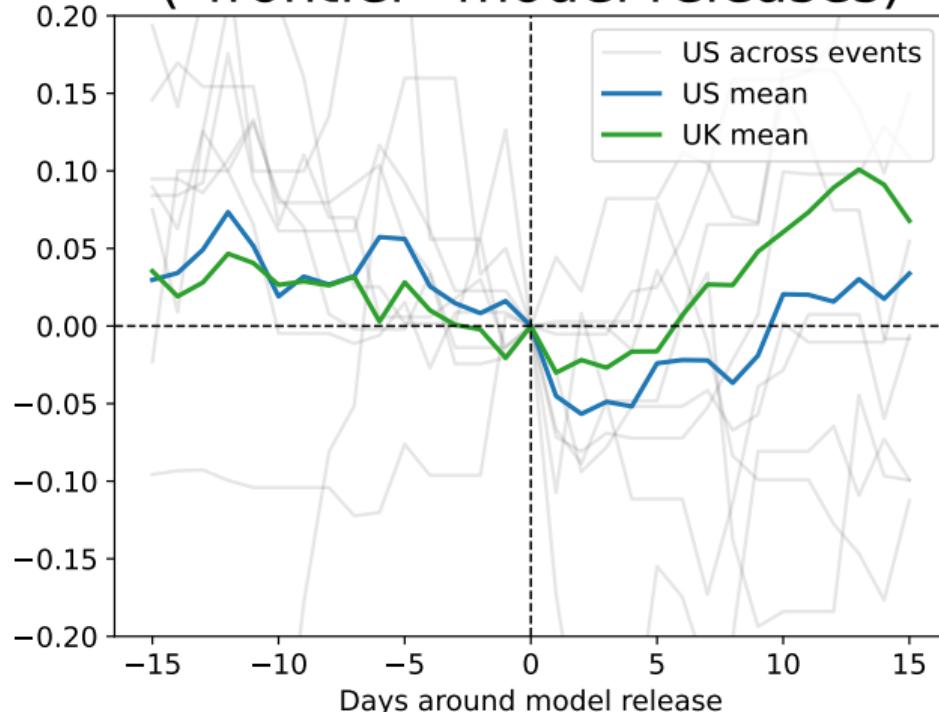
- ▶ From $[0, \hat{t}]$, you follow original dynamics
- ▶ At \hat{t} , you have to be on new saddle path
- ▶ Need to jump between old and new saddle path, otherwise you explode

Andrews and Farboodi (2025): event study around LLM releases



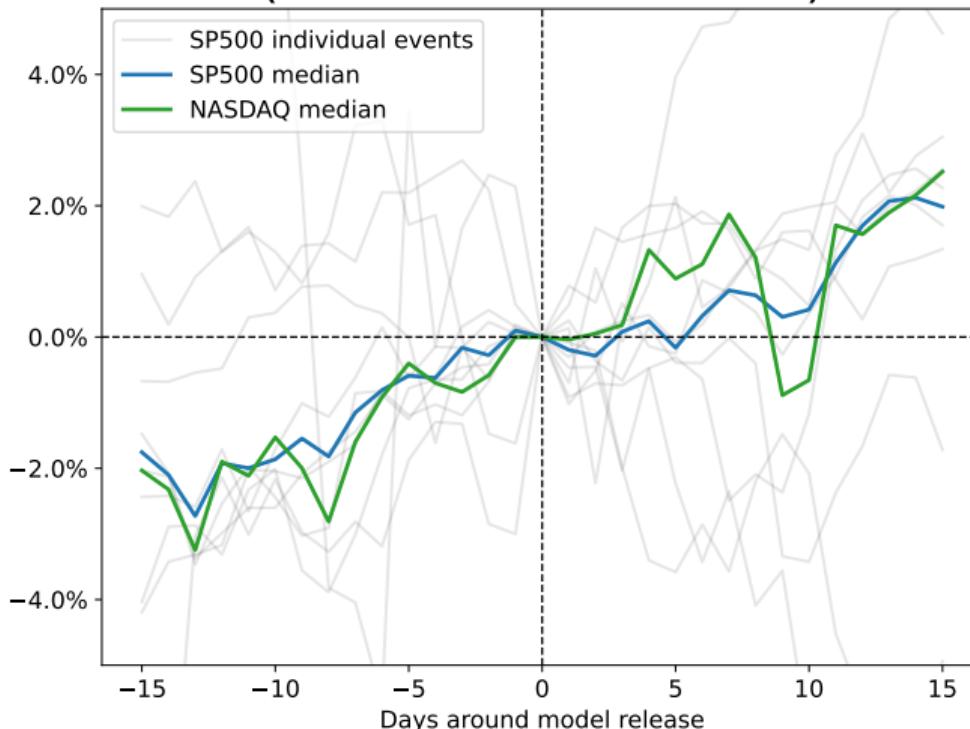
Andrews and Farboodi (2025): event study around LLM releases

10y real rate around model releases
("frontier" model releases)



Stock market event study around Andrews-Farboodi dates

Median equities around model releases
("frontier" model dates)



The world's simplest structural model

- ▶ Take the (nonlinear) Euler equation
- ▶ Plug in smoothed **Cotra (2022) probabilities for transformative AI over the next 30 years:** a 2% yearly chance until 2030, a 3% yearly chance through 2036, and a 4% yearly chance through 2052
- ▶ We use the **FTX Future Fund's median estimate of 15% for the probability that AI is unaligned** conditional on the development of transformative AI.
- ▶ Arrival of aligned AI: 30% annual consumption growth.
- ▶ Extinction: utility goes to zero.
- ▶ In absence of TAI: 1.8% growth rate
- ▶ Calibrate the pure rate of subjective time preference to 0.01
- ▶ Calibrate EIS = 1

Play with the model yourself! [here](#)

Euler equation model: baseline results

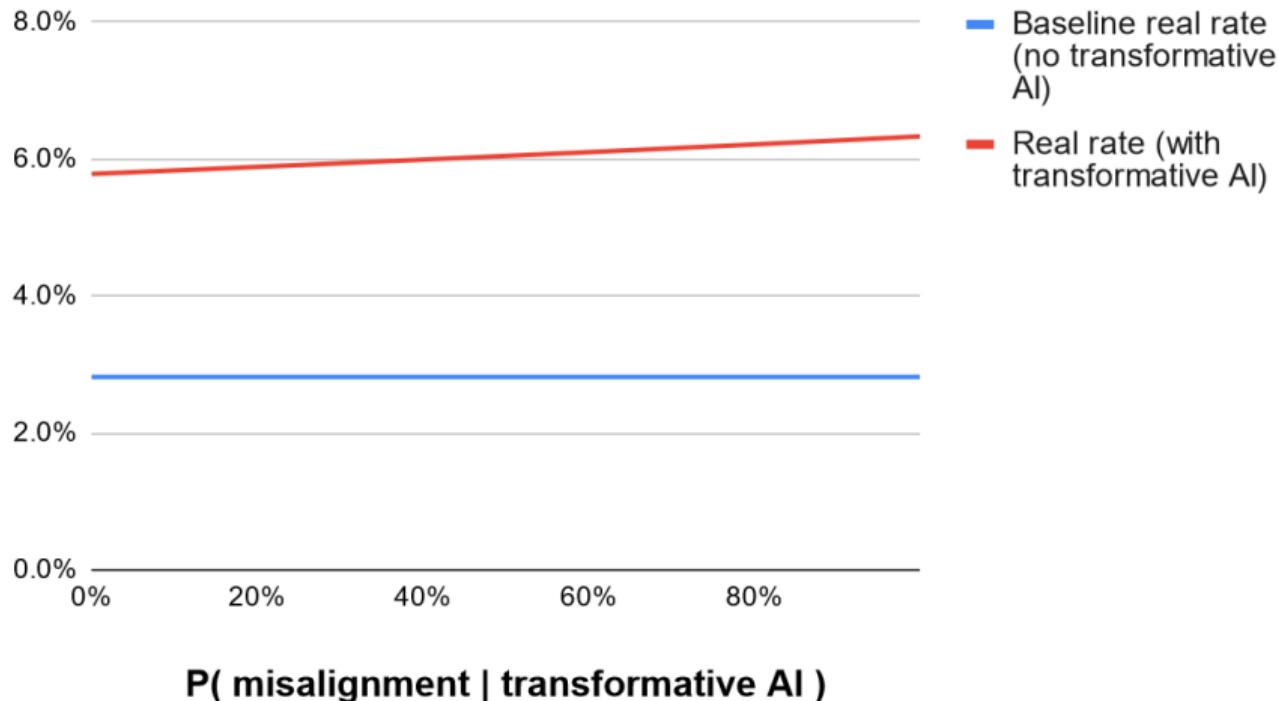
By default, $g = 1.8\%$ per year. Every year, some probability (based on Cotta) TAI developed. If developed, 15% probability world ends + 85% chance $g = 30\%$

Baseline Euler: $r = 2.8\%$

Euler + stochastic TAI: $r_{0,30} = 5.9\%$

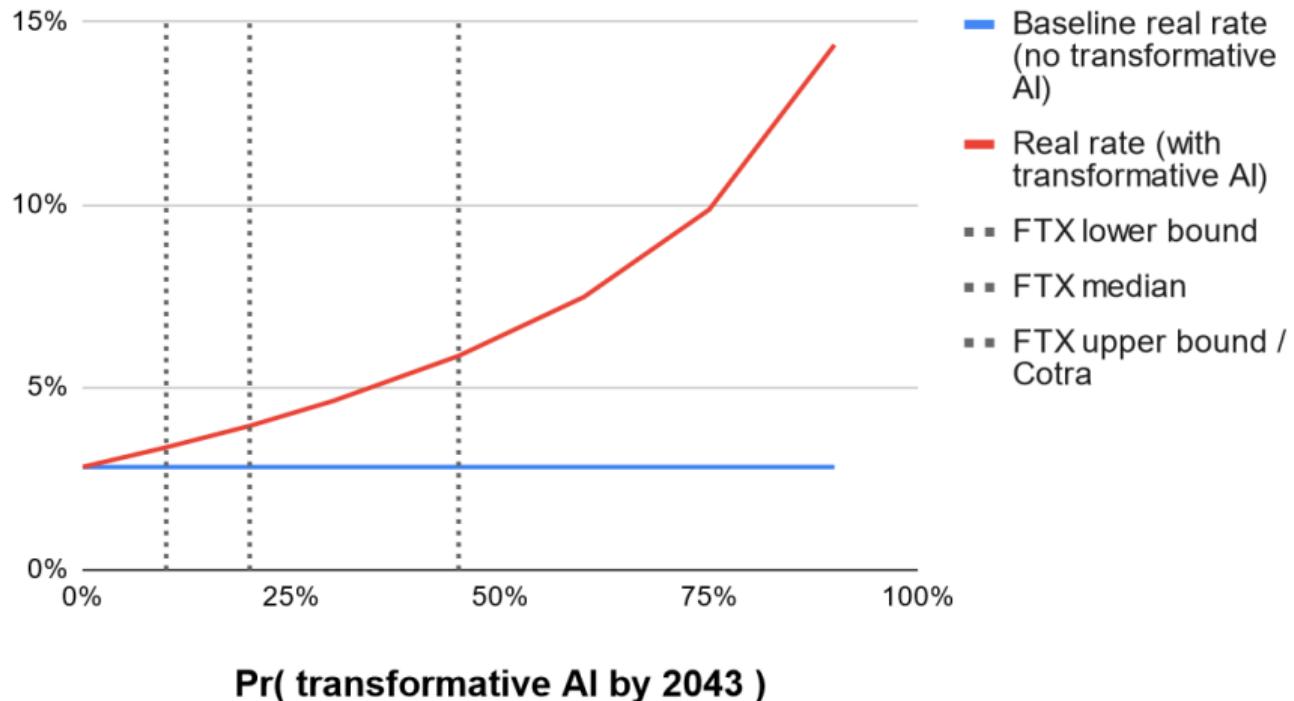
Euler model: varying alignment risk

Real rate model: varying alignment risk



Euler model: varying timelines

Real rate model: varying timelines



Euler model: markets are decisively rejecting AI 2027

1. Five year timelines: With a 50% probability of transformative AI by 2027, and the same yearly probability thereafter, the model predicts 13.0pp higher 30-year real rates today!
2. Ten year timelines: With a 50% probability of transformative AI by 2032, and the same yearly probability thereafter, the model predicts 6.5pp higher 30-year real rates today.