

Transformative AI, existential risk, and real interest rates*

Trevor Chow
Stanford

Basil Halperin
University of Virginia

J. Zachary Mazlish
Oxford and GPI

August 2025

Abstract

We study how financial market prices can be used to forecast the likelihood of transformative artificial intelligence. Transformative AI is a double-edged sword: while advanced AI could lead to rapid economic growth, some researchers argue that superintelligence misaligned with human values could pose an existential risk to humanity. Theoretically, we show that either possibility would predict a large increase in *long-term real interest rates*, due to consumption smoothing. We then use rich cross-country data on real rates and growth expectations to show that, contrary to other recent findings, higher long-term growth expectations are indeed associated with higher long-term real interest rates. We conclude that monitoring real interest rates is a promising framework for forecasting AI timelines.

*tmychow@stanford.edu, basilh@virginia.edu, john.mazlish@economics.ox.ac.uk

Thanks to Seth Benzell, Tamay Besiroglu, Daniele Caratelli, Matt Clancy, Tyler Cowen, Joel Flynn, Caleb Maresca, Jackson Mejia, Laura Nicolae, Jett Pettus, Karthik Tadepalli, and Phil Trammell for feedback; as well as participants at ETGP 2022 and the 13th Oxford Workshop on Global Priorities Research. We are grateful for generous financial support to the George and Obie Shultz Fund at MIT and to the Institute for Humane Studies under grant number 017435.

First version: January 2023

1 Introduction

Background. Recent rapid progress in generative artificial intelligence has highlighted the possibility that humanity may soon develop “transformative AI”: AI technology that precipitates a transition comparable to the agricultural or industrial revolutions. Leading research labs like OpenAI and Google DeepMind bluntly declare their mission to build “artificial general intelligence” (AGI) that can perform at or above human level on all tasks (OpenAI 2023; DeepMind 2023). The possibility of relatively short timelines for AGI is taken seriously by leading machine learning researchers, who in a 2023 survey gave a 10% chance that by 2027 AI will outperform humans at all tasks and a median forecast for such capability by 2047 (Grace, Stewart, et al. 2024).

The prospect of such transformative AI is a “double-edged sword”, in the language of Jones (2023). On the one hand, continued AI innovations like those which have occurred in protein folding or text generation could accelerate economic growth and improve well-being. In the same way that growth increased by roughly an order of magnitude with the industrial revolution, some have predicted that transformative AI automating all tasks would increase growth by another order of magnitude, with GDP growth rising to 30% or more per year (Davidson 2021). Indeed, standard models of economic growth extended to include human-level AI can predict even economic singularities: infinite output in finite time (Aghion, Jones, and Jones 2018; Trammell and Korinek 2020).

On the other hand, many in the AI research community and in the broader public are concerned that such powerful AI technology could create severe risks, even an “existential risk” for the human species. This concern is driven by the challenge of ensuring that smarter-than-human AI technology pursues goals matching human values, rather than pursuing unintended and undesirable goals: the “AI alignment problem” (Ngo 2022; Yudkowsky 2016). The 2023 survey of machine learning researchers found that – among those who chose to respond – the median believed there to be a 5% chance that human-level AI results in “human extinction or similarly permanent and severe disempowerment of the human species” (Grace, Stewart, et al. 2024). This scenario is referred to as *unaligned* AI, in contrast to the growth-enhancing scenario with *aligned* AI.

Most economists, meanwhile, have been notoriously less likely to agree that transformative AI will be developed soon, less optimistic that aligned AI would radically accelerate economic growth, and less pessimistic that unaligned AI could pose an existential risk to human survival, on average (Korinek *forthcoming*).

This paper. We study the implications of transformative AI for asset prices and show how financial market prices can be used to forecast the arrival of such technology. In particular, we show that the prospect of transformative AI would predict a large increase in *long-term real interest rates*, and would do so under expectations of either aligned or unaligned AI. As a result, to the extent that financial markets are efficient information aggregators, the level of long-term real interest rates can be used to help forecast the development of transformative AI.

This predicted rise in real interest rates is a basic implication of all modern asset pricing models, and is simply an application of the logic of consumption smoothing. The key insight is that in both scenarios, future consumption becomes less valuable relative to present consumption. Consider the case of *aligned* transformative AI: growth-induced abundance would lead to low marginal utility of future consumption. Similarly, if the market were forecasting future AI to be *unaligned* and to extinguish humanity, future consumption would have zero value. In either case, it is less valuable to save resources for future consumption, which pushes up interest rates at the relevant horizon.

Empirical results on real rates. We offer new empirical evidence confirming that indeed, in the data, higher long-term growth expectations increase real interest rates. This challenges a recent literature arguing for a weak or nonexistent relationship between real rates and growth (Rogoff, Rossi, and Schmelzing 2024; Lunsford and West 2019; Schmelzing 2019; Hamilton et al. 2016; Borio et al. 2022).

Measuring real interest rates is challenging. Existing work estimates real interest rates by using the nominal yields on nominal bonds and attempting to construct a measure of expected inflation to subtract from the nominal yields. The estimation of expected inflation needed for this, however, is difficult. We tackle this difficulty in two ways.

First, we use real yields from *inflation-linked bonds*, which provide a cleaner direct measurement of real rates compared to using nominal yields with estimated inflation expectations. To our knowledge, prior literature on the topic has not used real rates from inflation-linked bonds only because these bonds are comparatively new, with 20 or 30 years of data available.

Using these real yields directly from inflation-linked bonds, we show that higher real rates today indeed predict higher future GDP growth. Figure 1 shows the correlations for the US, UK, Australia, and Canada at the 10-, 15-, and 20-year horizons, comparing real interest rates over the relevant horizon with future GDP growth at the same horizon. While this data is merely correlational, and the data points are not independent of each other, it is suggestive evidence that growth and real interest rates are significantly linked.

Second, we use rich survey data on long-term inflation expectations from across 59 countries over the last 35 years to construct real interest rates from nominal bonds, together with long-term growth expectations for the same sample. The survey data is a unique dataset of forecasts from professional forecasters collected by Consensus Economics. By using forward-looking forecasts of inflation – rather than backward-looking statistical measures of expected inflation, as in much of the literature – we are able to construct a large panel of real interest rate data. We then conduct a battery of tests and find a strong relationship between long-term real interest rates and long-term growth expectations. To the best of our knowledge, these exercises are the cleanest available evidence on the link between *ex ante* real rates and expected aggregate growth.

Other asset prices. We also briefly discuss the implications of transformative AI for other asset prices. We highlight that the implications of transformative AI for *equity* prices

Real rate vs. future real GDP growth

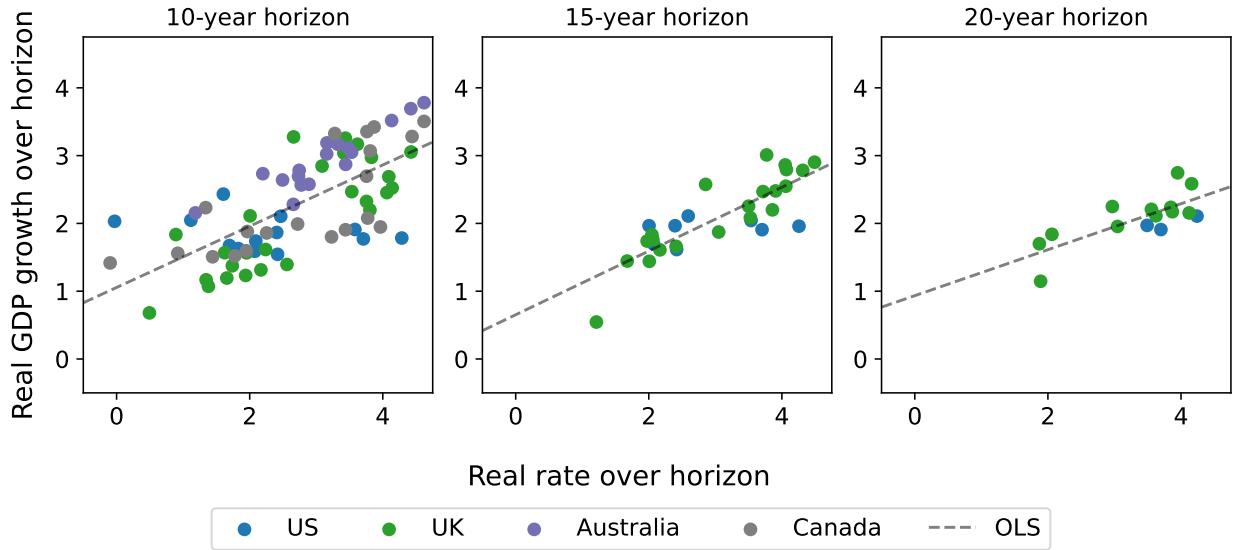


Figure 1: Real interest rates from inflation-linked bonds versus future GDP growth. Each subfigure plots a scatterplot of real interest rates of the titular maturity on the x-axis versus *future* annual GDP growth over the same horizon on the y-axis. Real interest rates are measured using yields on inflation-linked bonds on the last trading day of each year. The scatter plots show all available data up through 2022, for the US (since 1999), the UK (since 1985), Australia (since 1995), and Canada (since 1991), where the end date of the data depends on the time horizon. More details on data sources are given in section 3.

are much more ambiguous than for real interest rates. Among other issues, while the prospect of *aligned* AI leading to rapid growth may increase equity valuations, expectations of *unaligned* AI on the other hand would lower valuations. The net effect is qualitatively ambiguous, making stocks more difficult to use as a barometer for market expectations for AI timelines without an accurate equilibrium asset pricing model. Moreover, even whether higher expected future growth from aligned AI raises or lowers overall equity valuations is itself unclear, and depends critically on the intertemporal elasticity of substitution: larger future cashflows due to economic acceleration may be more than offset by the higher discount rate previously discussed. Even setting these two issues to the side, additionally it is not obvious that AI companies will capture *profits* from developing advanced AI – which is necessary for the expectation of AI to show up in equity prices – or that any companies which do capture profits are currently publicly traded. Finally, we also touch on the implications of transformative AI for land and commodity prices.

Related literature. Two recent papers build on our work. Andrews and Farboudi (2025) perform an event study analysis, examining the behavior of interest rates around 15 AI model releases. They find a large decline – 12 basis points – around each of these model releases on average, an effect that is statistically significant as long as model release tim-

ing is as good as random. They interpret this as evidence that investors do take the possibility of transformative AI seriously, but have updated against its prospect. Maresca (2025) builds an equilibrium model with incomplete markets and idiosyncratic income risk, where labor is fully automated after the development of transformative AI. He also finds, quantitatively, that the prospect of transformative AI pushes up real interest rates substantially.

Looking instead at equity prices, Eisfeldt, Schubert, and Zhang (2023) study the cross-sectional equity price implications of the ChatGPT release, and Korinek (2025) estimates the equity value of OpenAI. An example of using a full equilibrium asset pricing models to forecast technological progress is Ward (2020), who uses equity valuations to forecast the duration of the information technology revolution.

In section 2, we briefly review the large literature outside economics forecasting the pace of AI progress as well as the literature on the economics of AI more broadly. Aside from this and the papers above, the paper is most closely related to the empirical literature measuring the relationship between aggregate growth and real interest rates discussed in 4. Section 5 reviews related literature on the relationship between mortality risk, savings behavior, and real rates.

Outline. The structure of this paper is as follows. In section 2, we define the “transformative AI” scenario under consideration, and provide a brief overview of related forecasting work, which may be less familiar to many readers with a background in economics. We also briefly review the burgeoning literature on the economics of AI. In section 3, we demonstrate the simple result that growth and death risk raise real interest rates in a very broad set of models. Section 4 is our core empirical contribution: we use new data to present evidence that higher growth expectations raise real rates today. We also offer some commentary on existing analysis of this topic. Section 5 reviews relevant literature finding that mortality and savings behavior is related. Section 6 discusses the implications of transformative AI for equity, land, and commodity prices. Section 7 concludes.

2 Defining transformative AI and relevant literature

In this section, we define the “transformative AI” scenario under consideration and provide context on existing research on the topic.

2.1 Defining transformative AI

For the purposes of this paper, we consider the prospect of “transformative AI” as defined informally by Karnofsky (2016): artificial intelligence technology that has at least as profound an impact on the human trajectory as did the industrial revolution or agricultural revolution. As Karnofsky (2016) discusses, this term is similar to other concepts such as “artificial general intelligence” and “superintelligence”, but is intended to be more in-

clusive – capturing technology which is transformative, even if such technology is not able to match all human abilities.

We operationalize this definition of transformative AI by dividing two cases.

Definition (Aligned transformative AI). Aligned transformative AI is technology that causes growth in global GDP in excess of 30% per year.

Definition (Unaligned AI). Unaligned AI is technology that causes the extinction of humanity.

Our definition of aligned transformative AI follows Davidson (2021), who defines “explosive growth” as growth in gross world product of at least 30%, i.e. an increase in growth rates by an order of magnitude.¹ He discusses the possibility that transformative AI could cause such explosive growth. We take this as our benchmark for the effect of aligned AI, though given the unprecedented magnitude under consideration, these numbers clearly should be taken as rough approximations rather than as precise predictions.

A small economics literature has analyzed the economics of transformative AI, and does not reject the possibility that advances in artificial intelligence technology could radically accelerate growth. The seminal contribution to this literature is Aghion, Jones, and Jones (2018), who consider a range of possible scenarios for the effects of artificial intelligence on economic growth. Of particular interest is their result that if AI automates tasks in the *ideas* production function (rather than the goods production function), then growth could accelerate without bound.²

Our definition of the unaligned AI scenario follows the literature on the topic.³ The basic concern is that it may be technically challenging to successfully program artificial intelligence technology in such a way that it behaves in line with human values. Just as software bugs can have large negative consequences in more mundane computer systems, software bugs in very powerful artificial intelligence systems could have correspondingly impactful negative consequences (Yudkowsky 2016; Bostrom 2014; Ngo 2022; Karnofsky 2021). There is limited analysis of the AI alignment problem from an economics perspective. Four exceptions are Hadfield-Menell and Hadfield (2019), Gans (2018), Ely and Szentes (2023), and Chen, Ghersengorin, and Petersen (2024).⁴

¹See also Hanson (2000).

²Clancy (2022) offers a readable summary. Korinek and Stiglitz (2018) analyze how the development of AI could affect the income distribution (see also Korinek 2019; Korinek and Suh 2024). Trammell and Korinek (2020) and Erdil and Besiroglu (2023) review different ways of modeling AI’s role in growth models. Acemoglu (2025) and Erdil, Potlogea, et al. (2025) quantitatively forecast the growth effects of contemporaneous AI capabilities using Hulten’s theorem and an integrated assessment model, respectively.

³Concern over risks from artificial intelligence technology are widespread not just among the public and in fiction, but also among many scientists across many fields. This has recently been captured by the “Statement on AI Risk”, signed by a long list of AI scientists and public figures, stating, “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war” (Center for AI Safety 2023).

⁴A larger set of papers in economics has analyzed how we should think about the tradeoffs between technology which brings positive benefits but creates existential risks (Trammell and Aschenbrenner 2024; Jones 2023; Acemoglu and Lensman 2023; Guerreiro, Rebelo, and Teles 2023; Gans 2024; Beraja and Zorzi 2022; Lehr and Restrepo 2022).

2.2 Forecasting transformative AI

Forecasts of progress in artificial intelligence have a long history. Good (1965) originated the concept of an “intelligence explosion”, a hypothesized phenomenon where AI systems gain the ability to improve their own algorithms and architectures, leading to recursive improvement and rapid increases in intelligence and power. Vinge (1993) originated and Kurzweil (2005) popularized the related concept of a “technological singularity”, referring to an acceleration in technological progress occurring so quickly that it would be difficult to predict *ex ante* how the world would look after. While these earlier analyses were mostly speculative, rapid progress in machine learning over the last decade has resulted in analysis more grounded in the reality of modern AI.

Cotra (2020) provides an influential benchmark forecast for the development of transformative AI. Her framework is based on estimating the number of computations the human brain can perform per second. She then forecasts forward trends in the computational power of computers, using long-run trends like Moore’s Law. She combines these to estimate the date by which computing power could match that of the human brain. Her analysis generates a distribution of estimates, with Cotra (2020) estimating a median arrival date of 2050 for transformative AI, and the updated analysis in Cotra (2022) forecasting a median of 2040. These estimates, however, are highly uncertain: the analysis of Cotra (2020) showed a 10% probability of transformative AI before 2030 and a 20% probability that transformative AI is not developed until after 2100.

Surveys of machine learning researchers are not too far off from the Cotra (2020) estimates. Grace, Salvatier, et al. (2018) survey 352 AI researchers on “when unaided machines can accomplish every task better and more cheaply than human workers” and find a median of 2061. Stein-Perlman, Weinstein-Raun, and Grace (2022) run an updated version of this survey with 738 respondents, and find a median of 2058 for the same question; Grace, Stewart, et al. (2024), in the latest iteration of the same survey with 2,778 published researchers, find a median of 2047. These results again come with significant dispersion and depend on question phrasing (Weinstein-Raun 2024).

Davidson (2023) uses a large-scale semi-endogenous growth model, a la Jones (1995), to forecast timelines for the development of transformative AI, and has a median forecast of 2043 for the development of transformative AI. This approach to forecasting the path of AI is analogous to the “dynamic integrated-climate economy” (DICE) modeling approach used in the climate literature: it is a computational integrated assessment model with an economics foundation.

Economists generally have been more cautious about forecasting the development of transformative AI. Korinek (forthcoming) surveys economists and AI researchers about the probability of the development of “human-level machine intelligence”. The median response of AI researchers in this survey was before 2050; for economists, the median response was after 2070, though results were sensitive to how the question was asked.

3 Real interest rates, growth, and mortality in theory

This section does three things. First, it briefly states the standard Euler equation logic that motivates looking at real rates to predict transformative AI. Second, it motivates the exact specification we take to the data. Third, it describes the modifications to the canonical consumption-savings problem that would be necessary to overturn or weaken the baseline prediction.

3.1 The basic logic

The relationship between real interest rates, growth, and mortality risk follows from standard intertemporal optimization. With time-separable utility over consumption $u(C_t)$, the canonical Euler equation is:

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

where β is the subjective rate of time preference, $\delta \in [0, 1]$ is the survival probability, and r_t is the real interest rate. First, observe that higher death risk δ causes a higher real rate. Second, higher consumption growth, all else equal, also raises the real rate under diminishing marginal utility.

In the case of perfect foresight, to a first-order approximation, this simplifies to the Ramsey rule:

$$r = \rho + \frac{1}{\sigma} g \quad (2)$$

where ρ combines the rate of time preference and survival probability, σ is the elasticity of intertemporal substitution, and g is consumption growth. With σ typically calibrated between 0.2 and 2, transformative AI would raise real rates dramatically. In the benchmark case of log utility ($\sigma = 1$), our definition of aligned transformative AI (30% annual growth) would imply real rates above 30%. For context, 10-year real rates in our developed country sample have never gone above 4%.

The result that real rates rise with both higher mortality and higher consumption growth holds in a wide variety of models.

1. **Preferences.** First, beyond the representative agent model with separable utility described above, the same comparative statics hold under recursive utility (Flynn, Schmidt, and Toda 2023) and under internal habit formation (Bhamra and Uppal 2014; Hamilton et al. 2016; Dennis 2009). Furthermore, in these models, the slope of the relationship between the real rate and growth continues to be determined by the elasticity of intertemporal substitution.
2. **Incomplete markets.** Second, in a baseline incomplete markets model with acyclical income risk, the same relationships hold (Werning 2015), and allowing for cyclical income risk changes the magnitude of the slope but not the sign. Relatedly, the same

comparative statics hold in the overlapping generations model (Baker, De Long, and Krugman 2005; Acemoglu 2009).

3. **Belief heterogeneity.** Third, Buraschi and Whelan (2022) show that under heterogeneous beliefs about future growth, the real rate reflects the *wealth-weighted average* of expectations about future growth, plus an additional “speculative demand” term due to agents betting against one another. This speculative demand pushes the real rate up further if and only if the elasticity of intertemporal substitution is less than than unity. Since our estimates in section 4 suggest an EIS below unity, belief disagreement would push real rates up further. See additionally Xiong and Yan (2010) and Molavi, Tahbaz-Salehi, and Vedolin (2025).

Online appendix C provides more explicit formulas and discussion.

3.2 Distinguishing the short run and long run

With continuous-time data on real interest rates and consumption levels, the Euler equation (1) holds exactly at every instant. In practice, however, economic data is measured at discrete intervals, creating a time aggregation problem that can obscure – or even reverse – the relationship between real rates and growth.

The time aggregation issue arises because nominal rigidities affect the relationship between real rates and growth in the short run, while nominal rigidities dissipate at a long-enough horizon. A substantial body of empirical literature provides evidence that when real interest rates are “too high”, this causes lower growth (Ramey 2016; Nakamura and Steinsson 2018; Bauer and Swanson 2023), supported by a vast theoretical literature on nominal rigidities (e.g. Woodford 2003).

In the long run, on the other hand, nominal rigidities relax as prices and wages have time to adjust, and it is higher consumption growth that causes higher real rates. Thus, in the short run, too-high real rates cause low growth; in the long run, high growth causes high real rates.

To see the problem of time aggregation, consider the following example without any time discounting or death risk. Suppose the level of consumption is initially constant, so the Euler equation implies a real rate of zero. Then, a monetary tightening at time t^* raising the real interest rate would cause the *level* of consumption to drop and the *growth rate* of consumption to rise above zero, in a benchmark model. The Euler equation holds both before and after t^* . However, suppose we only measure at an annual frequency and t^* occurs mid-year. The measured annual growth rate includes *both* the jump down and the subsequent higher growth. If the level drop dominates, we observe high real rates associated with *negative* measured consumption growth: the opposite of the Euler equation’s prediction.

Thus, due to this flipping of signs, empirical work must carefully distinguish between short-run analysis and long-run analysis. This is an important issue with much existing literature, and is a motivation for our focus in our empirical analysis on *long-term* real interest rates. This motivates our empirical approach in section 4 where we focus on

long-term real interest rates and growth. In particular, we focus on the five-year-five-year forward horizon, to strip out short-term confounding factors from nominal rigidities in the five-year window.

3.3 When higher growth does not increase real rates (as much)

Existential risk from *unaligned* AI robustly predicts higher real rates. In the case of *aligned* transformative AI, two mechanisms could break or dampen the prediction that higher growth would push up real rates.

Precautionary savings term. If a higher level of expected growth is associated with a higher volatility of expected growth, this can depress real interest rates. Mathematically, this follows from Jensen's inequality in (1); intuitively, higher risk leads to a higher desire for savings, pushing down real rates. What matters is expected growth in marginal utility, which weights bad states more heavily.

For example, if consumption growth is lognormally distributed with a mean g and variance Var under isoelastic utility, then:

$$r = \rho + \frac{1}{\sigma}g - \frac{1}{2\sigma^2}\text{Var} \quad (3)$$

This motivates our empirical approach in section 4, where in our preferred specification we control for the volatility in expected growth.

Persistent high marginal utility. The natural assumption of diminishing marginal utility of consumption creates a consumption-smoothing motive: if there will be high growth in the future, resources in that future are less valuable, and thus there is less reason to save today. However, consumer preferences may not be so simple, and marginal utility could remain high even in the face of 30% year-on-year growth.

One case is external habit formation ("keeping up with the Joneses"). Consider the extreme case with external habit where utility is determined entirely by the difference between individual consumption and average consumption. In such a world, a rapid acceleration in growth that lifts the consumption of all equally would not lower future marginal utility at all, and would not provide any incentive to save less or borrow more today. The real interest rate then would be unaffected by the prospect of aligned transformative AI. However, to the extent that preferences are not exclusively based on external habit, then rapid growth caused by transformative AI would still raise the real rate.

A second case where marginal utility could remain high is if growth is accompanied by the introduction of new products. Scanlon (2019) and Trammell (2023) both show that the introduction of new goods can keep marginal utility perpetually high, even as consumption grows without bound. Guerrieri et al. (2022) discuss a similar result in the context of temporary product unavailability during the COVID-19 pandemic.

Our empirical results in section 4 provide evidence that in the available historical data, across a wide variety of countries and macroeconomic environments, growth and real rates are robustly related. This suggests that, on average, growth has been associated with diminishing marginal utility of consumption and higher real rates.

4 Empirical evidence on real rates versus growth: r vs. g

In the last section, we presented theoretical motivation for the claim that higher expected growth results in higher interest rates. In this section, we provide empirical evidence that the predicted relationship holds in the available data. We first preview the empirical approach, then discuss our data sources, and then show the results.

4.1 Empirical approach

r vs. g in levels. The baseline specification is a panel regression of the following form:

$$r_{i,[t+5,t+10]} = \alpha + \beta_1 \mathbb{E}_t(g_{i,[t+5,t+10]}) + \varepsilon_{i,t} \quad (4)$$

The dependent variable is the level of country i 's five-to-ten-year real interest rate; the independent variable $\mathbb{E}_t(g_{i,[t+5,t+10]})$ is expected GDP growth from five to ten years ahead.⁵

Our analysis compares the five-to-ten year real interest rate with five-to-ten year expected growth. The motivation for this choice of horizon is that the short-term relationship can be confounded by monetary factors, as discussed in section 3.2. We isolate the relationship between long-term growth and long-term real rates by using the five-to-ten year horizon. Online appendix D shows the full suite of results with 10-year rates and 10-year growth expectations, where the results largely go through.

Again motivated by the discussion in sections 3, we can also consider the regression with a vector of controls X and country fixed effects:

$$r_{i,[t+5,t+10]} = \alpha_i + \beta_1 \mathbb{E}_t(g_{i,[t+5,t+10]}) + \beta_2 X_{i,t} + \eta_{i,t} \quad (5)$$

where we consider three controls:

- (i) Credit default swap (CDS) rates on the country's ten-year debt. Using CDS rates allows us to control for country default risk, an important issue even for advanced economies like the US, as shown by Chernov, Schmid, and Schneider (2020), which many other papers in the literature have neglected.⁶ CDS rates come from either Bloomberg or Longstaff et al. (2011).

⁵The choice to focus on GDP growth rather than consumption growth is for reasons of data availability, and is discussed further below and in appendix A.1.

⁶Due to heterogeneity country-specific recovery rates and liquidity, we simply use CDS as a control, rather than subtracting it from the real rate.

- (ii) The standard deviation of the consensus five-to-ten-year-ahead growth forecast. Controlling for dispersion across forecasts, within a given country, is motivated by equation (3) and the discussion of the term premium.
- (iii) Average expected growth from zero to five years. Controlling for the short-run expected growth is motivated by business cycle considerations discussed above and in section 3.2.

Throughout our analyses, standard errors are Driscoll-Kraay to account for cross-sectional and longitudinal error correlation.

***r* vs. *g* in differences.** We also consider a first-differenced version of the same regression:

$$\Delta r_{i,[t+5,t+10]} = \beta_1 \Delta \mathbb{E}_t(g_{i,[t+5,t+10]}) + \beta_2 \Delta X_{i,t} + \epsilon_{i,t} \quad (6)$$

We use Δ to denote the change in a variable's value across Consensus survey dates and present results below using one-, three-, and five-year changes.⁷

One potential advantage of estimating in changes is that it more directly reflects our paper's question: how is a *change* in growth expectations reflected in changes in real rates? Another potential advantage is that it avoids stationarity concerns. The issue with estimating in changes is that it reduces our sample size and is potentially biased by other sources of noise. For example, short-term liquidity issues in the bond market during times of crisis could cause measured real rates to rise while growth expectations are falling.

There are also tradeoffs to choosing between one-, three-, or five-year windows for differencing. An advantage of differencing with a shorter horizon is a larger sample size. An advantage of looking at longer horizon changes is that such a window is more likely to purge the short-term noise issues just mentioned. We show results for each window.

4.2 Measuring real rates and expected growth

Constructing real rates. We measure ex ante real interest rates as nominal rates less inflation expectations as measured in the Consensus Economics survey. Consensus Economics data covers 59 countries and directly asks professional forecasters for their inflation forecasts.⁸ As such, these survey expectations are a direct measure of inflation expectations at the appropriate horizon, unlike the existing literature – which econometrically estimates inflation expectations – as we return to in section .⁹

⁷As a misspecification test, we can run the regression (6) without imposing an intercept of 0. Across all specifications shown in table 5, the intercept is statistically indistinguishable from zero.

⁸Consensus surveys ask for annual inflation at the one-year horizon through the five-year horizon, and for a single forecast for the average over the five-to-ten year horizon.

⁹The Consensus Economics data has been used to study other topics, and only using a strict subset of both countries and time – for example, Engel and Rogers (2009) – but as far as we know, no other paper using this data has included a similarly large time span and country sample.

Nominal bond data comes primarily Global Financial Data, supplemented for some countries with data from the OECD, Eikon, or Bloomberg. Appendix B describes the precise breakdown.

Consensus surveys of 10-year expectations – for GDP, consumption, and inflation – are conducted twice a year before 2014 and quarterly since then. We work with the unbalanced panel to maximize power; results are robust to using the biannual data throughout. The Driscoll-Kraay standard errors are adjusted for cross-sectional and serial correlation using Driscoll-Kraay (1998) with X lags to account for the mixed-frequency structure

When subtracting inflation expectations from our measures of 10-year nominal rates, the dates of Consensus surveys do not always perfectly align with the dates on which we have 10-year nominal rate data. We always subtract our inflation expectations from the closest possible measured rate, and only keep data points where the gap between survey and rate measure is less than one month.

Expected growth. In addition to the inflation forecasts already discussed, Consensus Economics also asks for GDP growth and consumption growth expectations. For the results we present in the main body of this text, we use GDP growth expectations, rather than consumption growth. This is because the sample of GDP growth expectations is 34% larger than the sample of consumption growth expectations and the two measures track each other closely (0.95 correlation). The fact that the two measures track each other so closely is consistent with a failure of international risk-sharing while an own-country aggregate Euler still holds.

4.3 Real rates and expected growth

Results in levels. Table 1 shows the results of the regression (5) comparing the level of real interest rates versus the expected growth rate. Column 1 shows the relationship absent any controls or fixed effects, while the rest of the table adds these. Figure 2 shows a raw scatterplot of the same data, providing a visualization of the relationship absent any controls or fixed effects. Figure 3 plots country averages to show the relationship solely using across-country variation; figure 4 shows the scatterplot after demeaning by country to show the relationship exclusively using within-country variation.¹⁰

The primary result is that the coefficient on long-term growth expectations is uniformly positive and highly significant, with a magnitude greater than one in all specifications with controls or fixed effects. A coefficient of one would imply that when long-term GDP growth is expected to be one percentage point higher, real rates are correspondingly one percentage point higher. Further, the coefficient on expected growth maps to the inverse of the intertemporal elasticity of substitution. Thus, our preferred estimate – the specification including fixed effects and the three controls motivated by theory – implies

¹⁰Outlier observations where real rates are greater than 10% are removed from figures as they are from our regressions.

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)
Observations	3080	3080	2113	2113
Overall R^2	0.16	0.10	0.39	0.34
Country FE	No	Yes	No	Yes

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

an elasticity of intertemporal substitution of 0.74. This estimate is in line with the micro-level evidence of Crump et al. (2022) and Marenčák and Nghiem (2024), who estimate an elasticity of 0.5 and 0.7 respectively.

All controls are also highly significant. The coefficient on the standard deviation of the long-term growth forecast is statistically indistinguishable from zero but the sign matches the prediction of the model in equation (3), where higher expected consumption volatility pushes down real rates. The coefficient on 0-to-5-year growth expectations is negative, which is consistent with short-run monetary factors, as previously discussed. Finally, the coefficient on the CDS spread implies that a 100 basis point higher CDS rate yields a 20-30 basis point higher real rate. The R^2 values are also meaningfully large. Note that in column (3) we do not use any country fixed effects, but still explain 39% of the variation in ex ante real rates, across the 59 countries and 2113 observations.

In appendix A.2, we run country-by-country regressions rather than estimating as a panel. In the specification with controls, the median coefficient (across 59 countries) on long-term growth is 1.34, in comparison to our point estimate of 1.38 in the panel regression. 83% of individual country regression coefficients are positive. Many individual country samples are quite small, so we do not expect perfectly consistent results. Appendix A.2 presents more details on these results.

Results in differences. Table 5 shows the results of the regression (6) comparing the *change* in real interest rates with the *change* in the expected growth rate during the same period. The first column of the table presents results where independent variables are one-year changes; the second column with three-year changes; and the third column with

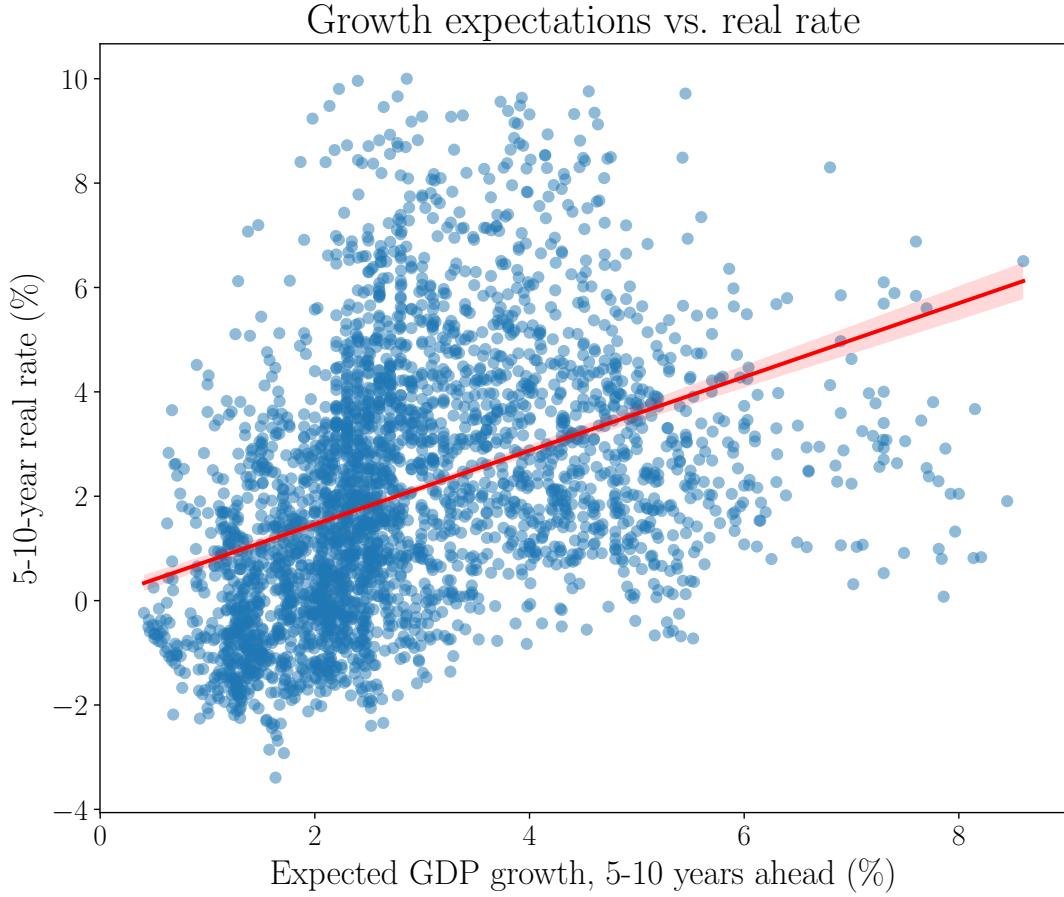


Figure 2: Ex ante real interest rates versus expected GDP growth, both at the 5-to-10-year horizon. Real interest rates are measured using nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics, as is expected GDP growth. Observations are biannual before 2014 and quarterly thereafter. More details on data construction are given in the text.

five-year changes. Figure 5 shows a raw scatterplot of the same data for five-year changes.

The results in changes are noisier than the results in levels, but once again the coefficient on long-term growth expectations is always positive, and significant for three- and five-year changes. The coefficient's magnitude is noticeably but not statistically smaller, and increasing in horizon. Both the three- and five-year change specifications include a point estimate of $\beta_1 = 1$ in their 95% confidence interval. It is not surprising that analysis in changes is noisier: long-term expected growth varies much less than, for example, the CDS spread, and therefore explains less of the variation in real rate changes (see also [Cochrane 2012](#)).

Though not statistically significant, the point estimate for the coefficient on the change in the standard deviation of long-term growth forecasts is always less than 0, consistent

Country averages: Growth expectations vs. real rate

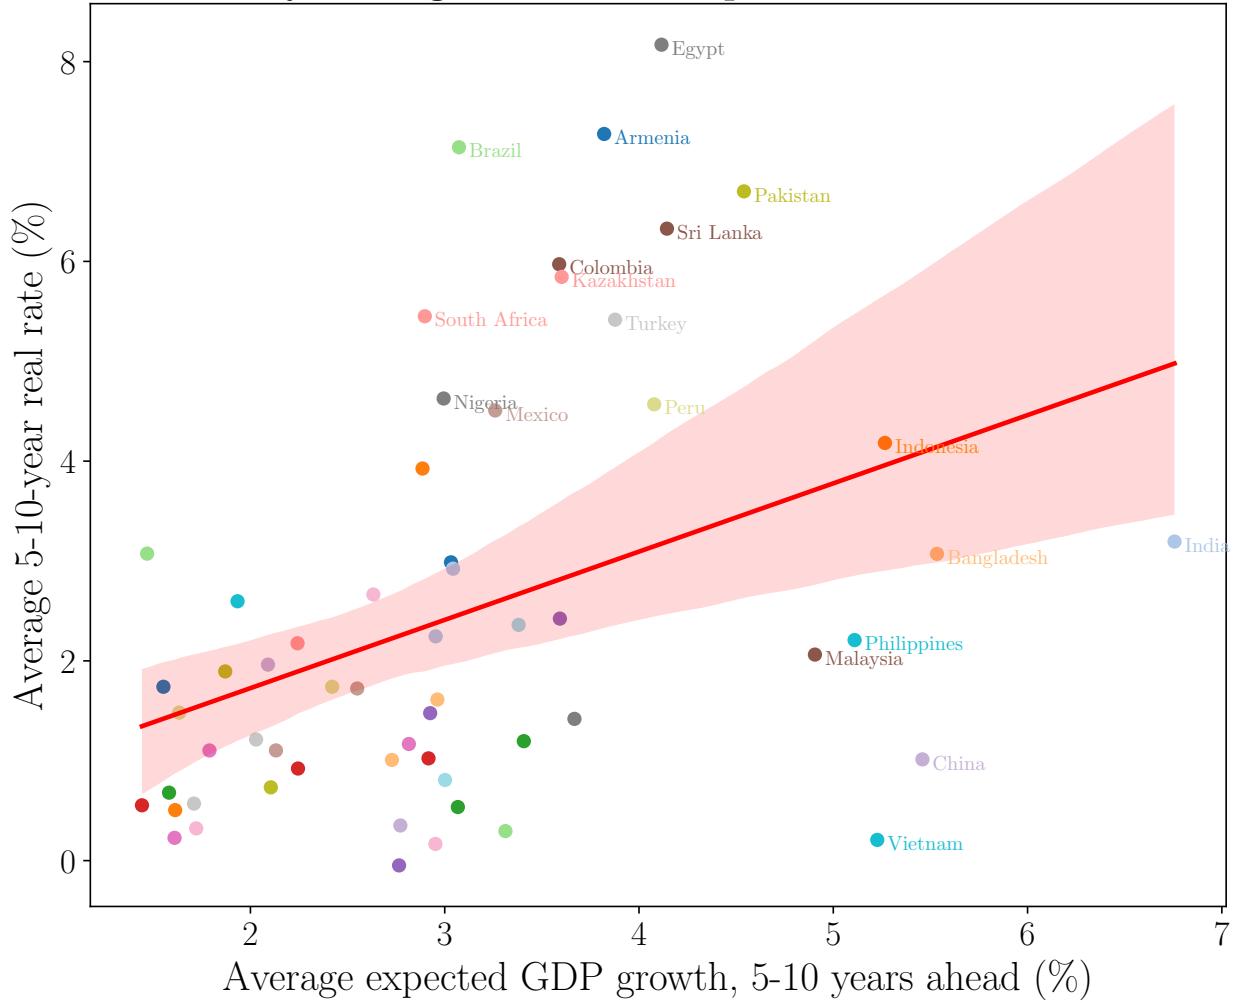


Figure 3: Average by country: ex ante real interest rates versus expected GDP growth, both at the 5-to-10-year horizon. Real interest rates are measured using nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics, as is expected GDP growth. Observations are biannual before 2014 and quarterly thereafter. More details on data construction are given in the text.

with theory. The coefficient on the change in zero-to-five-year growth expectations is again negative, also consistent with theory. We view these results as a potential explanation of the “puzzle” in Duffee (2023) where upward changes in one-year US GDP forecasts (from the Fed’s Greenbook) are associated with downward changes in interest rates. Monetary factors account for this short-term inverse relationship, while traditional consumption smoothing logic dominates on longer-horizons. The coefficient on CDS also remains highly significant, though of smaller magnitude than the regression in levels. Since we are regressing in changes, we do not use country fixed effects, but still achieve

Demeaned growth expectations vs. demeaned real rate

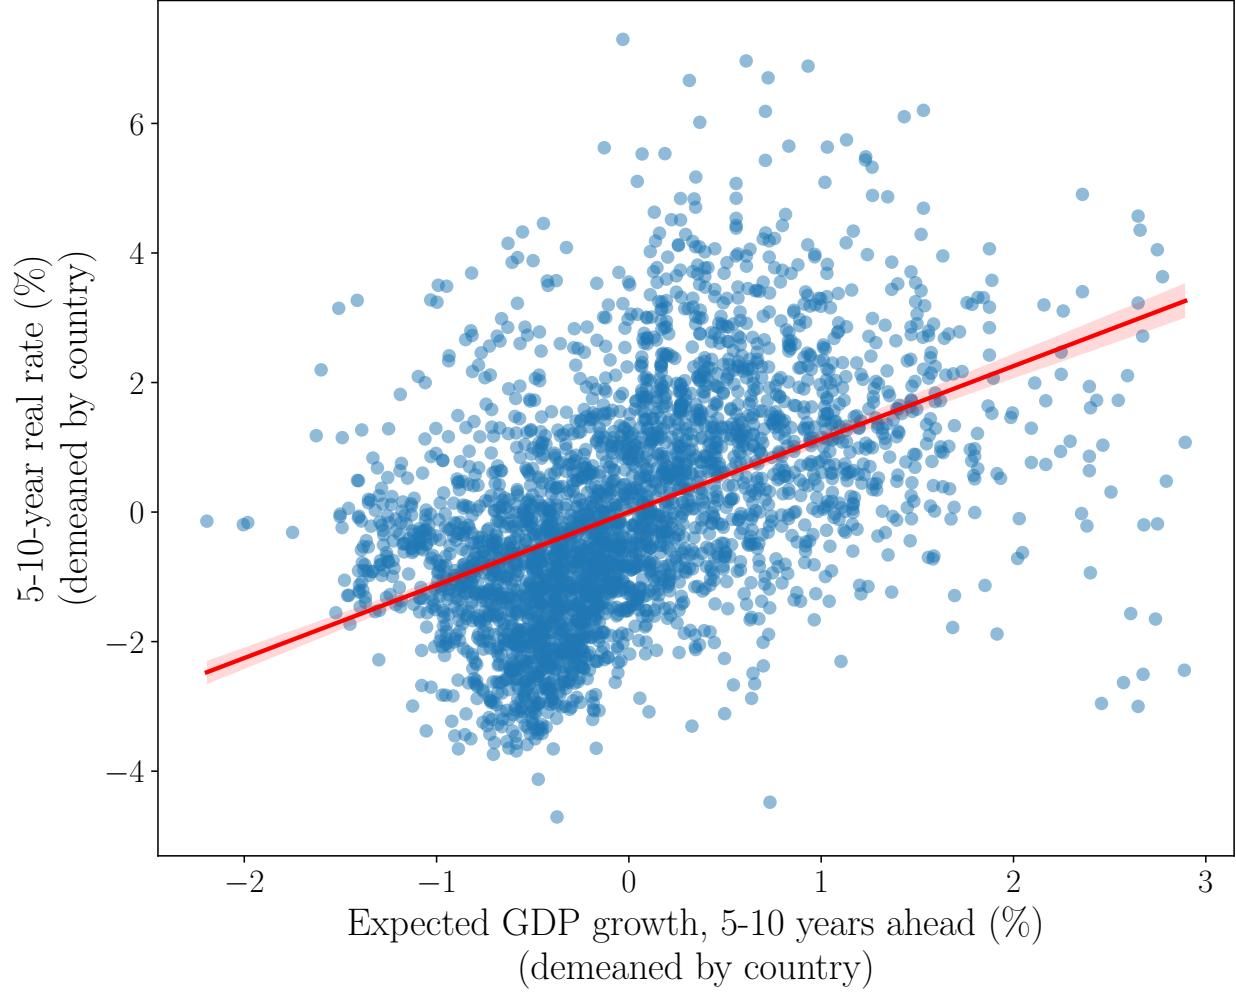


Figure 4: Demeaned by country: ex ante real interest rates versus expected GDP growth, both at the 5-to-10-year horizon. Real interest rates are measured using nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics, as is expected GDP growth. Observations are biannual before 2014 and quarterly thereafter. More details on data construction are given in the text.

meaningfully large R^2 values.

Robustness. Appendix A.3 shows that all results above about the sign and magnitude of the β_1 coefficient are robust to only looking at G7 countries, where the results imply an elasticity of substitution less than 0.5. Appendix A.1 shows that the results hold when using consumption growth expectations instead of GDP growth expectations – in both levels and changes – though with a smaller sample size and lower precision.

Altogether, our results show a clear and reliable connection between higher long-term

Five-year changes:
Growth expectations vs. real rate

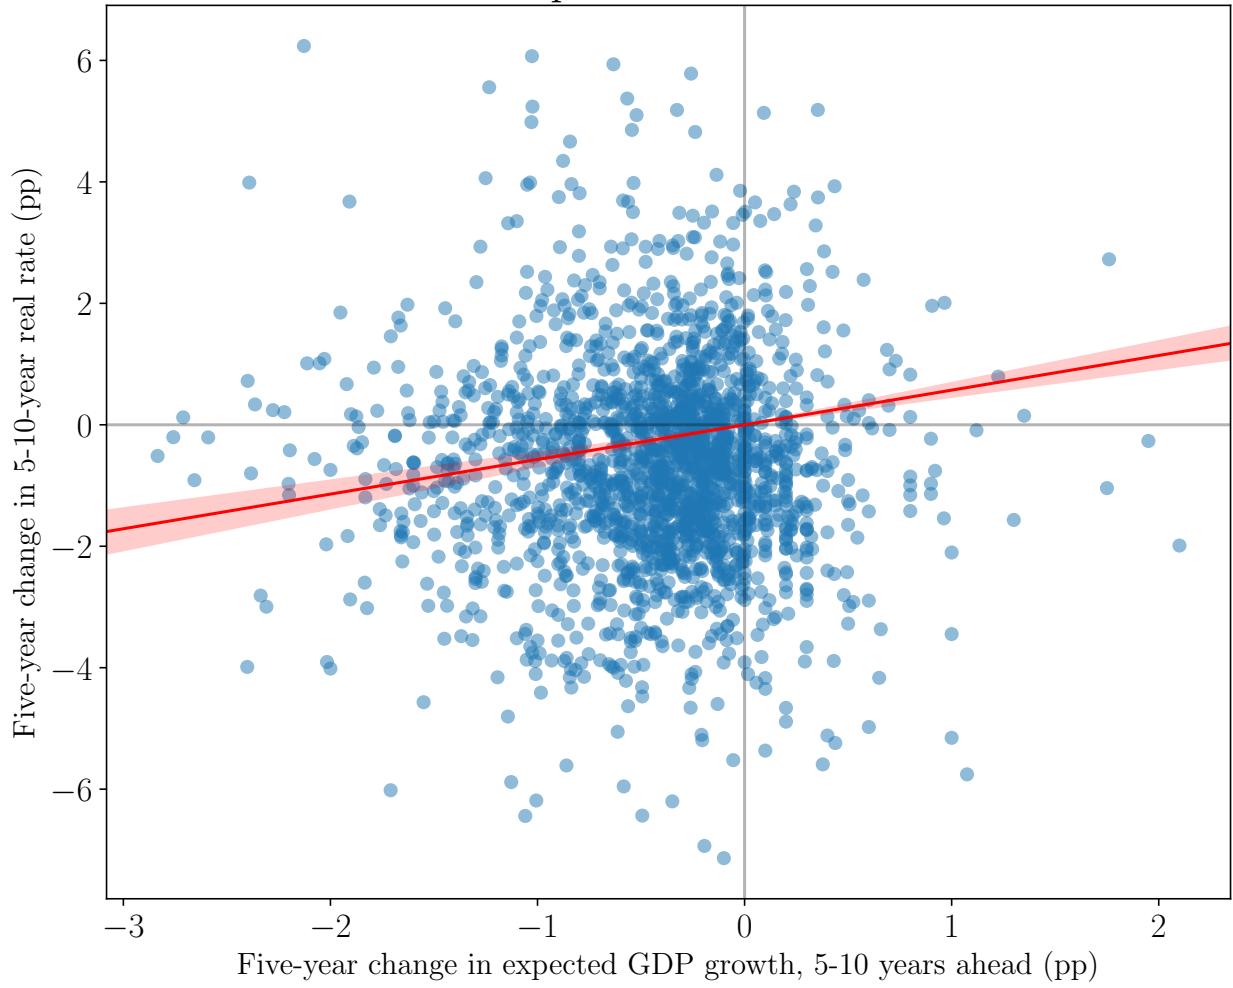


Figure 5: Five-year change in ex ante real interest rate versus five-year change in expected GDP growth, both at the 5-to-10-year horizon. Real interest rates are measured using nominal interest rates minus expected inflation; expected inflation is measured using the consensus of professional forecasters from Consensus Economics, as is expected GDP growth. Observations are biannual before 2014 and quarterly thereafter. More details on data construction are given in the text.

Table 2: Change in expected growth vs. change in real rates

	Dependent variable: $\Delta 5\text{-}10\text{-year real rate}$		
	$\Delta 1\text{yr}$	$\Delta 3\text{yr}$	$\Delta 5\text{yr}$
$\Delta(5\text{-}10\text{-year GDP growth forecast})$	0.21 (0.32)	0.80*** (0.27)	0.99*** (0.34)
$\Delta(\text{SD}(5\text{-}10\text{-year GDP growth forecast}))$	-0.35 (0.25)	-0.31* (0.18)	-0.11 (0.23)
$\Delta(5\text{-year GDP growth forecast})$	-0.19 (0.26)	-0.41* (0.24)	-0.31 (0.22)
$\Delta(\text{CDS spread})$	0.081** (0.041)	0.142** (0.062)	0.730*** (0.120)
Observations	1854	1470	1107
Overall R^2	0.04	0.10	0.32

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

growth expectations and higher long-term real rates. We do not believe such a robust relationship between r and g at the aggregate level has been shown before empirically, and we see our wide cross-country sample as the best available evidence on this foundational macroeconomic relationship.

4.4 Existing literature on r vs. g

A recent literature argues that there is little or no relationship between real interest rates and aggregate growth, in contrast to our findings above. We argue that this is due to data limitations in (1) estimating ex ante inflation expectations and (2) controlling for credit risk.

Estimating ex ante inflation expectations. Measures of historical inflation expectations do not exist for many countries or only have short histories – especially for measures of historical *long-term* inflation expectations. Therefore, most papers in this literature have attempted to construct ex ante inflation expectations using available data, rather than using a direct measure of expectations.

Papers with this approach typically construct inflation expectations using a backward-looking statistical model – often simply lagged, realized inflation. This is the approach used in the careful archival work of Schmelzing (2019) and Rogoff, Rossi, and Schmelzing (2024). Similarly, the analyses of Lunsford and West (2019), Borio et al. (2022), and Hamilton et al. (2016) use rolling AR(1) forecasts based on past inflation as their measure of expected inflation.

However, these approaches are inherently backward-looking and fail to capture the

forward-looking nature of inflation expectations. For example, consider the case of the US at the start of 2023, when inflation was falling rapidly from its highs of the previous year. Under the approach used in Rogoff, Rossi, and Schmelzing (2024), 10-year expected inflation would be calculated as 4.2%.¹¹ This backward-looking approach is quite high, because the 2022 CPI inflation rate was atypically high at 6.4%. However, more direct measurements of inflation expectations at the start of 2023 showed substantially lower inflation expectations: for example, the Survey of Professional Forecasters showed a consensus inflation forecast of 2.4% for the subsequent 10 years. In short, crude backward-looking statistical models often diverge sharply from more direct measures of inflation expectations for time periods when such ground truth is available, especially around turning points.

Credit risk. Additionally, another problem with using historical bonds to measure real rates is credit risk. While modern sovereign bonds from countries like the US are closer to risk-free, this is not the case for all sovereign bonds, and especially was not always so historically. This is relevant, for example, in the long-run historical trends estimated by Schmelzing (2019). He estimates a steady long-run decline in real rates using historical sovereign nominal bonds. Besides also finding an explanation in declining time preference (Clark 2007; Stefanski and Trew 2022), this plausibly reflects a long-run decline in credit risk.

For example, the estimates of Schmelzing show a sharp rise in real rates during the Napoleonic Wars. The estimates for this period come from the yields of British perpetuities, and the United Kingdom is termed a “safe asset provider”. It seems natural to suspect that the measured jump in real rates during this period this reflects, at least in part, heightened credit risk during the conflict, rather than a true increase in risk-free real interest rates.¹² This interpretation is consistent with the post-war normalization of real rate estimates.

4.5 Real rates and *realized* growth

In this subsection, we present some brief evidence showing that real rates and future *realized* growth are also linked, to complement the earlier results linking real rates and future *expected* growth as motivated by the theory. The link between real rates and *realized* growth relies on growth expectations representing rational forecasts. Therefore, given the above evidence that real rates respond to changes in expected growth, the evidence we now provide is evidence that those growth expectations were indeed rational.

As motivating evidence in the introduction, we showed in figure 1 an evident relationship between the real rate today and future realized GDP growth. There, real rates were measured using inflation-linked bonds from the small sample of countries with liquid inflation-linked bond markets.

¹¹This approach uses a seven-year weighted average of lagged inflation, with declining weights.

¹²As well as measurement error in inflation expectations following suspension of the gold standard.

Now, to be consistent with our main analysis, we compare our real rates constructed using Consensus inflation forecasts with future realized growth. In table 3, we run regressions of realized GDP growth five-to-ten years ahead on the five-to-ten-year real interest rate today:

$$g_{i,[t+5,t+10]} = \alpha_i + \beta_3 r_{i,t} + \beta_4 X_{i,t} + \epsilon_{i,t} \quad (7)$$

Row 1 shows that, indeed, higher real rates today are significantly associated with higher realized long-term (5-10 years ahead) GDP growth, whether or not we include our previous batch of controls. The magnitude of the coefficient is smaller than that between real rates and expected growth. Such a difference is consistent with the fact that in a Euler equation framework, this β_3 coefficient is the elasticity of intertemporal substitution, while the coefficient β_1 in equations (4)-(6) is the inverse of the same elasticity. However, we do not put too much emphasis on this interpretation, as the point estimates of the coefficients are by no means inverses of each other, and the fact that we use 5-10 year ahead growth makes this relationship more complicated to disentangle.

Table 3: Real rates vs. realized growth

	<i>Dependent variable: Realized 5-10-year GDP growth</i>	
	(1)	(2)
5-10-year real rate	0.28*** (0.04)	0.22*** (0.05)
SD(5-10-year GDP growth forecast)		0.31* (0.17)
5-year GDP growth forecast		-0.22* (0.12)
CDS spread		-0.208*** (0.052)
Observations	1118	461
Overall R^2	0.00	-0.26
Country FE	Yes	Yes

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

5 Empirical evidence on real rates versus mortality risk

In section 3, we presented theoretical intuition for why higher expected mortality or existential risk would result in higher real interest rates: a heightened probability of death tomorrow would lead agents to want to save less and borrow more today. In this section, we present a review of already-existing work from a disparate set of literatures which provide evidence in support of the theory.

As a preliminary comment, we clarify that we study the relationship between real interest rates and the probability of truly *existential* risks – the probability of human extinction. We contrast this with the large literature on “rare disasters”, which studies events that have a differential impact on risky assets like equities versus on risk-free bonds. “Disaster risk” thus provides a potential explanation for the equity premium puzzle (Rietz 1988; Barro 2006; Gourio 2008; Gabaix 2012). While disaster risk is about events that *differentially* affect stocks versus bonds, existential risk is about events that eliminate agents, thus affecting the return on stocks and bonds equally (and therefore cannot contribute to explaining the equity premium puzzle): existential risk sets the return on both to -100%.^{13,14}

5.1 Mortality risk and savings behavior

In the theory reviewed in section 3, the mechanism by which higher expected mortality risk increased the real interest rate is by reducing savings. With higher probability of nonexistence in the future, agents have lower incentive to save for the future, and this reduced supply of savings increases the real interest rate.

In this subsection we provide evidence on the mechanism: we review existing work showing that reduced mortality risk causally increases savings (or equivalently, increases investment). While this does not provide direct evidence that extinction risk increases real interest rates, it does provide evidence for the hypothesized *mechanism* through extinction risk would increase interest rates.

One example comes from testing for Huntington’s disease, a disease which causes a meaningful drop in life expectancy to around 60 years, in Oster, Shoulson, and Dorsey (2013). Using variation in when people are diagnosed with Huntington’s, the authors find that those who learn they carry the gene for Huntington’s earlier are 30 percentage points less likely to finish college, which is a significant fall in their human capital investment – i.e., savings in the form of human capital investment decrease.

A second example comes from the informational experiment, in Malawi, of Ciancio et al. (2020). The authors provide information to correct pessimistic priors about life expectancy, and find that higher life expectancy directly caused more savings, via investment in agriculture and livestock.

Another set of papers study how the rollout of medical innovations, increasing life expectancy, led to increased savings and investment. Baranov and Kohler (2018) study the

¹³There is also a literature on the relationship between violent, non-existential conflict and asset prices. Hirshleifer, Mai, and Pukthuanthong (2023) use natural language processing techniques to study war discourse in newspaper articles and the relationship with equity prices. Ferguson (2008) and Bialkowski and Ronn (2017) provide narrative evidence of the effect of the world wars on financial markets. Rexer, Kapstein, Rivera, et al. (2022) study the relationship between violent conflict and sovereign nominal bonds. Leigh, Wolfers, and Zitzewitz (2003) as well as Wolfers and Zitzewitz (2009) study the relationship between the Iraq War and financial markets. He et al. (2024) study the impact of political risk in Hong Kong on property markets.

¹⁴Another literature that studies a non-existential disaster risk and financial markets is the climate literature. Giglio, Kelly, and Stroebel (2021) provide a review.

provision of a new AIDS therapy (also in Malawi) which caused a 13-year increase in life expectancy. Using spatial and temporal variation in where and when these therapeutics were rolled out, they find that increased life expectancy results in more financial savings and more human capital investment. Jayachandran and Lleras-Muney (2009) study the sudden drop in maternal mortality in Sri Lanka between 1946 to 1953. They find that for every additional year of life expectancy, years of education increase by 0.11 – i.e., savings in the form of human capital investment increased. Hansen (2013) and Hansen and Strulik (2017) argue that difference-in-difference evidence shows that improvements in antibiotics and cardiovascular disease treatment led to increased human capital investment, with a similar elasticity to the other studies.

Finally, there is suggestive correlational evidence from surveys during the Cold War that a higher perceived risk of nuclear war was associated with a higher savings rate.¹⁵ Russett and Slemrod (1993) find this in a 1990 survey data based on 431 American respondents. Slemrod (1982) as well as Russett, Cowden, et al. (1994) look at the timeseries correlation over the course of the Cold War between the U.S. private savings rate and the average of public opinion surveys on nuclear war risk (as well as the correlation with the Bulletin of Atomic Scientists “doomsday clock”) and find positive correlations. Finally, Slemrod (1990) finds a suggestive negative correlation between the national savings rate and the survey average of perceived nuclear war risk in a cross-section of 19 OECD countries in the 1980s.¹⁶ In contemporary times, Heimer, Myrseth, and Schoenle (2019) find, cross-sectionally in US survey data, that pessimistic survival beliefs are correlated with a lower savings rate. This is true even after controlling for risk preferences, cognitive ability, and socioeconomic factors.

6 Other asset prices

In this section, we consider the possibilities for how transformative AI may affect asset prices other than real interest rates. Our main message is that the sign of the impact on real rates is much clearer than the sign of the impact on other asset prices.

¹⁵We note that even a full-blown nuclear war, while the gravest catastrophe in history, need not be a true *existential* risk in the sense of wiping out the entirety of the human population. Besides a two-sided nuclear exchange possibly being limited, it is still a matter of scientific debate just how much damage such a war and the resulting nuclear winter would cause. Reisner et al. (2018) provides a full-scale analysis; Rodriguez (2019) offers an opinionated summary of the literature. As a result, the literature reviewed here on nuclear war risk is not necessarily comparable to the truly *existential* risk postulated by unaligned AI. It *may* be closer in nature to the “rare disasters” literature mentioned above.

¹⁶There is also work on the relationship between nuclear war risk and *equities*, with particular focus on the Cuban Missile Crisis. Finer (2021) studies the cross-section of US equities during the Cuban Missile Crisis. He compares companies with headquarters that are more or less exposed to Cuban missiles, as assessed by secret (at the time) intelligence assessments. He finds that the more exposed stocks fell by more during the crisis. Burdekin and Siklos (2022) study the Cuban Missile Crisis. They hand collect data on daily equity prices in Canada and Mexico, and together with US data, conclude “markets assigned a very small risk to the crisis leading to the use of nuclear arsenals”. Section 6.1 below discusses transformative AI and equities.

6.1 Transformative AI and equity prices

To forecast AI timelines, it may be tempting to use the market capitalization of companies like Alphabet (owner of DeepMind, a leading AI research lab) or that of chipmakers like Nvidia and TSMC. However, extracting AI-related expectations from stock prices is a challenging exercise for four reasons.

Aligned versus unaligned AI. First, and most importantly, AI-related companies will only have the possibility of high profits if transformative AI is aligned. Under *unaligned* AI where humanity is extinguished, the value of stocks along with everything else is converted to zero.

Profiting versus not. Second, it is not obvious that even in the aligned case that these companies will earn high profits. For instance, OpenAI has committed to a capped profit model, and other AI labs may sign on to a similar ‘Windfall Clause’ promising *ex ante* to donate profits beyond some threshold (OpenAI 2023; O’Keefe et al. 2020). Beyond corporate altruism, it is plausible that if a private company develops truly transformative AI technology, then the local government may nationalize and expropriate it (or at least attempt to do so) to distribute the benefits more broadly, preventing profits Aschenbrenner (2024).

Public versus private companies. Third, when considering equity valuations, there is the question of which stock or stocks to consider. Critically, even if one takes a basket of tech companies and averages over them, then this only includes existing public companies. If the market expects transformative AI very soon, but only because it will be developed by a company which is not traded publicly (e.g. leading labs OpenAI or Anthropic) then this will not necessarily show up in any index of publicly-traded equities, depending on the effect of such technology on the distribution of firm profits.¹⁷

Higher growth may lower stock prices. Fourth, and quite importantly, it is not obvious whether expectations of transformative AI would raise or lower average equity prices. This is because stock prices reflect the present-discounted value of future profits; and transformative AI may raise those future profits, but – as emphasized throughout this paper – transformative AI would also raise the interest rate used to discount those profits. The net effect on average stock prices is ambiguous, without making more assumptions.

In particular, higher growth causes lower average stock prices if the *intertemporal elasticity of substitution* is less than one, rather than greater than one. This parameter – denoted as σ in section 3 – is subject to significant debate. In particular, while macroeconomics papers often calibrate to $\sigma < 1$, typically asset pricing papers calibrate to $\sigma > 1$.

¹⁷For example, the development of the automobile transformed the United States. However, it has been argued that an investor in the year 1900 would have lost money betting on the then-available publicly-traded automobile manufacturers, since these firms went bankrupt (Locke 2021).

For example, Best et al. (2020) use bunching at mortgage notches to estimate $\sigma = 0.1$, and Crump et al. (2022) use directly-measured subjective expectations data to estimate $\sigma = 0.5$. An elasticity greater than unity, on the other hand, would imply that stock prices fall, on average, with news about higher future growth.

6.2 The price of land and commodities

To the extent that advanced AI is able to substitute for labor but not for land or commodities in production, then the value of land and commodities could skyrocket in the case of aligned AI. However, this does require the auxiliary assumption about the shape of the production function – regarding the non-substitutability with land or commodities – which was not needed previously, and which is highly uncertain.

Additionally, again the value of land and commodities are (of course) directly sensitive themselves to real interest rates.¹⁸ This complicates interpretation of their valuation for the same reason as stock valuations.

Finally, the value of land and commodities are hurt by the prospect of *unaligned* AI. As with equities, the net effect of higher valuation from the prospect of aligned AI versus lower valuation from the prospect of human extinction makes the prices of these assets difficult to use for forecasting AI timelines.

7 Conclusion

In this paper, we do not use any detailed inside knowledge of artificial intelligence technology to provide a forecast of the likely timeline for the development of transformative AI. That is, we do not present an ‘inside view’ on AI timelines (Kahneman 2011).

Instead, we argue that market efficiency provides an ‘outside view’ for forecasting AI timelines. The straightforward economic logic of intertemporal optimization, backed up by simple empirical evidence, shows that the prospect of transformative AI would predict high long-term real interest rates. Such rates can be measured using the yields on long-term inflation linked bonds or by subtracting a measure of expected inflation from nominal bonds, and used to inform forecasts of transformative AI.

References

Acemoglu, Daron (2009). *Introduction to Modern Economic Growth*. Princeton University Press.
— (2025). “The simple macroeconomics of AI”. In: *Economic Policy* 40.121, pp. 13–58.

¹⁸Relatedly, Giglio, Maggiori, and Stroebel (2015) estimate 999-year risky, nominal discount rates using features of housing market contracts. See also Andersen (2022), Bracke, Pinchbeck, and Wyatt (2018), Feselmeyer, Liu, and Salvo (2016), and Giglio, Kelly, and Stroebel (2021).

Acemoglu, Daron and Todd Lensman (2023). *Regulating Transformative Technologies*. Tech. rep. National Bureau of Economic Research.

Aghion, Philippe, Benjamin F Jones, and Charles I Jones (2018). "Artificial intelligence and economic growth". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 237–282.

Andersen, Carsten (2022). "Do distant rent flows matter? Inferring discount rates from leasehold apartments in Denmark". In: *Economics Letters* 219, p. 110798.

Andrews, Isaiah and Maryam Farboodi (2025). "Pricing Transformative AI".

Aschenbrenner, Leopold (2024). "Situational awareness". In: *The decade ahead. situational-awareness. ai*.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2018). *The intertemporal keynesian cross*. Tech. rep. National Bureau of Economic Research.

Baker, Dean, J Bradford De Long, and Paul R Krugman (2005). "Asset returns and economic growth". In: *Brookings Papers on Economic Activity* 2005.1, pp. 289–330.

Baranov, Victoria and Hans-Peter Kohler (2018). "The impact of AIDS treatment on savings and human capital investment in Malawi". In: *American Economic Journal: Applied Economics* 10.1, pp. 266–306.

Barro, Robert J (2006). "Rare disasters and asset markets in the twentieth century". In: *The Quarterly Journal of Economics* 121.3, pp. 823–866.

Bauer, Michael D and Eric T Swanson (2023). "An alternative explanation for the "fed information effect"". In: *American Economic Review* 113.3, pp. 664–700.

Beraja, Martin and Nathan Zorzi (2022). *Inefficient automation*. Tech. rep. National Bureau of Economic Research.

Best, Michael Carlos et al. (2020). "Estimating the elasticity of intertemporal substitution using mortgage notches". In: *The Review of Economic Studies* 87.2, pp. 656–690.

Bhamra, Harjoat S and Raman Uppal (2014). "Asset prices with heterogeneity in preferences and beliefs". In: *The Review of Financial Studies* 27.2, pp. 519–580.

Bialkowski, Jędrzej and Ehud I Ronn (2017). "Financial markets facing apocalypse in World War II". In: *The Journal of European Economic History* 46.2, pp. 9–49.

Borio, Claudio et al. (2022). "Why so low for so long? A long-term view of real interest rates". In: *International Journal of Central Banking* 18.3, pp. 47–87.

Bostrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.

Bracke, Philippe, Edward W Pinchbeck, and James Wyatt (2018). "The time value of housing: Historical evidence on discount rates". In: *The Economic Journal* 128.613, pp. 1820–1843.

Buraschi, Andrea and Paul Whelan (2022). "Speculation, sentiment, and interest rates". In: *Management Science* 68.3, pp. 2308–2329.

Burdekin, Richard CK and Pierre L Siklos (2022). "Armageddon and the stock market: US, Canadian and Mexican market responses to the 1962 Cuban missile crisis". In: *The Quarterly Review of Economics and Finance* 84, pp. 112–127.

Center for AI Safety (2023). *Statement on AI Risk*.

Chen, Eric, Alexis Ghersengorin, and Sami Petersen (2024). "Imperfect recall and AI delegation". In.

Chernov, Mikhail, Lukas Schmid, and Andres Schneider (2020). "A macrofinance view of US sovereign CDS premiums". In: *The Journal of Finance* 75.5, pp. 2809–2844.

Ciancio, Alberto et al. (2020). "Mortality risk information, survival expectations and sexual behaviors". In.

Clancy, Matt (2022). *What if we could automate invention? Growth theory in the shadow of artificial general intelligence*. URL: <https://www.newthingsunderthesun.com/pub/2ek4d4s3/release/3>.

Clark, Gregory (2007). *A Farewell to Alms: A brief economic history of the world*. Princeton University Press.

Cochrane, John H (2012). "A brief parable of over-differencing". In: *Manuscript*. Chicago: University of Chicago, Booth School of Business.

Cotra, Ajeya (2020). *Forecasting TAI with biological anchors*. <https://www.lesswrong.com/posts/KrJfoZppSDpnrv9va/draft-report-on-ai-timelines>.

— (2022). *Two-year update on my personal AI timelines*. <https://www.lesswrong.com/posts/AfH2oPHCApdKicM4m/two-year-update-on-my-personal-ai-timelines>.

Crump, Richard K et al. (2022). "Subjective intertemporal substitution". In: *Journal of Monetary Economics* 126, pp. 118–133.

Davidson, Tom (2021). "Could advanced AI drive explosive economic growth". In: *Open Philanthropy*.

— (2023). *What a compute-centric framework says about AI takeoff speeds*. <https://www.lesswrong.com/posts/Gc9FGtdXhK9sCSEYu/what-a-compute-centric-framework-says-about-ai-takeoff>. Accessed on July 7th, 2023.

DeepMind (2023). *About DeepMind*. Accessed: 2023-07-07. URL: <https://www.deepmind.com/about>.

Dennis, Richard (2009). "Consumption habits in a new Keynesian business cycle model". In: *Journal of Money, Credit and Banking* 41.5, pp. 1015–1030.

Duffee, Gregory R (2023). "Macroeconomic news in asset pricing and reality". In: *The Journal of Finance* 78.3, pp. 1499–1543.

Dupraz, Stéphane, Hervé Le Bihan, and Julien Matheron (2022). "Make-up strategies with finite planning horizons but forward-looking asset prices". In.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang (2023). *Generative AI and firm values*. Tech. rep. National Bureau of Economic Research.

Ely, Jeffrey C. and Balazs Szentes (2023). *Natural Selection of Artificial Intelligence*. Working Paper.

Engel, Charles and John H Rogers (2009). "Expected consumption growth from cross-country surveys: implications for assessing international capital markets". In: *IMF Staff Papers* 56.3, pp. 543–573.

Epstein, Larry G and Stanley E Zin (1991). "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis". In: *Journal of political Economy* 99.2, pp. 263–286.

Erdil, Ege and Tamay Besiroglu (2023). "Explosive growth from AI automation: A review of the arguments". In: *arXiv preprint arXiv:2309.11690*.

Erdil, Ege, Andrei Potlogea, et al. (2025). "GATE: An Integrated Assessment Model for AI Automation". In: *arXiv preprint arXiv:2503.04941*.

Ferguson, Niall (2008). "Earning from history?: Financial markets and the approach of world wars". In: *Brookings Papers on Economic Activity* 2008.1, pp. 431–477.

Fesselmeyer, Eric, Haoming Liu, and Alberto Salvo (2016). "How do households discount over centuries? Evidence from Singapore's private housing market". In.

Finer, David Andrew (2021). *No Shock Waves through Wall Street? Market Responses to the Risk of Nuclear War*. The University of Chicago.

Flynn, Joel P, Lawrence DW Schmidt, and Alexis Akira Toda (2023). "Robust comparative statics for the elasticity of intertemporal substitution". In: *Theoretical Economics* 18.1, pp. 231–265.

Gabaix, Xavier (2012). "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance". In: *The Quarterly journal of economics* 127.2, pp. 645–700.

Gans, Joshua S (2018). *Self-regulating artificial general intelligence*. Tech. rep. National Bureau of Economic Research.

— (2024). *How Learning About Harms Impacts the Optimal Rate of Artificial Intelligence Adoption*. Tech. rep. National Bureau of Economic Research.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel (2021). "Climate finance". In: *Annual Review of Financial Economics* 13, pp. 15–36.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel (2015). "Very long-run discount rates". In: *The Quarterly Journal of Economics* 130.1, pp. 1–53.

Good, IJ (1965). "Franz L. Alt and Morris Rubinoff, ed.," Speculations Concerning the First Ultraintelligent Machine". In: *Advances in computers*, pp. 31–88.

Gourio, Francois (2008). "Disasters and recoveries". In: *American Economic Review* 98.2, pp. 68–73.

Grace, Katja, John Salvatier, et al. (2018). "When will AI exceed human performance? Evidence from AI experts". In: *Journal of Artificial Intelligence Research* 62, pp. 729–754.

Grace, Katja, Harlan Stewart, et al. (2024). "Thousands of AI authors on the future of AI". In: *arXiv preprint arXiv:2401.02843*.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles (2023). *Regulating artificial intelligence*. Tech. rep. National Bureau of Economic Research.

Guerrieri, Veronica et al. (2022). "Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages?" In: *American Economic Review* 112.5, pp. 1437–1474.

Hadfield-Menell, Dylan and Gillian K Hadfield (2019). "Incomplete contracting and AI alignment". In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 417–422.

Hamilton, James D et al. (2016). "The equilibrium real funds rate: Past, present, and future". In: *IMF Economic Review* 64, pp. 660–707.

Hansen, Casper Worm (2013). "Life expectancy and human capital: Evidence from the international epidemiological transition". In: *Journal of Health Economics* 32.6, pp. 1142–1152.

Hansen, Casper Worm and Holger Strulik (2017). "Life expectancy and education: evidence from the cardiovascular revolution". In: *Journal of Economic Growth* 22, pp. 421–450.

Hanson, Robin (2000). "Long-term growth as a sequence of exponential modes". In: *Working manuscript*.

He, Zhiguo et al. (2024). "Valuing long-term property rights with anticipated political regime shifts". In: *American Economic Review* 114.9, pp. 2701–2747.

Heimer, Rawley Z, Kristian Ove R Myrseth, and Raphael S Schoenle (2019). "YOLO: Mortality beliefs and household finance puzzles". In: *The Journal of Finance* 74.6, pp. 2957–2996.

Hirshleifer, David, Dat Mai, and Kuntara Pukthuanthong (2023). *War Discourse and the Cross Section of Expected Stock Returns*. Tech. rep. National Bureau of Economic Research.

Jayachandran, Seema and Adriana Lleras-Muney (2009). "Life expectancy and human capital investments: Evidence from maternal mortality declines". In: *The Quarterly Journal of Economics* 124.1, pp. 349–397.

Jones, Chad (2023). "The AI Dilemma: Growth versus Existential Risk". In.

Jones, Charles I (1995). "R & D-based models of economic growth". In: *Journal of political Economy* 103.4, pp. 759–784.

Kahneman, Daniel (2011). *Thinking, fast and slow*. macmillan.

Karnofsky, H (2021). "The “most important century” blog post series". In: *Cold Takes*.

Karnofsky, Holden (2016). *Some Background on Our Views Regarding Advanced Artificial Intelligence*. <https://www.openphilanthropy.org/research/some-background-on-our-views-regarding-advanced-artificial-intelligence/>.

Korinek, Anton (2019). *The rise of artificially intelligent agents*. Tech. rep. working paper, University of Virginia.

— (2025). "Valuing AGI". In.

— (forthcoming). "Economists' Views on the Future of AI". In.

Korinek, Anton and Joseph E Stiglitz (2018). "Artificial intelligence and its implications for income distribution and unemployment". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 349–390.

Korinek, Anton and Donghyun Suh (2024). *Scenarios for the Transition to AGI*. Tech. rep. National Bureau of Economic Research.

Kurzweil, Ray (2005). "The singularity is near". In: *Ethics and emerging technologies*. Springer, pp. 393–406.

Lehr, Nils Haakon and Pascual Restrepo (2022). *Optimal Gradualism*. Tech. rep. National Bureau of Economic Research.

Leigh, Andrew, Justin Wolfers, and Eric Zitzewitz (2003). *What do financial markets think of war in Iraq?*

Locke, Taylor (May 2021). "3 investing lessons Warren Buffett shared at the 2021 Berkshire Hathaway meeting". In: URL: <https://www.cnbc.com/2021/05/03/investing-lessons-from-warren-buffett-at-berkshire-hathaway-meeting.html>.

Longstaff, Francis A et al. (2011). "How sovereign is sovereign credit risk?" In: *American Economic Journal: Macroeconomics* 3.2, pp. 75–103.

Lunsford, Kurt G and Kenneth D West (2019). "Some evidence on secular drivers of US safe real rates". In: *American Economic Journal: Macroeconomics* 11.4, pp. 113–139.

Marenčák, Michal and Giang Nghiem (2024). "Elasticity of intertemporal substitution in the euro area". In: *Available at SSRN* 4913717.

Maresca, Caleb (2025). "Strategic Wealth Accumulation Under Transformative AI Expectations". In.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Andrea Vedolin (2025). "Expectations and the Term Structure of Interest Rates". Working paper, July 2025.

Nakamura, Emi and Jón Steinsson (2018). "Identification in macroeconomics". In: *Journal of Economic Perspectives* 32.3, pp. 59–86.

Ngo, Richard (2022). "The alignment problem from a deep learning perspective". In: *arXiv preprint arXiv:2209.00626*.

O'Keefe, Cullen et al. (2020). "The windfall clause: Distributing the benefits of AI for the common good". In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pp. 327–331.

OpenAI (2023). *About OpenAI*. Accessed: 2023-07-07. URL: <https://openai.com/about/>.

Oster, Emily, Ira Shoulson, and E Ray Dorsey (2013). "Optimal expectations and limited medical testing: Evidence from Huntington disease". In: *American Economic Review* 103.2, pp. 804–830.

Ramey, Valerie A (2016). "Macroeconomic shocks and their propagation". In: *Handbook of macroeconomics* 2, pp. 71–162.

Reisner, Jon et al. (2018). "Climate impact of a regional nuclear weapons exchange: An improved assessment based on detailed source calculations". In: *Journal of Geophysical Research: Atmospheres* 123.5, pp. 2752–2772.

Rexer, Jonah M, Ethan B Kapstein, Andres F Rivera, et al. (2022). *Pricing Conflict Risk: Evidence from Sovereign Bonds*. Tech. rep. Empirical Studies of Conflict Project.

Rietz, Thomas A (1988). "The equity risk premium a solution". In: *Journal of monetary Economics* 22.1, pp. 117–131.

Rodriguez, Luisa (2019). "How many people would be killed as a direct result of a US–Russia nuclear exchange?" In: *Rethink Priorities, June 29*.

Rogoff, Kenneth S, Barbara Rossi, and Paul Schmelzing (2024). "Long-run trends in long-maturity real rates, 1311–2022". In: *American Economic Review* 114.8, pp. 2271–2307.

Russett, Bruce, Jonathan Cowden, et al. (1994). "Did Americans' expectations of nuclear war reduce their savings?" In: *International Studies Quarterly* 38.4, pp. 587–603.

Russett, Bruce and Joel Slemrod (1993). "Diminished expectations of nuclear war and increased personal savings: Evidence from individual survey data". In: *The American Economic Review* 83.4, pp. 1022–1033.

Scanlon, Paul (2019). "New goods and asset prices". In: *Journal of Financial Economics* 132.3, pp. 140–157.

Schmelzing, Paul (2019). "Eight centuries of global real interest rates, RG, and the 'suprasecular' decline, 1311–2018". In: *RG, and the 'Suprasecular' Decline*, pp. 1311–2018.

Slemrod, Joel (1982). *Post-war capital accumulation and the threat of nuclear war*. Tech. rep. National Bureau of Economic Research.

— (1990). "Fear of nuclear war and intercountry differences in the rate of saving". In: *Economic inquiry* 28.4, pp. 647–657.

Stefanski, Radoslaw and Alex Trew (2022). *Selection, Patience, and the Interest Rate*. Adam Smith Business School.

Stein-Perlman, Zach, Benjamin Weinstein-Raun, and Katja Grace (2022). "expert survey on progress in AI". In: *AI Impacts*. Available online at: <https://aiimpacts.org/2022-expert-survey-on-progress-in-ai> (accessed December 7, 2022).

Trammell, Philip (2023). "New Products and Long-term Welfare". In.

Trammell, Philip and Leopold Aschenbrenner (2024). "Existential Risk and Growth". In.

Trammell, Philip and Anton Korinek (2020). "Economic growth under transformative AI". In: *Global Priorities Institute*. <https://globalprioritiesinstitute.org/philip-trammell-and-anton-korinek-economic-growth-under-transformative-ai>.

Vinge, Vernor (1993). "Technological singularity". In: *VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute*, pp. 30–31.

Ward, Colin (2020). "Is the IT revolution over? An asset pricing view". In: *Journal of Monetary Economics* 114, pp. 283–316.

Weil, Philippe (1989). "The equity premium puzzle and the risk-free rate puzzle". In: *Journal of monetary economics* 24.3, pp. 401–421.

Weinstein-Raun, Ben (Dec. 2024). "Reanalyzing the 2023 Expert Survey on Progress in AI". In: URL: <https://blog.aiimpacts.org/p/reanalyzing-the-2023-expert-survey>.

Werning, Iván (2015). *Incomplete markets and aggregate demand*. Tech. rep. National Bureau of Economic Research.

Wolfers, Justin and Eric Zitzewitz (2009). "Using markets to inform policy: The case of the Iraq war". In: *Economica* 76.302, pp. 225–250.

Woodford, Michael (2003). "Interest and prices: Foundations of a theory of monetary policy". In: *Macroeconomic Dynamics* 9.3, pp. 462–468.

Xiong, Wei and Hongjun Yan (2010). "Heterogeneous expectations and bond markets". In: *The Review of Financial Studies* 23.4, pp. 1433–1466.

Yudkowsky, Eliezer (2016). "The AI alignment problem: why it is hard, and where to start". In: *Symbolic Systems Distinguished Speaker* 4.

A Appendix

A.1 Consumption Growth Expectations

Table 4: Expected consumption growth vs. real rate

	<i>Dependent variable: 5-10-year real rate</i>			
	(1)	(2)	(3)	(4)
5-10-year consumption growth forecast	0.66*** (0.09)	0.97*** (0.18)	0.90*** (0.19)	0.87*** (0.24)
SD(5-10-year consumption growth forecast)			-0.69*** (0.15)	0.08 (0.13)
5-year consumption growth forecast			-0.38** (0.16)	-0.15 (0.23)
CDS spread			1.312*** (0.173)	0.869*** (0.163)
Observations	2534	2534	1770	1770
Overall R^2	0.13	0.10	0.50	0.44
Country FE	No	Yes	No	Yes

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 5: Change in expected consumption growth vs. change in real rates

	<i>Dependent variable: Δ5-10-year real rate</i>		
	Δ 1yr	Δ 3yr	Δ 5yr
Δ (5-10-year consumption growth forecast)	0.12 (0.15)	0.41** (0.20)	0.65** (0.26)
Δ (SD(5-10-year consumption growth forecast))	0.00 (0.10)	-0.00 (0.11)	-0.01 (0.16)
Δ (5-year consumption growth forecast)	-0.02 (0.18)	0.00 (0.21)	-0.10 (0.18)
Δ (CDS spread)	0.801*** (0.172)	0.964*** (0.225)	0.799*** (0.147)
Observations	1566	1281	996
Overall R^2	0.16	0.23	0.25

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

A.2 Country-by-country regressions

The first parenthesis reports the percent of coefficients with the “correct” sign; the second reports the percent that are correctly signed and significant.

The “median” columns report median coefficients, observations per country regression, and adjusted R^2 per regression, and similarly the “mean” columns report means.

Table 6: By country: expected growth vs. real rate

	<i>Dependent variable: 5-10-year real rate</i>			
	Median	Mean	Median	Mean
5-10-year GDP growth forecast	1.36 (64%) (81%)	1.54 (64%) (81%)	1.23 (57%) (75%)	1.90 (57%) (75%)
SD(5-10-year GDP growth forecast)			-0.694 (15%) (32%)	-0.923 (15%) (32%)
5-year GDP growth forecast			0.091 (8%) (64%)	0.376 (8%) (64%)
CDS spread			0.620 (47%) (74%)	0.421 (47%) (74%)
Observations	52	52	47	40
Adjusted R^2	0.34	0.34	0.50	0.47

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

In the below table, 58 countries appear in the 1-year change regressions, while 46 countries appear in the 5-year change regressions. The first-two columns report medians and means for 1-year changes; the last-two columns for 5-year changes.

Table 7: By country: change in expected growth vs. change in real rate

	<i>Dependent variable: Δ5-10-year real rate</i>			
	Median Δ_1	Mean Δ_1	Median Δ_5	Mean Δ_5
Δ (5-10-year GDP growth forecast)	0.41 (23%) (66%)	0.71 (23%) (66%)	1.18 (46%) (63%)	1.34 (46%) (63%)
Δ (SD(5-10-year GDP growth forecast))	-0.627 (25%) (72%)	-0.666 (25%) (72%)	0.110 (20%) (46%)	0.797 (20%) (46%)
Δ (5-year GDP growth forecast)	-0.003 (8%) (51%)	0.221 (8%) (51%)	0.012 (24%) (50%)	0.052 (24%) (50%)
Δ (CDS spread)	0.496 (51%) (79%)	0.569 (51%) (79%)	0.472 (50%) (72%)	0.566 (50%) (72%)
Observations	42	35	28	24
Adjusted R^2	0.15	0.20	0.35	0.34

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

A.3 G7 regressions

Table 8: Expected growth vs. real rate (G7 only)

	<i>Dependent variable: 5-10-year real rate</i>			
	(1)	(2)	(3)	(4)
5-10-year GDP growth forecast	2.09*** (0.27)	2.82*** (0.19)	1.33*** (0.29)	2.96*** (0.54)
SD(5-10-year GDP growth forecast)			-1.58 (0.98)	-1.47* (0.80)
5-year GDP growth forecast			-0.40 (0.38)	-0.27 (0.27)
CDS spread			1.109*** (0.219)	0.480*** (0.163)
Observations	591	591	290	290
Overall R^2	0.40	0.35	0.32	-0.37
Country FE	No	Yes	No	Yes

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 9: Change in expected growth vs. change in real rates (G7 only)

	<i>Dependent variable: Δ5-10-year real rate</i>		
	Δ 1yr	Δ 3yr	Δ 5yr
Δ (5-10-year GDP growth forecast)	1.21* (0.65)	1.59*** (0.56)	1.76*** (0.52)
Δ (SD(5-10-year GDP growth forecast))	-0.87** (0.40)	-0.13 (0.83)	-2.06** (0.80)
Δ (5-year GDP growth forecast)	-0.11 (0.29)	-0.16 (0.23)	0.10 (0.21)
Δ (CDS spread)	0.638*** (0.244)	0.656** (0.304)	0.535*** (0.170)
Observations	251	194	146
Overall R^2	0.13	0.17	0.28

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Online Appendix: For Online Publication Only

B Details on data construction

To construct five-year-five-year forward nominal interest rates, we use nominal five- and ten-year sovereign yields generally taken from Global Financial Data (GFD) at a daily frequency. Real rates are the five-year-five-year forward minus the corresponding five-year-five-year Consensus Economics inflation forecast, matched to the first GFD quote within 30 days *after* each Consensus survey.

For some countries, we take nominal interest rate data from an alternate, richer source. Table 10 lists every country for which either maturity comes from a different source.

Table 10: Countries with non-GFD bond-yield sources

Country	Source(s)
Chile	10-year: OECD LTIR (monthly); 5-year: GFD
Colombia	10-year: OECD LTIR (monthly); 5-year: GFD
United Kingdom	5- and 10-year real yields: Bank of England fitted real yield curve ¹⁹
Brazil	5- and 10-year: LSEG Refinitiv Eikon (monthly)
China	5- and 10-year: Bloomberg (monthly)
Poland	5-year: Bloomberg; 10-year: GFD
Slovakia	5-year: Eikon (preferred) + Bloomberg; 10-year: GFD
Taiwan	5- and 10-year: Eikon
Georgia	5- and 10-year: Eikon
Kazakhstan	5- and 10-year: Eikon
Indonesia	5- and 10-year: Eikon
Nigeria	5- and 10-year: Eikon
Russia	5- and 10-year: Eikon
Serbia	5- and 10-year: Eikon
Peru	5-year: Eikon; 10-year: GFD

¹⁹Downloaded from <https://www.bankofengland.co.uk/statistics/yield-curves>. For this data, we do not subtract Consensus expected inflation, as it directly measures a real interest rate.

C Real rates, growth, and mortality in theory: detailed exposition

In this appendix, we demonstrate that real interest rates are connected to both expected economic growth and mortality across a broad range of modeling frameworks. The connection is driven by the same, simple economic logic across all modeling frameworks: consumption smoothing. We show this logic holds in the three classes of models, covering the modern asset pricing modeling frameworks:

- (i) Representative agent models
- (ii) Incomplete markets models
- (iii) Overlapping generations models

Since these results are known in the literature, we focus on results and intuition, and refer interested readers to relevant papers for full derivations. We also consider extensions with behavioral frictions (rule-of-thumb behavior and household myopia); nonstandard preferences (recursive preferences and habit formation); and models with new goods, which create nonstationary utility functions.

We conclude the section by emphasizing how the relationship between real rates and growth depends critically on the time horizon. In the short run – at business cycle horizons – *real interest rates that are too high cause low growth*, due to nominal rigidities. In the long run, nominal rigidities fade, and *high growth causes high real rates*. Hence, our empirical analyses in sections 4 and 5 focus on long-term real rates.

C.1 Representative agent models

It is well-known that in the canonical infinitely-lived representative agent model that the real interest rate is closely tied to growth and death.

The Ramsey rule. In the deterministic case with time-separable utility over the level of consumption, the real interest rate has a particularly simple expression – the canonical “Ramsey rule”:

$$r = \rho + \frac{1}{\sigma} g \tag{8}$$

Here, r is the real interest rate over some time horizon, $\rho > 0$ is the rate of pure time preference, $\sigma > 0$ is the elasticity of intertemporal substitution, and g is the growth rate of consumption. With σ usually calibrated somewhere between 0.2 and 2 – an issue to which we return – we see that higher growth implies a higher real rate. Higher existential risk shows up here as a higher rate of time discounting ρ , thus also implying a higher real rate.

Benchmark calibration under transformative AI. Consider briefly a benchmark calibration at the annual frequency with $\sigma = 1$ (log utility) and $\rho = 0.01$ for a back-of-the-envelope calculation. Then, a growth rate of 1% per capita would imply a real interest rate of 2% under the Ramsey rule – not far off the level seen in advanced economies today. Meanwhile, a transformative AI-induced growth explosion causing the growth rate g to rise to 30% (as defined in section 2.1) would raise real interest rates to 31%. This would be an unprecedentedly high level.

The Euler equation. The Ramsey rule analysis above importantly assumed away uncertainty, among other things. Consider an infinitely-lived household with expected, discounted, time-separable utility over the level of consumption. Denote the period utility function $u(C_t)$ where C_t is consumption and u has diminishing marginal utility, $u' > 0$ and $u'' < 0$. Denote the subjective discount rate as β , and the probability of death in period t as δ . In this representative agent framework, the probability of death δ is equivalent to the probability of extinction.

The resulting intertemporal optimality condition is the well-known Euler equation:

$$1 = \beta\delta\mathbb{E}_t \left[\frac{u'(C_{t+1})}{u'(C_t)} \right] (1 + r_t) \quad (9)$$

Suppose the path of consumption does not adjust, as in an endowment economy, for simplicity.

First, observe that higher death risk causes a higher real rate. A higher death risk is a lower probability of surviving until the next period, δ . A lower δ in (9) requires a higher real rate r_t . The intuition is that a higher probability of death shifts in the willingness to supply savings.

Second, observe that higher consumption growth, all else equal, also raises the real rate. Consider a shock that increases next-period consumption in at least one state of the world and shrinks it in none of them. Then, due to diminishing marginal utility, expected marginal utility $\mathbb{E}_t[u'(C_{t+1})]$ decreases, requiring a higher real rate r_t .

A risky shock – one which increases next-period consumption in some states but lowers it in others – does not unambiguously increase the real interest rate. The Euler equation (9) shows that what matters is *expected growth in marginal utility* – i.e. growth expectations taken over the risk-neutral measure. Thanks to diminishing marginal utility, this means that low-growth states of the world are weighted more highly.

For example, if consumption growth is lognormally distributed around the mean g and variance Var , then:

$$r = \rho + \frac{1}{\sigma}g - \frac{1}{2\sigma^2}\text{Var} \quad (10)$$

Here, $\rho \equiv -\ln(\beta\delta)$ and σ is once again the intertemporal elasticity of substitution.

The fact that the real rate is decreasing in the variance term shows how a shock which increases expected consumption growth could still push down the real rate if the shock

also increases the variance of growth sufficiently. This again is due to the fact that what matters is expected growth in marginal utility. Diminishing marginal utility ensures that low-growth states of the world are weighted more highly.

C.2 Incomplete markets models and heterogeneous agents

The analysis above of the representative agent model demonstrates the importance of savings and borrowing decisions for understanding the effect of expected growth and mortality on real interest rates. This suggests it is worth considering how including realistic borrowing frictions affects the analysis.

Werning (2015) provides a benchmark analysis. In a world where idiosyncratic income risk does not covary with aggregate output, assuming isoelastic utility, and taking the “zero-liquidity limit” so that all agents are hand-to-mouth, then the *slope* of the relationship between growth and the real interest rate is the same, but the *level* is lowered. The analog to the Ramsey equation (8) is:

$$r = \rho + \frac{1}{\sigma} g - \gamma_1 \quad (11)$$

All the terms are as before, with the addition of $-\gamma_1$. The term $\gamma_1 > 0$ reflects the idiosyncratic risk facing the “marginal saver”, which is the agent who *most* wants to save. The slope of the relationship between real rates and growth is still governed by the inverse of the intertemporal elasticity of substitution. Thus the real rate still increases with growth g and the existential risk probability embedded in ρ .

Moving away from the Werning (2015) benchmark, if idiosyncratic risk does covary with aggregate output, the relationship is more complicated. Auclert, Rognlie, and Straub (2018) show that an analog of the Ramsey equation can be written, for a particular form of idiosyncratic risk, as:

$$r_t = \rho + \frac{1}{\sigma \gamma_2} [C_{t+1} - \gamma_3 C_t] - \gamma_1 \quad (12)$$

Recall C_t is aggregate consumption at time t . If $\gamma_2 = \gamma_3 = 1$, then (12) is the same as (11), but they diverge due to cyclical income risk. γ_2 is the cyclicity of income of the marginal saver: the elasticity of individual income to aggregate income. γ_3 is the ratio of average cyclicity of income across all types, relative to the cyclicity of the marginal saver. Holding all else equal, higher C_{t+1} unambiguously increases the real rate.

Summarizing. While the relationship between output or consumption growth and real rates in these models is more complicated, a positive shock to growth still causes higher real rates. An increase in mortality risk has the same effect on the real rate as previously.

C.3 Overlapping generations models

The overlapping generations (OLG) framework is closely related to the incomplete markets framework of the prior section. Consider a simple version of this framework, where each agent lives for two periods and has log utility. There is Cobb-Douglas production technology with capital share α , population growth of n , and exogenous Hicks-neutral productivity growth of g . Then it can be shown that the analog of the Ramsey rule is:

$$r = \rho + g + \gamma_4 \quad (13)$$

Here, the coefficient on growth is 1, since log utility implies that the elasticity of intertemporal substitution is 1. The new term is γ_4 , which is a function of the capital share α and population growth n .²⁰ Once again, the slope of the relationship between the real interest rate and growth is governed by the intertemporal elasticity of substitution; and the relationship between the real rate and mortality risk is direct.²¹

C.4 Recursive preferences and habit formation

Flynn, Schmidt, and Toda (2023) study the relationship between consumption growth and real interest rates under recursive preferences, such as the form studied in Epstein and Zin (1991) and Weil (1989). They show that the relationship is again determined by the elasticity of intertemporal substitution, where this elasticity must be defined appropriately given the recursive nature of preferences. The relationship between real rates and existential risk is unaffected by recursive preferences.

Bhamra and Uppal (2014), Hamilton et al. (2016), and Dennis (2009) study the relationship between consumption growth and real rates under habit formation. They show that with internal habits, the real rate is increasing in consumption growth. On the other hand, consider the extreme case with external habit where utility is determined entirely by the difference between individual consumption and average consumption. In such a world, a rapid acceleration in growth that lifts the consumption of all equally would not lower future marginal utility at all, and would not provide any incentive to save less or borrow more today. The real interest rate would be unaffected by the prospect of *aligned* transformative AI under this assumption, though it would still rise under the prospect of misaligned, extinction-causing AI. However, to the extent that preferences are not purely based on external habit, then rapid growth caused by transformative AI would still raise the real rate. This discussion emphasizes the importance of whether transformative AI will decrease marginal utility, rather than growth rates per se.

²⁰Population growth does not affect the real rate in the canonical representative agent model, unlike the OLG model. Baker, De Long, and Krugman (2005) discuss how under imperfect altruism, population growth increases the real rate even in the representative agent model.

²¹In general without more assumptions than we have made so far, the OLG framework can lead to multiple or degenerate equilibria (see Acemoglu 2009, ch. 9).

C.5 Myopic consumers

If all agents in the economy are fully myopic and do not recognize an impending acceleration in growth or extinction event, then real interest rates are unaffected by such prospects. However, even if *consumers* are fully myopic, as long as *financial markets* can foresee these events, then these prospects will be priced in to real interest rates. Dupraz, Le Bihan, and Matheron (2022) consider a model where consumers are myopic but financial markets are fully forward-looking.

C.6 New goods

Scanlon (2019) and Trammell (2023) both show that the introduction of new goods can keep marginal utility perpetually high, even as consumption grows without bound. In this case, there would not be any incentive to save less or borrow more today in response to higher expected growth. The real interest rate would be unaffected by the prospect of *aligned* transformative AI under this assumption, though it would still rise under the prospect of misaligned, extinction-causing AI.

C.7 Summarizing

Real rates and growth. The common thread across models is: if growth lowers the future marginal utility of consumption, then growth increases real interest rates today. We showed that this holds broadly across models, and highlighted two ways in which it might not. First, for a shock which increases consumption in some states of the world but lowers it in others, real interest rates could fall depending on how these net out. Second, marginal utility could stay high even with rapid consumption growth if utility is a function of relative consumption (i.e. external habit) or if the introduction of new goods keeps marginal utility high.

Real rates and mortality risk. Across all models, higher expected mortality risk raises real interest rates.