Sovereign default and the decline in interest rates*

Max Miller[†] James D. Paron[‡] Jessica A. Wachter[§]

April 28, 2021

Abstract

Sovereign debt yields have declined dramatically over the last half-century. Standard explanations, including aging populations and increases in asset demand from abroad, encounter difficulties when confronted with the full range of evidence. We propose an explanation based on a decline in inflation and default risk, which we argue is more consistent with the long-run nature of the interest rate decline. We show that a model with investment, inventory storage, and sovereign default captures the decline in interest rates, the stability of equity valuation ratios, and the recent reduction in investment and output growth coinciding with the binding zero lower bound.

Keywords: Savings glut, Inflation expectations, Rare disasters, Secular stagnation

JEL codes: E31, E43, G12

---

*We thank Cristina Arellano, Jules van Binsbergen, Sylvain Catherine, John Cochrane, François Gourio, Narayana Kocherlakota, Karen Lewis, Carolin Pflueger, Jonathan Wachter, and seminar participants at the NBER Summer Institute Capital Markets Meeting, at the Conference on Macroeconomics and Monetary Policy at the San Francisco Federal Reserve, and at the Wharton School for helpful comments.

[†]Department of Finance, The Wharton School, University of Pennsylvania. Email: maxmil@wharton.upenn.edu.

[‡]Department of Finance, The Wharton School, University of Pennsylvania. Email: jparon@wharton.upenn.edu.

[§]Department of Finance, The Wharton School, University of Pennsylvania and NBER. Email: jwachter@wharton.upenn.edu.

1 Introduction

Over the last three decades, interest rates across the developed world have starkly declined. Low interest rates, together with low output growth following the Great Recession of 2009, has evoked, for some, the possibility of "secular stagnation," a term coined by Hansen (1939) to describe a persistent period of low investment, employment, and growth. Summers (2015) and Gordon (2015) argue for the relevance of Hansen's concept from two angles: demand-side—an increase in demand for savings arising from changing demographics or growing inequality—and supply-side—arising from a decline in the ideas and dynamism that have fueled the economic growth of the last half-century. A complementary idea is that of a "global savings glut" (Bernanke, 2005): there is too great a supply of savings, mainly from patient investors outside the United States, compared to demand arising from the need to fund productive activities (ideas for which may be lacking).

However, is a greater desire for savings, in fact, what lies behind the decline in interest rates? On some level, the link appears too obvious to be worth questioning. Yet any explanation based on a greater desire for savings runs into a significant problem when one also considers evidence from the U.S. stock market. A greater desire for savings should have pushed up stock prices to a similar degree as bond prices, but it did not. From the point of view of the literature on increased desire for savings, low interest rates, and low growth, the behavior of the aggregate stock market is a puzzle. For a careful explanation of this puzzle, one can look to Farhi and Gourio (2018), who jointly consider growth, interest rates, and stock valuations in a neoclassical growth model that allows for rare disasters of the type considered by Gourio (2012). Farhi and Gourio (2018) show that a significant increase in the risk of rare disasters is necessary to jointly reconcile the level of interest rates and stock prices.

While Farhi and Gourio (2018) succeed in jointly explaining stock prices and interest rates, an explanation based on increased fears of a disaster runs into its own problems. First, the nature of rare disasters means that increased disaster fears are hard to falsify, but one would expect to find evidence in option prices. Yet evidence from options suggests remarkable stability in fears of rare disasters. Second, implications of increased risk of disasters are fragile in that they depend directionally on whether the elasticity of intertemporal substitution (EIS) is above or below one. If the EIS is below, rather than above one, increased risk of rare disasters, together with low growth, require that agents become less patient, not more. Yet the arguments in Summers (2015) and Caballero et al. (2008) point unambiguously toward an increase in patience.

We therefore propose a different explanation, one based on a decline in the risk of inflation. There is substantial evidence for a steady decline in inflation expectations, spanning the 30 years over which interest rates have declined. More recently, evidence from options markets suggests that inflation expectations have become "anchored"—that is, investors do not fear either very high or very low inflation (Reis, 2020).[1] When one takes this evidence into account, it is not difficult to jointly explain the decline in interest rates and the stability of stock valuation ratios. Because the true real rate has not declined, valuation ratios are unchanged, and there is no need to assume a large increase in the probability of a rare disaster to explain the evidence.

One may wonder: if it is simply inflation expectations that have declined, why is it that *measured* real rates, namely nominal rates minus ex post realized inflation, have also declined? But this apparent disconnect disappears if one accounts for inflation *risk*. Indeed, if inflation were perfectly forecastable, then

[1]Further, there is evidence to suggest that the sign of the correlation of inflation on the output gap has changed since 2001, as noted by Campbell et al. (2020).

a change in inflation expectations should not impact ex post real rates. But historically, inflation has on occasion come as a surprise. A decline in inflation risk will lead investors to require a lower premium to hold nominal securities. Interest rates will decline if this premium declines, even if measured in real terms ex post. This effect is more pronounced if investors fear inflation that, in sample, does not occur. From the point of view of cash flows, and given that the sovereign has control over the money supply, inflation risk is essentially risk of default.[2] A decline in inflation risk is thus a decline in the probability of default, and thus should also be expected to affect rates on securities that are said to be inflation-protected. The first contribution of our paper is to show that a model with rare disasters and a decline in inflation risk can explain the decline in interest rates and the stability of valuation ratios. Because sovereign risk depends on institutions that have altered substantially over the centuries, this explanation could account for the striking fact that current rates are low, not just relative to the last 30 years, but to the last 300 years.

We also show that fear of rare disasters, together with low inflation risk, leads to nominal rates that are at or below zero without the need to assume an increased desire for savings. In practice, the existence of cash creates an effective lower bound on interest rates. Such a lower bound is absent in traditional asset pricing models. Thus a second contribution of our paper is to augment a traditional asset pricing model with cash. We introduce cash in a way that does not require any change to preferences, or demand for liquidity. In our model, cash is a storage technology (inventory).[3] When interest rates are sufficiently low, agents have an

[2] We say "essentially" as there is a literature that examines the possibility of outright U.S. government default. Credit default swaps (CDS) traded on U.S. government debt imply a default probability of 0.2%, a phenomenon examined in depth by Chernov et al. (2020).

[3] The theory of cash as inventory dates to Baumol (1952), who applies an inventory control analysis to the theory of money.

incentive to hold cash, which becomes a positive-net-supply asset.

When we consider storage of output, together with productive technologies in a general equilibrium model, we can jointly match low growth, as well as flat valuation ratios and low interest rates. Including inventory in this framework allows us to obtain dynamics even within a framework with independent and identically distributed shocks.

The model with inventory also allows us to match an additional puzzle: the decline in the investment-capital ratio over the last 40 years. In standard models with production, if interest rates are low due to an increased demand for savings, investment relative to capital should rise, and can only be offset by lower productivity or high levels of risk (Farhi and Gourio, 2018). Neither mechanism is needed in a model with riskfree storage of consumption goods. In low-interest-rate regimes, resources that would have been spent on capital are endogenously funneled into non-productive inventory, crowding out private investment.

The crowding out mechanism comes into play when expected returns on the risky asset do not rise fast enough as real rates fall; in the absence of sufficiently attractive investment opportunities, investors look to hoard funds through unproductive avenues. Our approach provides a new mechanism through which low growth can be compounded by low interest rates in a general equilibrium environment.

The remainder of this paper is organized as follows. In Section 2, we briefly summarize the empirical evidence. Section 3 considers the ability of an endowment economy to match this evidence, either with changes in the probability of disaster, or changes in the probability of default. In Section 4, we solve the model with an inventory technology and show additional implications. Section 5 concludes.

2 Summary of the data

Figure 1: Nominal government rates

Panel A shows a five-year moving average of long-term nominal sovereign yields in the United Kingdom, Holland, Germany, Italy, and the United States from 1311–2018. The solid black line represents an average of all of the plotted series. Yields are from Schmelzing (2020) and are in annual terms. Yields come from a variety of archival, primary, and secondary sources. Panel B shows the nominal lending rate for the Bank of England expressed in annual terms.

A. Long-term sovereign yields **B. Bank of England lending rate**

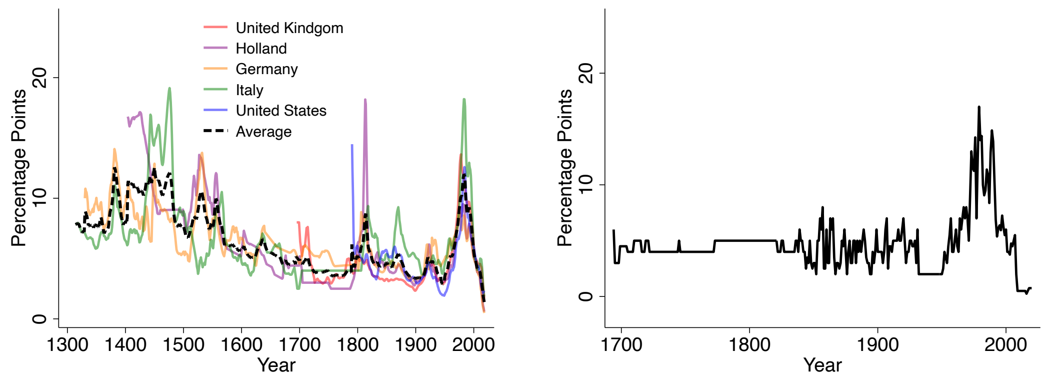

Panel A of Figure 1 shows nominal government rates in a seven-century-long dataset collected by Schmelzing (2020). Interest rates are highly volatile, as Jordà et al. (2019) emphasize.[4] Periods of extreme spikes, and also low rates, occurred around the American Revolution, Napoleonic Wars, and World War II, reflecting a tension between an increase in risk of sovereign default and precautionary savings around disasters. High rates in the 1970s and 1980s clearly stand out. Nonetheless, the figure shows a steady decline. Perhaps a more dramatic demonstration comes from Figure 1, Panel B, which shows the Bank of England lending rate, from the start of when the series was available. Only in the very

[4] Jordà et al. (2019) note that prior observations of a real rate of zero are not unusual. However, these are observations after subtracting ex post realized inflation, not ex ante yields. While it is true that both returns are zero from an investor's perspective, one was a realization of zero because of high inflation, whereas the other is an expected value of zero.

most recent period did this rate reach a zero lower bound.

Figure 2 zooms in on the last thirty years, the focus of much of the literature. The Federal Funds Rate in the U.S. declined sharply from 10% to 2% at present (Panel A). On the other hand, the price-dividend ratio has gone from around 20 to 50, implying a dividend yield of approximately 5% going to 2%—a smaller decline (Panel B). Moreover, the last row of Figure 2 displays the decline in the investment-capital ratio (Panel C) and real GDP growth (Panel D) from 1984–2016. Investment as a percentage of the capital stock went from an average of 7.7% to 6.9%, while real GDP growth declined from an average of 3.7% to 1.9%.

Figure 2: Various data moments, United States from 1984–2018

The figure shows the effective federal funds rate (shown in annual percentage points), the annual price-dividend ratio for the United States on the value-weighted CRSP index, the investment-capital ratio, and the annual real GDP growth rate for the United States.

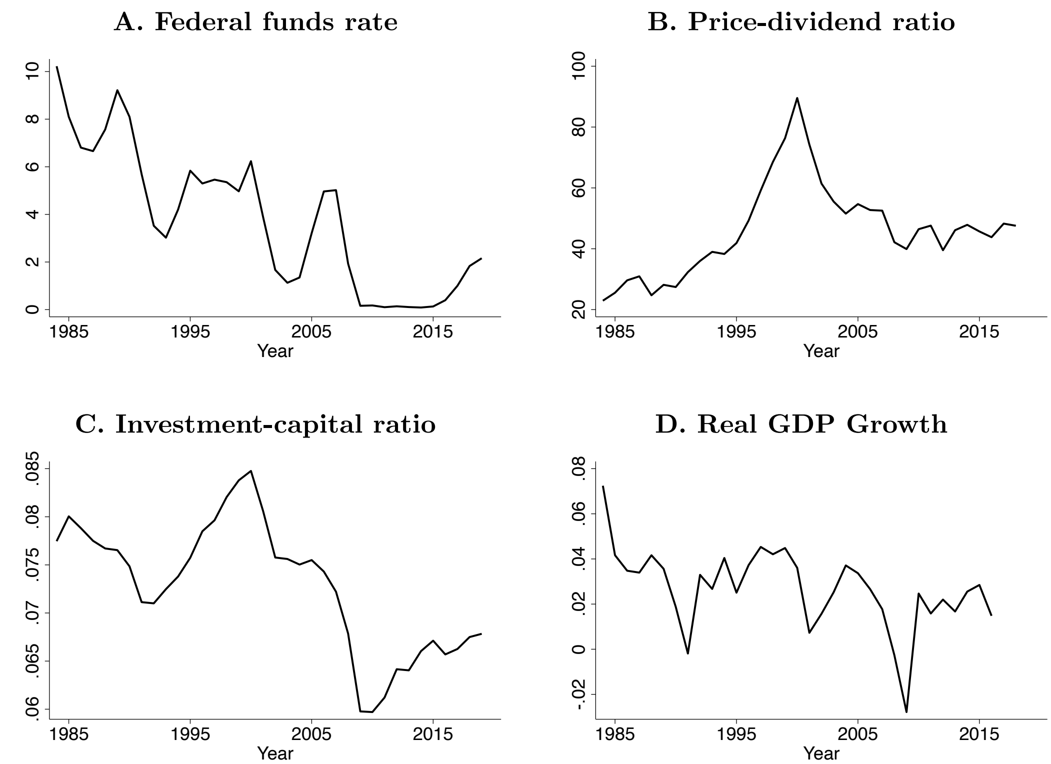

Figure 3 shows a longer time series of the price-dividend ratio, and also includes the cyclically-adjusted price-earnings (CAPE) ratio and the price-dividend ratio from the United Kingdom. It shows that the price-dividend ratio shifted upward in the late 1990s. This pattern does not appear in the CAPE ratio, nor in the U.K., and therefore may reflect a use of repurchases rather than cash payments as a means of returning cash to shareholders, and not a decline in interest rates (Boudoukh et al., 2007, Fama and French, 2001). For more information on the data and sources, see Appendix A.

Figure 3: Price-dividend and price-earning ratios: United States and United Kingdom

The figure shows the price-dividend ratio for the United States and United Kingdom since 1870 and the U.S. cyclically-adjusted price-earnings (CAPE) ratio. The black, solid line shows data for the United States price-dividend ratio and the red, solid line shows data for the price-dividend ratio of the United Kingdom. Price-dividend ratios are the end of year price divided by the aggregate dividends from the preceding year. The blue dashed-dotted line shows the CAPE ratio.

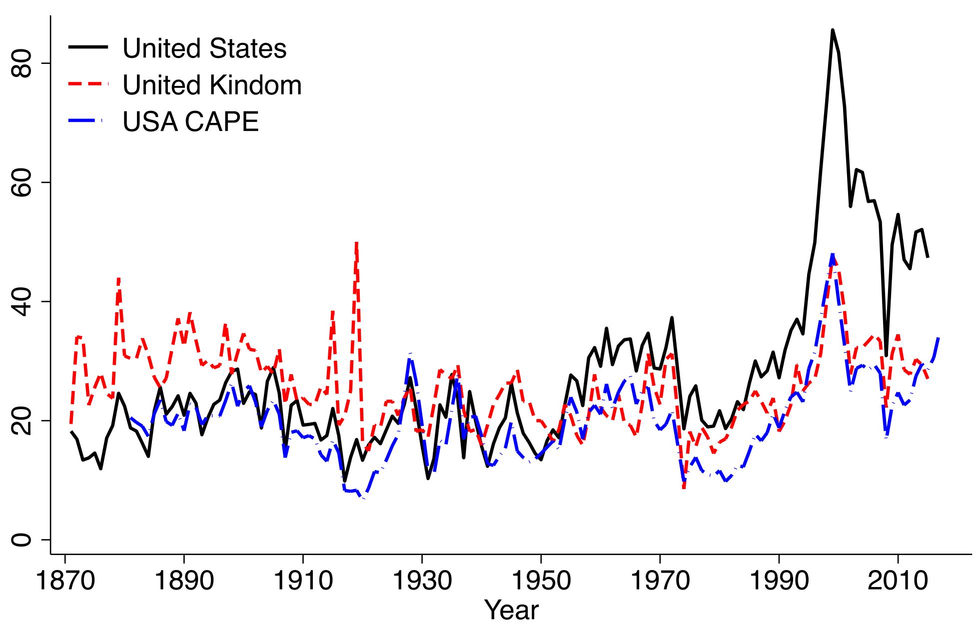

3 Endowment economy model with rare disasters

To interpret the data, we first turn to a standard endowment economy with a representative agent. Following Farhi and Gourio (2018), the model is calibrated separately to two sample periods (1984–2000 and 2001–2016), assuming constant parameters. While this approach means that certain features of the data (such as volatility of prices and interest rates) remain outside the scope of the analysis, it allows us to consider the possibility of long-run unforeseen structural changes. Farhi and Gourio assume a neoclassical growth model. We will return to such a model in the next section, but for the analysis at hand the extra degree of complication is not necessary. As far as prices and interest rates are concerned, and in this i.i.d. growth rate setting, the production model and the endowment model yield the same predictions.

Aggregate endowment evolves according to

$$C_{t+1} = C_t e^{\mu}(1 - \chi_{t+1}), \tag{1}$$

where χ_{t+1} represents an occurrence of rare disaster:

$$\chi_{t+1} = \begin{cases} 0 & \text{with probability} \quad 1 - p \\ \eta & \text{with probability} \quad p, \end{cases} \tag{2}$$

for $\eta \in (0, 1)$. Note that p represents the probability of a disaster and η its magnitude. We assume the representative agent has Epstein-Zin-Weil recursive preferences (Epstein and Zin, 1989, Weil, 1990) with risk aversion γ, elasticity of intertemporal substitution (EIS) ψ, and discount factor β. Let W_t denote the representative agent's wealth, here assumed to be the cum-dividend value of the consumption claim. Let $R_{W,t+1} \equiv W_{t+1}/(W_t - C_t)$ denote the return on wealth from time t to $t+1$. The stochastic discount factor (SDF) then equals

$$M_{t+1} \equiv \beta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} R_{W,t+1}^{\theta-1}, \tag{3}$$

for $\theta \equiv (1-\gamma)/(1-1/\psi)$.

In this section, we assume that the aggregate stock market equals aggregate wealth (ex-dividend) and that the ex post real return on the Treasury bill equals the riskfree rate. We relax these assumptions in the sections that follow. In equilibrium, $R_{W,t+1}$ must satisfy:

$$\mathbb{E}_t \left[M_{t+1} R_{W,t+1} \right] = 1. \tag{4}$$

Our assumptions and the endowment and preferences imply a constant price-dividend ratio $(W_t - C_t)/C_t$, which we denote by κ. Standard arguments (see Appendix B) then imply that

$$\kappa = \frac{\beta e^{(1-\frac{1}{\psi})\mu} \left[1 + p((1-\eta)^{1-\gamma} - 1) \right]^{\frac{1}{\theta}}}{1 - \beta e^{(1-\frac{1}{\psi})\mu} \left[1 + p((1-\eta)^{1-\gamma} - 1) \right]^{\frac{1}{\theta}}}. \tag{5}$$

Given the return on the wealth portfolio, the Euler equation provides the return on the one-period riskless bond:

$$\log R_f = -\log \beta + \frac{1}{\psi}\mu - \log(1 + p((1-\eta)^{-\gamma} - 1)) + \left(\frac{\theta - 1}{\theta} \right) \log(1 + p((1-\eta)^{1-\gamma} - 1)). \tag{6}$$

Equations (5) and (6) constitute a system of two equations in two unknowns, which can in principle be solved for p and β.[5] Combining (5) and (6) gives the

[5] There are some parameter combinations for which this is not possible. For example, when we try to set $\beta = .967$ for both 1984–2000 and 2001–2016 and then estimate μ and p, we are not able to obtain a solution for the 2001–2016 period.

equity premium:

$$\begin{aligned}\log \mathbb{E}_t[R_{W,t+1}] - \log R_f &= \log(1-p\eta) + \log(1+p((1-\eta)^{-\gamma}-1)) \\ &\qquad - \log(1+p((1-\eta)^{1-\gamma}-1)) \\ &\approx p\eta((1-\eta)^{-\gamma}-1)\end{aligned}$$

where the approximation is accurate for small p.

3.1 Increasing disaster probability

We calibrate this model assuming measured growth rates of $\mu = 0.0350$ from 1984 to 2000 and $\mu = 0.0282$ from 2001 to 2016. In what follows, we refer to μ as expected growth, even though it is in fact expected growth in the absence of disasters. Table 1 shows the results. For greatest comparability, we consider $\gamma = 12$, $\psi = 2$, and a disaster size $\eta = 0.15$, the same parameters used by Farhi and Gourio.[6] Similar to their findings, we match the data using a discount factor (β) of 0.967 in the early period and 0.979 in the later period, and a disaster probability (p) of 3.43% in the early period and 6.67% in the later period. We thus arrive at our first result: matching the combined stability of valuations and the decrease in riskfree rates requires a large increase in the disaster probability, even after accounting for decreased growth.

One can reasonably question the robustness of this interpretation, in particular, the role of the EIS. In fact, a change in the EIS changes the interpretation

[6] Farhi and Gourio estimate the growth rate within a neoclassical growth model with rare disasters. Their growth rate is the composition of three different growth rates: namely, the growth rate of TFP, the growth rate of the population, and the growth in investment prices. We instead take this growth rate as being exogenously set at the same level. We also use identical values for the price-dividend ratio and the riskfree rate, which are reported in Panel A of Table 1.

of the data, as Panel C of Table 1 shows. For an EIS below one—say, $\psi = 1/2$—matching (5) and (6) with p and β still requires an increase in p in the second period. However, β is now lower, implying that investors would need to have become less patient, not more, contradicting the demand-side intuition for the decline in interest rates (Summers, 2015).

Table 1: Accounting for the data with a change in disaster probability

This table shows parameters necessary to match the data, assuming an endowment economy with rare disasters and no inflation risk. Unless otherwise noted, we take average consumption growth from the data, and calibrate the disaster probability p and the subjective discount factor β to fit average interest rates and the price-dividend ratios in each of two sample periods. Because there is no inflation or inventory storage in the model, the riskfree rate proxies for the ex post real yield on the Treasury bill (Treasury bill yield minus realized inflation, or "inflation-adjusted Treasury yield"), and the wealth-consumption ratio proxies for the price-dividend ratio on the aggregate market. The table shows how p and β change depending on assumptions regarding elasticity of intertemporal substitution (EIS) and on growth. Treasury yields in the data, and parameters in the model, are annual.

		Values	
	Parameter	1984–2000	2001–2016
Panel A: Moments in the data			
Price-dividend ratio	κ	42.34	50.11
Inflation-adjusted Treasury yield	y_b	0.0279	-0.0035
Panel B: $\gamma = 12$, EIS $= 2$, $\eta = 0.15$			
Average consumption growth	μ	0.0350	0.0282
Discount factor	β	0.967	0.979
Probability of disaster	p	0.0343	0.0667
Panel C: $\gamma = 12$, EIS $= 0.5$, $\eta = 0.15$			
Average consumption growth	μ	0.0350	0.0282
Discount factor	β	0.997	0.983
Probability of disaster	p	0.0343	0.0667

To summarize: one needs an increase in the disaster risk to account for the data; but this explanation is fragile to a number of considerations, outlined in detail in Appendix Sections B.3 and B.4. van Binsbergen (2020) states the puzzle as follows: given the decrease in interest rates and the duration of the stock market, one would have expected a much larger capital gain if the risk premium were to remain constant.

3.2 Did the equity premium rise?

We now ask whether the equity premium did in fact rise. The literature studying long-run variation in the equity premium generally comes to the conclusion that the equity premium has declined over the postwar period, including from the first to the second periods that are our focus (Avdis and Wachter, 2017, van Binsbergen and Koijen, 2010, 2011, Fama and French, 2002, Lettau, Ludvigson and Wachter, 2008).

Options markets are another place to look for evidence of an increase in the equity premium. Virtually any explanation for an increase in the ex ante equity premium involves an increase in risk. While it is possible that such risk is not realized in sample, option prices incorporate the probability market participants assign to such risk materializing. Figure 4 shows the VIX, reported by the Chicago Board Options Exchange (CBOE). The VIX is the risk-neutral expectation of quadratic volatility, which is tightly tied to the equity premium. While the VIX is highly volatile, the average level of the VIX is remarkably stable between the two periods: equal to 21 in both. It is hard to reconcile this stability with an increase in the equity premium.

Given a model, one can say more. In Appendix C, we show how to go from the endowment economy model to a value of the VIX. A higher disaster probability implies a significantly higher VIX, not only because the ex ante volatility is higher (due to the realization of disasters), but because the risk-neutral volatility is higher still. If we ask the model to explain the level of the VIX in the earlier sample, and then modify the disaster probability as required, the VIX would need to rise from 21 to 23, rather than remain at 21 as it in fact did.[7] A test of whether the higher value is consistent with the data is rejected at the 1% level.

Figure 4: Chicago Board Options Exchange Volatility Index (VIX)

The figure plots the VIX series from 1986 to 2020 from the Chicago Board Options Exchange (CBOE). The long dashed red line is the average VIX from the beginning of the series to the end of the year 2000. The long dashed blue line shows the average VIX from the beginning of 2001 to 2016. Estimated averages in both samples are plotted with a two-standard-error confidence interval where standard errors are adjusted for heteroskedasticity and autocorrelation (Newey and West, 1987) with two lags on the monthly VIX.

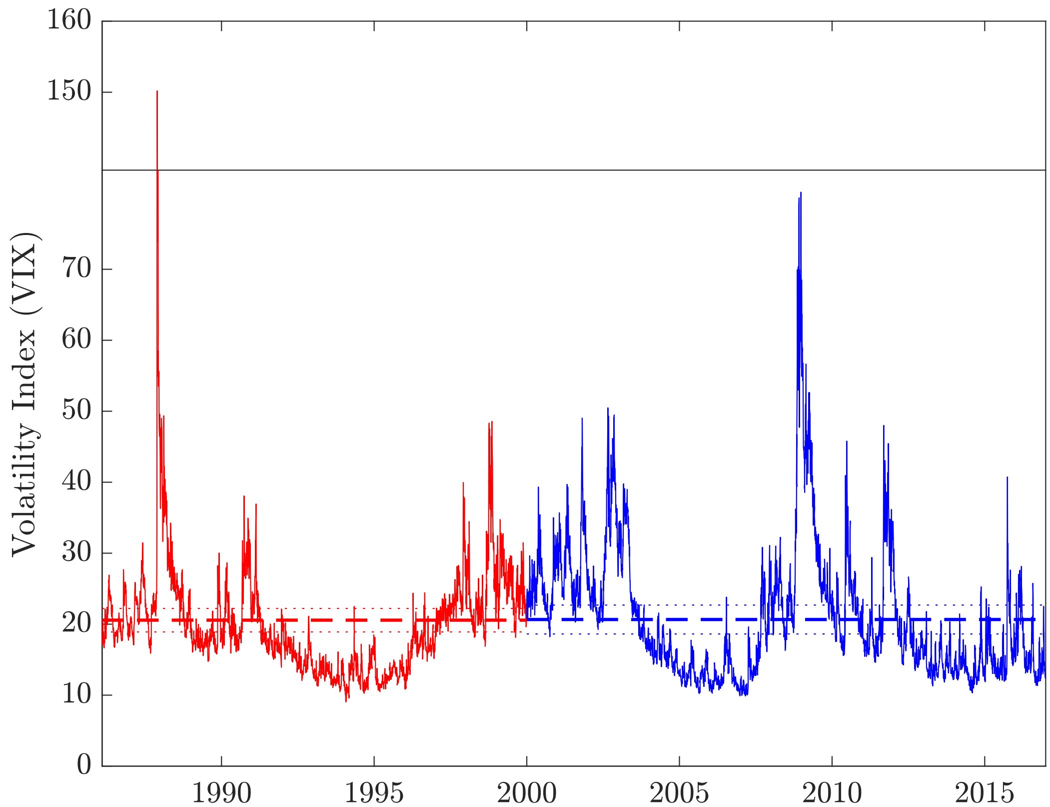

[7] In a similar way, Siriwardane (2015) and Seo and Wachter (2018) back out measures of disaster risk using options data and do not find an increase in the probability of disaster over this period.

3.3 Sovereign default risk

The typical empirical estimate of the equilibrium riskfree rate is the real return on short-term government debt; however, this return is not necessarily riskless, as the government can default either outright or through inflation. We now price this claim by including partial default that co-occurs with disasters. A decline in partial default risk can explain the secular trends in riskfree rates and valuation ratios since 1980.

Suppose, in a disaster, creditors lose a fraction $\lambda\eta$ relative to the face value of the bond. That is, a bond issued at time t pays $1 - L_{t+1}$ at time $t+1$, where loss $L_{t+1} = \lambda\chi_{t+1}$ represents a loss of zero if there is no disaster, and $\lambda\eta$ if a disaster should occur. If $\lambda = 1$, the loss to creditors is equal, in percentage terms, to the decline in consumption η. If $\lambda = 0$, then the bond is riskfree. Values of $\lambda < 0$ will correspond to deflation in disasters, as we describe below. Let Q_t be the price of the one-period bond, so that

$$Q_t = \mathbb{E}_t \left[M_{t+1}(1 - L_{t+1}) \right]. \tag{7}$$

Define the yield on the bond $y_b = y_{b,t} \equiv -\log Q_t$ as the log of the return on the bond in the case of no default—that is, the realized return when there is no disaster and the government makes its promised payment. Evaluating (7) implies:

$$y_b = \log R_f + \log\left(1 + p((1-\eta)^{-\gamma} - 1)\right) - \log\left(1 + p((1-\lambda\eta)(1-\eta)^{-\gamma} - 1)\right), \tag{8}$$

where R_f is the gross riskfree rate from (6). For $\lambda > 0$, the yield exceeds the riskfree rate. Letting $R_{b,t+1} \equiv (1 - L_{t+1})/Q_t$ denote the return on the defaultable bond, we have that the expected return

$$\begin{aligned}\log \mathbb{E}[R_{b,t+1}] &= \log R_f + \log(1 - p\lambda\eta) + \log\left(1 + p((1-\eta)^{-\gamma} - 1)\right) \\ &\qquad - \log\left(1 + p((1-\lambda\eta)(1-\eta)^{-\gamma} - 1)\right) \qquad (9) \\ &\approx \log R_f + p\lambda\eta((1-\eta)^{-\gamma} - 1).\end{aligned}$$

When $\lambda > 0$, the agent is compensated for the risk of partial default in a consumption disaster. Notice that the yield

$$y_b = \log \mathbb{E}[R_{b,t+1}] - \log(1 - p\lambda\eta) \approx \log \mathbb{E}[R_{b,t+1}] + p\lambda\eta$$

exceeds both the riskfree rate and the expected return on the bond when $\lambda > 0$. In a sample in which no disasters occur, the average ex post real return on the bond will correspond to the yield (8), not the expected return (9).

We have thus far assumed that the government defaults by failing to make part of its promised payments; a potentially more plausible way in which the government can default is through unexpected inflation. Let Π_t denote the price level, so that Π_{t+1}/Π_t is level inflation and $\Delta\pi_t = \log(\Pi_{t+1}/\Pi_t)$ is inflation in logs. The short-term nominal government bond has equilibrium price

$$Q_t^{\$} = \mathbb{E}_t\left[M_{t+1}\frac{\Pi_t}{\Pi_{t+1}}\right], \qquad (10)$$

with yield $y_{b,t}^{\$} \equiv -\log Q_t^{\$}$, and with nominal return $R_{b,t+1}^{\$} \equiv 1/Q_t^{\$}$. Note that the nominal return is known at time t with $R_{b,t+1}^{\$} = \exp\{y_{b,t}^{\$}\}$. A capital loss of L_{t+1} through default is equivalent, in real terms, to an inflation of $1/(1 - L_{t+1})$. To also allow for normal-times inflation, assume

$$\Pi_{t+1} = \Pi_t e^{\mu_{\pi,t}}(1 - L_{t+1})^{-1}, \qquad (11)$$

where $\mu_{\pi,t}$ represents growth in the price level that is locally deterministic (e.g. capturing lagged inflation). During disasters, the price level increases by a multiple $(1 - L_{t+1})^{-1}$ of its previous level, so that if consumption falls by 1/3, and $\lambda = 1$, inflation is 50% (that is, $1/(1 - \lambda\eta) = 1/(1 - 1/3) = 3/2$).

What is the relation between the nominal quantities and their real counterparts? Because $Q_t^{\$} = Q_t e^{-\mu_{\pi,t}}$, the yield on the nominal bond equals the real yield, plus an adjustment for normal-times inflation: $y_{b,t}^{\$} = y_b + \mu_{\pi,t}$. To compute the real return on the nominal bond, we take $R_{b,t+1}^{\$}$ and multiply by the inverse change in the price level. The expected real return on the nominal bond equals

$$\mathbb{E}_t\left[R_{b,t+1}^{\$}\frac{\Pi_t}{\Pi_{t+1}}\right] = \mathbb{E}_t[R_{b,t+1}],$$

namely, the same expected return as on the defaultable bond. Normal-times inflation is already priced in. Finally, the real yield y_b equals the average nominal yield minus inflation in samples in which no disasters occur:

$$\mathbb{E}[y_{b,t}^{\$} - \Delta\pi_{t+1} \,|\, \text{no disasters}] = \mathbb{E}[y_b + \mu_{\pi,t} - \mu_{\pi,t}] = y_b. \tag{12}$$

In what follows, we refer to (12) as the inflation-adjusted Treasury bill yield. It is the model counterpart of the average Treasury bill yield minus average realized inflation over the sample period of interest. One benefit of thinking of inflationary default, as opposed to outright default, is that it gives an intuitive interpretation of $\lambda < 0$. We can understand this case as corresponding to deflationary disasters, in which case the government bond becomes a hedge against disaster risk.

We now calibrate the model as we did above, but instead of varying p and β, we keep p constant at 3.43% (the calibrated value from 1984–2000) and allow λ to vary. The calibrated λ, therefore, gives the inflation risk premium such that no change in the disaster probability is needed to explain the decline in the riskfree rate and relative flatness of valuation ratios.

Table 2 shows the results. Panel B presents a calibration with risk aversion γ set to 5 and where consumption declines 30% in a disaster. We will use this calibration as our benchmark throughout the remainder of the paper, as it is more in line with estimates in the disaster risk literature. However, our points are qualitatively similar with higher γ and smaller disasters. The patience parameter

β does not rise as much as it did in the calibration where we allowed the disaster probability to vary, because it need only explain a slight increase in the price-dividend ratio. Further, there is a large inflation premium in the first half of the sample, and, essentially, no inflation premium in the second half. In Panel C, we consider an EIS equal to 1. Now neither μ nor p can affect the price-dividend ratio. Patience rises to account for the increase in the price-dividend ratio, and the default/inflation probability falls to explain the decline in the ex post real interest rate. Not surprisingly, these results are robust to alternative valuation measures, like the CAPE ratio, and to assumptions about the disaster probability.

Table 2: Accounting for the data with inflationary default risk

This table shows parameters necessary to match the data, assuming an endowment economy with rare disasters and inflationary default. We take average consumption growth from the data in each sample. We calibrate the discount factor β to match the average price-dividend ratio and the decline in bond value $\lambda\eta$ to match the average ex post real yield on the Treasury bill, assuming no disasters. We vary the elasticity of intertemporal substitution (EIS) as shown. We assume the disaster probability equals 3.43%, its benchmark value in Table 1. Parameters and yields are in annual terms.

		Values	
	Parameter	1984–2000	2001–2016
Panel A: Moments in the data			
Price-dividend ratio	κ	42.34	50.11
Inflation-adjusted Treasury yield	y_b	0.0279	-0.0035
Panel B: $\gamma = 5$, EIS $= 2$, $\eta = 0.30$			
Average consumption growth	μ	0.0350	0.0282
Discount factor	β	0.972	0.979
Fraction of bond value lost	$\lambda\eta$	0.163	0.016
Panel C: $\gamma = 5$, EIS $= 1$, $\eta = 0.30$			
Average consumption growth	μ	0.0350	0.0282
Discount factor	β	0.977	0.980
Fraction of bond value lost	$\lambda\eta$	0.163	0.016

The decline in the inflation risk premium is similar to the findings of Campbell, Pflueger and Viceira (2020), namely that inflation risk premia switched signs starting in 2001. While our estimates for λ vary across calibrations—we refer the reader to our final calibration in the next section for a quantitative estimate—each calibration shows that a decline in inflationary default risk is sufficient to explain the decline in real returns on short-term government debt.

3.4 Evidence from survey data

We have argued that a decline in the risk of unexpected inflation in a disaster is a more plausible explanation than a large and persistent increase in the probability of a consumption disaster. But is there evidence to support our explanation? Figure 5 shows that expected inflation in the U.S. has declined substantially over the past four decades. It is reasonable that the stabilization in inflation expectations would coincide with a stabilization in the risk of large and unexpected inflation shocks. Indeed, we also find independent evidence of declining inflationary default risk when we compare expected and realized inflation. Consider the inflation process implied by (11). Assuming agents incorporated a probability of inflationary default into their expectations that is never realized in sample, the difference between expected and realized log inflation is approximately given by $p\lambda\eta$.[8]

[8] For our given inflation process (11), a rational agent would predict inflation to be

$$\mathbb{E}_t[\Delta\pi_{t+1}] = \mu_{\pi,t} - p\log(1-\lambda\eta).$$

Now suppose there are no disasters in the sample, in which case the observed difference between realized and expected inflation is equal to

$$\mathbb{E}_t[\Delta\pi_{t+1}] - \Delta\pi_{t+1} = -p\log(1-\lambda\eta) \approx p\lambda\eta.$$

Figure 5: Expected inflation in the United States

The solid black line shows expected inflation from the Surveys of Consumers of University of Michigan. The dashed blue line shows the 10-year breakeven inflation rate computed from Treasury Inflation-Indexed Constant Maturity Securities. The dashed-dotted red line shows 10-year expected inflation from the Survey of Professional Forecasters.

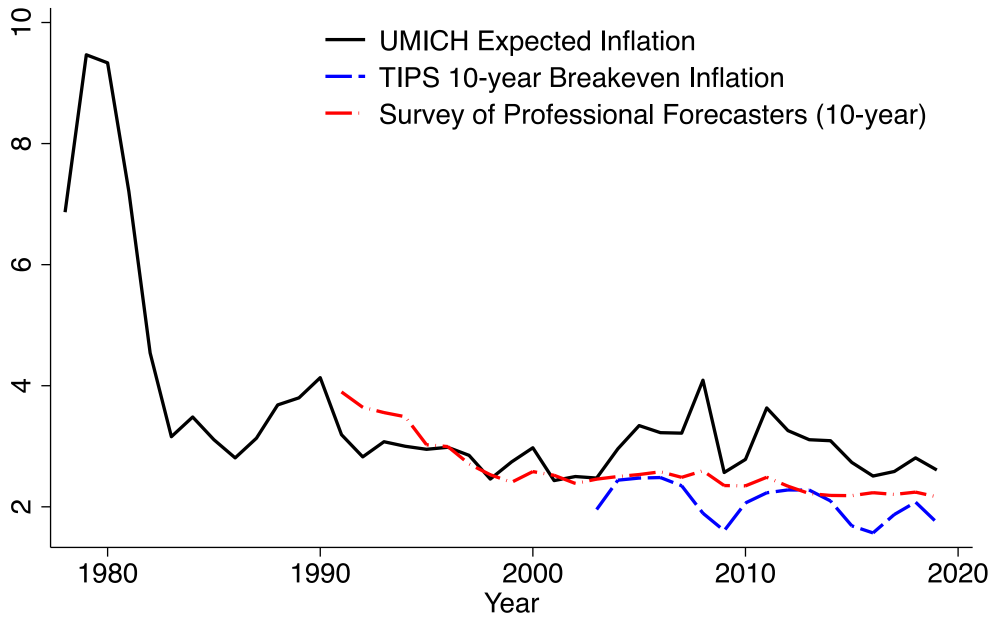

We estimate this difference using the one-year-ahead inflation forecast from the Survey of Professional Forecasters, which is plotted in Figure 6. The horizontal dashed lines show the average difference in each of our respective samples, along with two-standard-error confidence intervals, and can be interpreted as estimates of $p\lambda\eta$. If we assume a disaster probability of $p = 0.03$, then the implied bond value loss in the disaster state is roughly $\lambda\eta = 1/3$ from 1980 to 2000 and zero from 2000 to present. Quantitatively, this corresponds to an inflation disaster of around 50% and 0% in the first and second samples, respectively. There are other reasons for which expected and realized inflation could differ—for example, learning. We do not argue that there was a risk of inflation disasters as high as 50%; nonetheless, this exercise lends support to our explanation.

Our model and calibration imply that the *true* riskfree rate and equity risk premium have remained relatively constant over time, a conclusion that is consistent with evidence from valuation ratios, which have remained relatively flat; and the VIX, which suggests no substantial increase in risk. The price-dividend ratio, decomposed in (B.10), is unaffected by inflationary default risk. It is, however, common in the literature to use the return on the short-term government bond as a proxy for the true riskfree rate. Our calibration suggests that estimating the equity premium directly using this bond return implies an increase in the *measured* risk premium. In the model, this increase comes not from an increase in equity risk, but from a decline in the risk premium on government debt.

Figure 6: Expected versus realized one-year inflation

The figure plots the difference between expected and realized one-year inflation, where expectations are taken from the Survey of Professional Forecasters. The horizontal dashed lines show the average difference in each of our respective samples along with two-standard-error confidence intervals. These averages could be interpreted as estimates of $p\lambda\eta$ in our model, where p is the probability that a disaster occurs, and $\lambda\eta$ is the fraction of bond value lost when a disaster occurs.

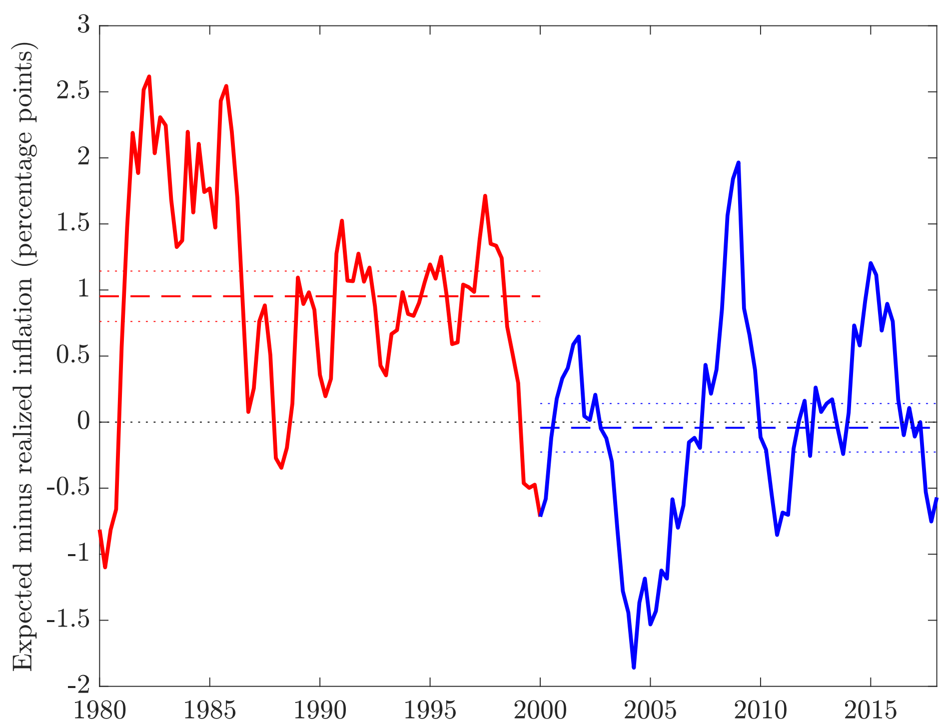

4 Production economy model

Section 2 shows that, from 1984 to 2016, the decline in interest rates and stability of valuation ratios correspond with a decline in investment relative to capital and GDP growth rates. To tie these four facts together, we introduce a model with a productive capital asset and a riskless inventory asset. We first solve the standard model with no inventory and show that, under any realistic calibration, it is not possible to simultaneously explain these secular trends. We then introduce inventory and show that its existence imposes an endogenous zero lower bound on the equilibrium riskfree rate and, at this lower bound, has real effects on investment and risk premia that can explain the data well.

4.1 No-inventory case

We consider a standard production model in which capital quality can decline suddenly and unpredictably.[9] Let K_t denote the quantity of productive capital at time t. Given K_t and constant productivity A, output equals

$$Y_t = AK_t. \tag{13}$$

Let δ denote depreciation, X_t investment. Capital evolves according to:

$$\tilde{K}_{t+1} \equiv X_t + (1-\delta)K_t \tag{14}$$

$$K_{t+1} \equiv \tilde{K}_{t+1}(1-\chi_{t+1}), \tag{15}$$

where χ_{t+1}, defined in (2), represents destruction of capital. We consider parameters such that $A > 1-\delta$, consistent with a growing economy. Following Gomes et al. (2019), we refer to $\tilde{K}$ as planned capital, the quantity of capital available if the disaster does not occur.

[9] See Barro (2009), Gabaix (2011), and Gourio (2012).

We can restate the agent's problem as a consumption-portfolio choice decision in which the agent allocates savings to capital and the riskfree bond. Let B_t denote the time-t dollar allocation to the riskfree asset. Define the agent's wealth at time t as

$$W_t \equiv C_t + B_t + \tilde{K}_{t+1}, \tag{16}$$

Then, if investment in capital grows at the stochastic rate $R_{K,t+1}$, wealth at time $t+1$ must equal

$$W_{t+1} = B_t R_{f,t+1} + \tilde{K}_{t+1} R_{K,t+1}. \tag{17}$$

What is $R_{K,t+1}$? Equations (13–15) indicate that, should a disaster not occur, a single unit of capital creates A units of output. A fraction δ is lost prior to the next period. Should a disaster occur, then a fraction χ_{t+1} is lost. Given the remaining capital, A units of output are created and an additional fraction δ is lost. Therefore, the return on capital

$$R_{K,t+1} = (1 - \delta + A)(1 - \chi_{t+1}). \tag{18}$$

We can rewrite the budget constraint in terms of flow variables. Applying (17) at time t implies:

$$W_t = B_{t-1} R_{f,t} + \tilde{K}_t R_{K,t}. \tag{19}$$

Equating (16) with (19) and substituting in for $R_{K,t}$ implies

$$C_t + B_t + \tilde{K}_{t+1} = B_{t-1} R_{f,t} + \tilde{K}_t (1 - \delta + A)(1 - \chi_t)$$

It follows from (15) that $\tilde{K}_t(1 - \chi_t) = K_t$. Using (14) and subtracting $(1-\delta)K_t$ from both sides implies

$$C_t + B_t + X_t = Y_t + B_{t-1} R_{f,t}. \tag{20}$$

That is, output from the capital stock, plus wealth in bonds can be used toward consumption, bond purchases at time t, or investment in the productive asset.

We can also rewrite the budget constraint in terms of the evolution of wealth. Define the share of savings invested in capital as

$$\alpha_t \equiv \frac{\tilde{K}_{t+1}}{W_t - C_t}.$$

Substituting in for B_t in (17) from (16) implies that

$$W_{t+1} = (W_t - C_t)(R_{f,t+1} + \alpha_t(R_{K,t+1} - R_{f,t+1})), \tag{21}$$

is an equivalent expression for the budget constraint. Let $R_{W,t+1} \equiv W_{t+1}/(W_t - C_t)$ denote the return on the wealth portfolio.

We assume Epstein and Zin (1989) and Weil (1990) preferences with unit EIS. The agent chooses consumption C_t and the capital portfolio share α_t to solve

$$\max_{C_t,\alpha_t} \left(C_t^{1-\beta} \left(\mathbb{E}_t \left[V(W_{t+1})^{1-\gamma} \right] \right)^{\frac{\beta}{1-\gamma}} \right), \tag{22}$$

subject to (21). Conjecturing that $V(W_t)$ equals a constant multiplied by W_t, and applying the first-order condition for optimal consumption implies the standard unit EIS result $C_t/W_t = 1 - \beta$.

In equilibrium, the bond is in zero net supply ($\alpha_t = 1$), and (20) reduces to

$$C_t + X_t = Y_t = AK_t. \tag{23}$$

Furthermore, $\alpha = 1$ and $C_t = (1-\beta)W_t$ imply that consumption is a fixed percentage of planned capital:

$$C_t = \frac{1-\beta}{\beta}\tilde{K}_{t+1} = \frac{1-\beta}{\beta}(X_t + (1-\delta)K_t). \tag{24}$$

The second equality follows from the capital accumulation equation (14).

What does this model imply for investment and for economic growth? Substituting in for C_t in (23) gives us the equilibrium investment-capital ratio with unit EIS:

$$\frac{X_t}{K_t} = \beta(1 - \delta + A) - (1 - \delta). \tag{25}$$

A savings glut unambiguously leads to an investment boom: (25) is strictly increasing in β. Evidently, an increased demand for savings coming from an increase in the β parameter implies an increase in the investment-capital ratio. Further, in the unit EIS case, risk does not affect the investment decision: lower investment relative to capital must come through either a reduction in β or the deterministic components of the return on capital A and δ. One may reconcile a decline in the riskfree rate with a decline in investment by arguing that productivity A or depreciation δ have declined. In order to match the decline in growth—a decline in μ in the endowment economy—one would need $A - \delta$ to decline as well. But even if this explanation succeeds at matching investment and interest rates, the puzzle of stable valuation ratios and the dependence of results on the EIS return. If the EIS were to exceed 1, increased risk could lead to a reduction in X/K, but this relies on scant evidence of increased risk and requires placing economically meaningful restrictions on the EIS.

Consumption, investment, and output grow at the same rate. First note that wealth grows at rate:

$$\frac{W_{t+1}}{W_t} = \frac{W_t - C_t}{W_t} \frac{W_{t+1}}{W_t - C_t} = \beta R_{K,t+1}. \tag{26}$$

(We have used the constant consumption-wealth ratio and the equilibrium condition $\alpha = 1$.) This must also be the growth rate of consumption. Substituting in for $R_{K,t+1}$ implies

$$\frac{C_{t+1}}{C_t} = \beta(1 - \delta + A)(1 - \chi_{t+1}). \tag{27}$$

This is then also the growth rate of planned capital, lagged one period. In equilibrium, all investment is in planned capital and so $\tilde{K}_{t+1}/\tilde{K}_t = W_t/W_{t-1}$. From (26), the relation between planned capital and actual capital, it follows that

$$\frac{K_{t+1}}{K_t} = \frac{\tilde{K}_{t+1}}{\tilde{K}_t} \frac{1 - \chi_{t+1}}{1 - \chi_t} = \beta R_{K,t} \frac{1 - \chi_{t+1}}{1 - \chi_t} = \beta(1 - \delta + A)(1 - \chi_{t+1}).$$

Note that we have used the fact that planned capital $\tilde{K}_{t+1}$ is a constant fraction of wealth W_t (in this model, this fraction is one), so that $\tilde{K}_{t+1}/\tilde{K}_t = W_t/W_{t-1} = \beta R_{K,t}$. The result for output then follows from $Y_t = AK_t$ and the result for investment follows from (23). As a consequence, if we let κ^Y denote the price-dividend ratio on the claim to output, $\kappa^Y = \kappa = \beta/(1-\beta)$, where κ denotes the price-dividend ratio on the consumption claim.

We now turn to the implications of this model for the interest rate and for stock returns. Given $V(W_t) \propto W_t$, the first-order condition with respect to α implies

$$\mathbb{E}_t \left[R_{W,t+1}^{1-\gamma}\right]^{-1} \mathbb{E}_t \left[R_{W,t+1}^{-\gamma}(R_{K,t+1} - R_{f,t+1})\right] = 0 \tag{28}$$

(see Appendix D for details). The equilibrium condition $\alpha = 1$ implies $R_{W,t+1} = R_{K,t+1}$. Substituting $R_{K,t+1}$ into (28) implies the riskfree rate equals:

$$\begin{aligned} R_f &= \mathbb{E}_t \left[R_{K,t+1}^{1-\gamma}\right] \mathbb{E}_t \left[R_{K,t+1}^{-\gamma}\right]^{-1} \\ &= (1-\delta+A)(1+p((1-\eta)^{1-\gamma}-1))(1+p((1-\eta)^{-\gamma}-1))^{-1}. \end{aligned} \tag{29}$$

Equations (28) and (29) imply the following expression for the SDF:

$$M_{t+1} = \mathbb{E}_t \left[R_{W,t+1}^{1-\gamma}\right]^{-1} R_{W,t+1}^{-\gamma}. \tag{30}$$

Furthermore, the risk premium equals

$$\begin{aligned} \log \mathbb{E}_t[R_{K,t+1}] - \log R_f &= \log(1-p\eta) \\ &\quad + \log(1+p((1-\eta)^{-\gamma}-1)) - \log(1+p((1-\eta)^{1-\gamma}-1)), \end{aligned} \tag{31}$$

exactly as in the endowment economy.

These asset pricing results are isomorphic to the endowment economy from Section 3. Indeed, equilibrium prices in the two models are identical if the parameters are such that the equilibrium consumption growth processes are the same.[10] The key difference between the models, however, is that there are two margins of adjustment in the production economy: quantities and prices. This is why, for example, the patience parameter β does not show up in (29). Instead, β influences quantities through the investment-capital ratio, which in turn affects prices. In the standard endowment economy, quantities cannot adjust, as the representative investor consumes whatever is produced in a given period.

Panel B of Table 3 calibrates the production model to match the cyclically-adjusted price-earnings (CAPE) ratio,[11] the short-term government bond rate, and the real GDP growth rate. Specifically, β, λ, and δ are chosen to match the data moments in the two samples. The calibration elucidates the puzzling nature of the reduction in the investment-capital ratio over the last four decades in light of falling interest rates and a stable marginal product of capital.[12] While the model matches rates and valuation ratios with reasonable estimates of β and λ, the calibration implies an increase in the investment-capital ratio (Table 4).

4.2 General case

Suppose now that, in addition to capital and a riskfree bond, the agent can put funds into inventory, namely a riskfree storage technology with a zero net return. If we impose the condition that riskfree storage be in zero supply, then

[10] In this setting, this occurs when $\beta^{-1}e^{\mu} = (1-\delta+A)$. One can verify this by comparing (6) and (29). In general, production and endowment economies can be mapped to one another by equating the consumption processes, a fact which is discussed in Chapter 2 of Cochrane (2001).

[11] We calibrate to the U.S. CAPE ratio because the U.S. price-dividend may be inflated by changes in the tendency of U.S. companies to pay dividends (Fama and French, 2001).

[12] Farhi and Gourio (2018) also note that the behavior of the investment-capital ratio represents a puzzle from the point of view of the standard model. They proxy for the marginal product of capital (MPK) using gross profitability, constructed as the ratio of (1−labor share) to the capital-output ratio. They also use the return on capital as constructed by Gomme et al. (2011). Both measures display a small increase in the MPK between the two samples.

the economy reduces to the one in the previous section. The innovation in this section is that the inventory can be in positive supply across the economy.

Why would one have a positive-supply riskfree asset? Conceptually, anything that is a store of value from one period to another could count as inventory, provided that it is in fact riskfree and can be frictionlessly interchanged between consumption and investment. Many consumption goods would not fit this description because they cannot easily be changed into something other than what they are. Money does fit this description provided that there is no unexpected inflation (in which case it ceases to be riskfree). To keep things simple, we will think of inventory as money.[13] Strictly speaking then, our analysis applies only to the second sample period, in which we estimate low, negative inflation risk. This turns out to make no difference—when the equilibrium real interest rate is greater than zero, inventory can exist but agents choose not to hold it.[14] Again, strictly speaking, if the asset is cash and there is non-zero expected inflation but no unexpected inflation, then we could specify a non-zero return on the inventory asset. However, expected inflation in the second sample period is small, and thus allowing for a slightly different return on inventory would make little difference. Likewise, we estimate unexpected inflation to be close to zero. Like all valuation equations, the existence of this riskfree storage is predicated on investors' (subjective) expectations about inflation. Evidence suggests (Reis, 2020) that investors believed inflation would be low and stable, and thus consistent with our assumptions on the existence of inventory. If, for example, government spending plays the role of inventory, as it does in Blanchard (2019), the fiscal theory

[13]This is not unlike the "social contrivance of money" as proposed by Samuelson (1958), which asserts that money can be used to obtain the socially optimal allocation in an overlapping generations framework in which the storage of consumption goods is impossible.

[14]Liquidity services could be such a reason investors choose to hold inventory in the presence of a positive riskfree interest rate, but for simplicity we do not assume these.

of the price level (Cochrane, 2021) could provide a foundation for these beliefs. However, inventory might not be synonymous with government spending.

Consider the agent's problem in Section 4.1, except here the agent can invest in a storage technology. The agent maximizes unit-EIS recursive utility by choosing consumption and B_t, I_t, and $\tilde{K}_{t+1}$. That is, the agent recursively solves

$$\max_{C_t, B_t, I_t, \tilde{K}_{t+1}} \left(C_t^{1-\beta} \left(\mathbb{E}_t \left[V(W_{t+1})^{1-\gamma} \right] \right)^{\frac{\beta}{1-\gamma}} \right), \tag{32}$$

subject to

$$W_t = C_t + B_t + I_t + \tilde{K}_{t+1} \tag{33}$$

$$W_{t+1} = B_t R_{f,t+1} + I_t + \tilde{K}_{t+1} R_{K,t+1} \tag{34}$$

$$I_t \geq 0. \tag{35}$$

The solution is still characterized by $V(W_t)$ a constant multiple of W_t, implying the standard unit EIS result of $C_t/W_t = 1 - \beta$.

We now characterize the equilibrium. Let R_f^* denote the equilibrium interest rate in the economy in Section 4.1, namely the economy with no inventory.

1. If $R_f^* > 1$, then in equilibrium $I_t = 0$, and the equilibrium is the same as in Section 4.1.

2. If $R_f^* < 1$, then $I_t > 0$. Investment in inventory crowds out investment in productive capital.

The argument is as follows (Appendix D gives an alternative and more formal argument). First assume parameters such that $R_f^* > 1$, and conjecture that $R_{f,t+1} = R_f^*$ in the problem (32–35) constitutes an equilibrium. This investor would never choose $I_t > 0$ because bonds offer superior returns; on the other hand (35) implies that the agent cannot short-sell inventory. Therefore $I_t = 0$, namely, the inventory asset is irrelevant, and thus $\alpha = 1$ is still the equilibrium condition. Equilibrium quantities and returns are the same as in Section 4.1.

Now assume that $R_f^* < 1$. Note that the only possible equilibrium value for R_f is unity; this is because $R_f < 1$ implies an arbitrage opportunity (the investor would borrow at R_f and invest the proceeds in the inventory asset), whereas if $R_f > 1$, the reasoning in the above paragraph implies the agent would hold no inventory. Then $R_f^* = R_f > 1$, contradicting the assumption. Intuitively, we can find an equilibrium with inventory for the following reason: if the agent does not hold inventory ($\alpha = 1$) and the riskfree rate equals $R_{f,t+1}^* < 1$, then the agent will wish to hold more inventory. Doing so, however, reduces the volatility of the return on the wealth portfolio and stochastic discount factor and thus increases the equilibrium riskfree rate. The agent will increase holdings of inventory until the equilibrium rate is equal to the return on inventory. Note the power of this reasoning: it implies we can proceed by analyzing the cases $R_f^* < 1$ and $R_f^* > 1$ separately.

We focus on the case of $R_f^* < 1$; as the above argument shows this is where inventory matters. We show it is also empirically relevant in that it prevails in the second sample period. Note that bonds are redundant and we can assume $B_t = 0$. The requirement $R_f = 1$ replaces $\alpha = 1$ as the equilibrium condition. Given that we have established that the equilibrium takes this form, for convenience we can rewrite the agent's optimization problem as

$$\max_{C_t,\alpha_t} \left(C_t^{1-\beta} \left(\mathbb{E}_t \left[V(W_{t+1})^{1-\gamma} \right] \right)^{\frac{\beta}{1-\gamma}} \right),$$

subject to

$$W_{t+1} = (W_t - C_t)(1 + \alpha_t r_{K,t+1}),$$

where $r_{K,t+1} = R_{K,t+1} - 1$ is the net return on capital and α has the same meaning as in Section 4.1. It is again the case that $V(W_t) \propto W_t$ and that $C_t/W_t = 1 - \beta$. The first-order condition for α continues to be (28), which, in the case with inventory, becomes

$$\mathbb{E}_t\left[\frac{1}{(1+\alpha r_{K,t+1})^{\gamma}} r_{K,t+1}\right] = 0. \tag{36}$$

Thus far we have not imposed distributional assumptions. Given our assumption on χ_t, we obtain:

$$\frac{p r_{K,\eta}}{(1+\alpha r_{K,\eta})^{\gamma}} + \frac{(1-p) r_{K,0}}{(1+\alpha r_{K,0})^{\gamma}} = 0, \tag{37}$$

where (with some abuse of notation), we let $r_{K,0} \equiv (1-\delta+A)-1$ and $r_{K,\eta} \equiv (1-\delta+A)(1-\eta)-1$ denote the net returns on capital in the non-disaster and disaster states, respectively. Solving for α implies:

$$\alpha = \min\left\{1, -\frac{((1-p)r_{K,0})^{1/\gamma} - (-p r_{K,\eta})^{1/\gamma}}{((1-p)r_{K,0})^{1/\gamma} r_{K,\eta} - (-p r_{K,\eta})^{1/\gamma} r_{K,0}}\right\}, \tag{38}$$

For future reference, we note that $R_{W,t} = \alpha R_{K,t} + (1-\alpha) = 1 + \alpha r_{K,t}$.

Because the consumption-wealth ratio is again $1-\beta$, we can apply the same reasoning used to show (27) to find:

$$\frac{C_{t+1}}{C_t} = \frac{W_{t+1}}{W_t} = \beta(1+\alpha r_{K,t+1}) = \beta\left(\alpha(1-\delta+A)(1-\chi_{t+1}) + 1 - \alpha\right). \tag{39}$$

Relative to the model in Section 4.1, consumption growth is less volatile because, in aggregate, agents use inventory to smooth out fluctuations. It is also, on average, lower, because less is invested in the productive asset. Output growth, however, is *more* volatile. Moreover, consumption growth is no longer tethered to output as in Section 4.1. While the relation between growth in the capital stock and growth in wealth remains the same:

$$\begin{aligned}\frac{K_{t+1}}{K_t} &= \frac{\tilde{K}_{t+1}}{\tilde{K}_t}\frac{1-\chi_{t+1}}{1-\chi_t} \\ &= \frac{W_t}{W_{t-1}}\frac{1-\chi_{t+1}}{1-\chi_t}.\end{aligned}$$

Substituting in from (39) now implies

$$\frac{Y_{t+1}}{Y_t} = \frac{K_{t+1}}{K_t} = \beta\left(\alpha(1-\delta+A)(1-\chi_{t+1}) + (1-\alpha)\left(\frac{1-\chi_{t+1}}{1-\chi_t}\right)\right), \tag{40}$$

Output growth is more volatile than consumption growth because it bears the full brunt of disasters: note that $1-\chi_{t+1}$ multiplies both the term with α (representing investment in the risky technology) and $1-\alpha$. By definition, the disaster applies to the entire existing capital stock. While this effect makes output growth more volatile than consumption in the present model, it does not, by itself, raise the volatility relative to the model in Section 4.1. There is, however, a second effect, represented by $1-\chi_t$ in the denominator. Coming out of a severe recession featuring capital destruction $\chi_t < 1$, output growth is higher, because agents invest more to get back to the optimal allocation. This raises the volatility of output growth relative to the model in Section 4.1.

What are the properties of investment? Rewriting the capital accumulation equation (14) so that X_t is on the left-hand side, and dividing by K_t implies

$$\begin{aligned}
\frac{X_t}{K_t} &= \frac{\tilde{K}_{t+1}}{\tilde{K}_t}\frac{\tilde{K}_t}{K_t} - (1-\delta) \\
&= \beta R_{W,t}(1-\chi_t)^{-1} - (1-\delta) \\
&= \beta(\alpha(1-\delta+A) + (1-\alpha)(1-\chi_t)^{-1}) - (1-\delta),
\end{aligned}$$

where we have used the fact that $\tilde{K}_{t+1}/\tilde{K}_t = \beta R_{W,t}$, due to the fact that the fraction invested in risky capital, and the consumption-wealth ratio are both constant. The investment-capital ratio is time-varying in this economy, despite the fact that shocks are i.i.d. and that there is a balanced growth path.

Figure 7: Investment capital ratio in the model

The figure shows how capital investment varies with the size of the consumption decline in a disaster for the production models with and without inventory. The figure plots the investment-capital ratio X/K in the model with inventory when there is and is not a disaster, and in the model without inventory. It also plots α, the share of savings invested in capital. Risk aversion $\gamma = 6$, the EIS $\psi = 1$, the patience parameter $\beta = 0.964$, depreciation $\delta = 0.057$, the probability of disaster $p = 0.0343$, and the marginal product of capital $A = 0.12$. The dotted black line represents the point at which the riskfree rate is equal to 0 in the model without inventory.

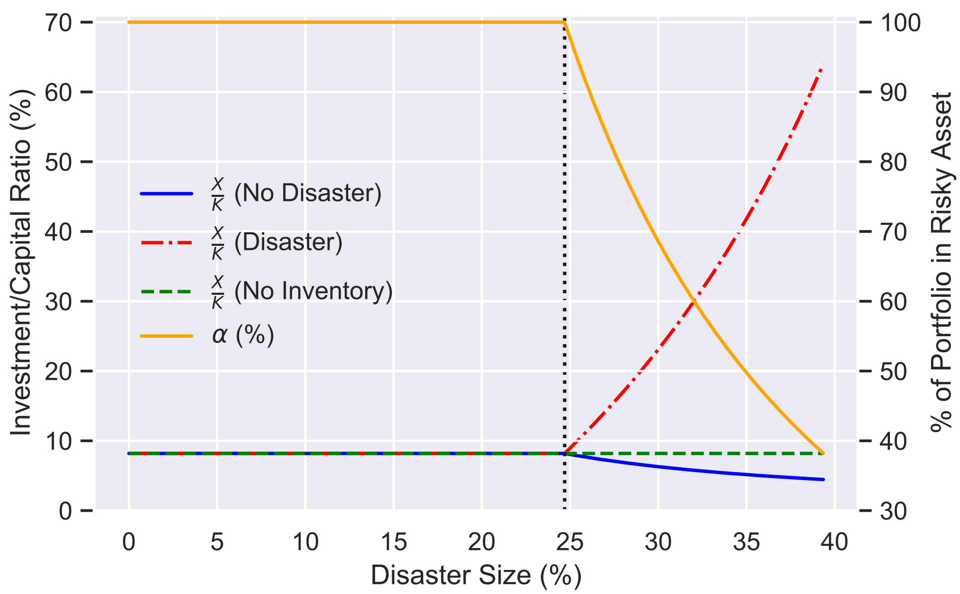

A disaster in the prior period increases investment in productive capital. This is because the disaster affected capital disproportionately, and the agent must re-invest to return capital back to its pre-crisis level. For an illustration, see Figure 7, which shows the investment-capital ratio for $\chi_t = 0$ (no disaster) and $\chi_t = \eta$ (disaster) for various values of the disaster size. The figure also shows the optimal planned capital to wealth ratio, which is constant. For comparison, the figure also shows quantities in the case of no inventory. Fixing other parameters, for disaster sizes of less than 25%, the gross riskfree rate is above one, implying that the economies with and without inventory are the same. As the size of the disaster increases, the equilibrium riskfree rate in the no-inventory economy falls sharply (we illustrate this in Figure 9). It becomes optimal to hold inventory and investment in productive assets falls. At that point, investment depends on the occurence of a disaster in the prior period. The greater the size of the disaster, the greater the increase in investment. In contrast, with no inventory, the investment-capital ratio is always the same.

We define the stock market as the claim to output Y_t in all future periods. As (40) shows, the growth rate of capital is no longer i.i.d. but depends on χ_t (note that χ_{t+1} is i.i.d. given time-t information). Therefore, the price-dividend ratio on the output claim is a function of χ_t and solves

$$\kappa^Y(\chi_t) = \mathbb{E}_t\left[M_{t+1}\left(1+\kappa^Y(\chi_{t+1})\right)\frac{Y_{t+1}}{Y_t}\right],$$

where the stochastic discount factor takes the same form as (30), with $R_{W,t+1}$ now given as above. Note that $R_{W,t+1}$ is i.i.d. Under our distributional assumptions:

$$\kappa^Y(0) = \frac{\hat{\beta}}{1-\beta}\left((1-p)(1+\alpha r_{K,0})^{1-\gamma} + p(1+\alpha r_{K,\eta})^{-\gamma}(1+\alpha r_{K,0})(1-\eta)\right), \tag{41}$$

$$\kappa^Y(\eta) = \kappa^Y(0)\left(\frac{1+\alpha r_{K,\eta}}{1+\alpha r_{K,0}}\right)(1-\eta)^{-1}, \tag{42}$$

where $\hat{\beta} \equiv \beta\left((1-p)(1+\alpha r_{K,0})^{1-\gamma} + p(1+\alpha r_{K,\eta})^{1-\gamma}\right)^{-1}$. See Appendix D for details. In the case where $\alpha = 1$, the price-dividend ratio is the constant $\kappa^Y = \beta/(1-\beta)$.

Figure 8 shows the price-dividend ratio for various levels of the disaster size, both in the economy with inventory and the economy without. The economy without inventory has a constant price-dividend ratio solely determined by β. When there is inventory, the price-dividend ratio rises in disasters because dividends are temporally depressed (they are also low because of the disaster). This increase is due to the endogenous investment response, whereby inventory is liquidated in a disaster to replace the capital that is destroyed.

Figure 8: Price-dividend ratio in the model

The figure shows how the price-dividend ratio varies with the size of the consumption decline in a disaster for the production models with and without inventory. The figure plots the price-dividend ratio in the model with inventory when there is and is not a disaster, and in the model without inventory. Risk aversion $\gamma = 6$, the EIS $\psi = 1$, the patience parameter $\beta = 0.964$, depreciation $\delta = 0.057$, the probability of disaster $p = 0.0343$, and the marginal product of capital $A = 0.12$. The dotted black line represents the point at which the riskfree rate is equal to 0 in the model without inventory.

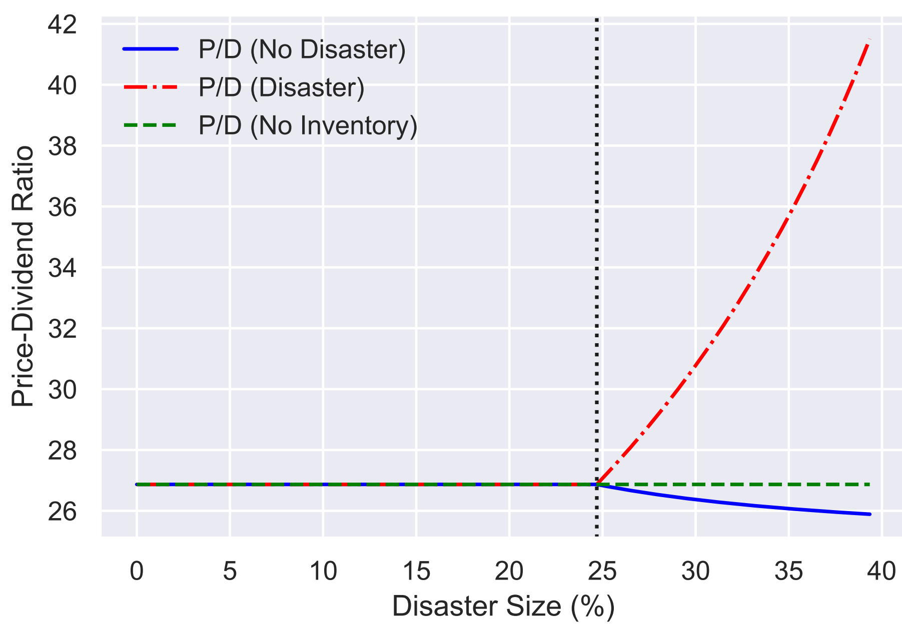

Note also that in contrast with standard production models, the price-dividend ratio in the case of no disaster declines (in a comparative statics sense) as a function of the disaster probability (see Figure 8). In the case without inventory, the price-dividend ratio is independent of the disaster probability. Models with production that seek to match business-cycle fluctuations in investment and valuation ratios require the EIS to be greater than 1. Endowment models achieve the same effect by imposing exogenous leverage (dividends more sensitive to shocks than consumption). In this model, leverage is endogenous, and qualitatively correct price-dividend ratio dynamics could in principle occur, even with an EIS of one. The magnitude of the decline in Figure 8 suggests that the effect is small.

Figure 9 shows that the equity risk premium in this economy loses its usual dependence on disaster risk. The equity premium equals $rp \equiv \log \mathbb{E}_t[R_{Y,t+1}] - \log R_f$, where the return on the output claim is

$$R_{Y,t+1} = \left(\frac{\kappa^Y(\chi_{t+1}) + 1}{\kappa^Y(\chi_t)}\right)\left(\frac{Y_{t+1}}{Y_t}\right).$$

The blue line in the figure shows the equity premium in the model without inventory: it is highly dependent on the disaster probability, as is the riskfree rate. However, the return on capital—which, in the economy with no inventory, is the equity return—is only very slightly decreasing. This is a standard result in disaster-risk economies: the full discount rate on the equity claim decreases slightly with the probability of a disaster.

While this might seem counterintuitive, it arises from the fact that, while the equity premium increases, the riskfree rate declines, and more than offsets the effect. Also recall that the continuously compounded return in a standard i.i.d. economy can be expressed as the log dividend yield plus the log growth in cash flows. When the EIS equals one, the dividend yield does not depend on the disaster probability, and so the only effect is the small one, through expected cash flows. In the economy with inventory, the return on capital is the same as in the economy without (this is defined by the production opportunities), and thus is slightly decreasing. The riskfree rate is constant, implying that the premium on capital is also slightly decreasing. The equity premium decreases slightly more in the disaster size as compared to $\log \mathbb{E}[R_K] - \log R_f$. This is because the increase in the price-dividend ratio counteracts the decline in output due to the disaster.

Figure 9: Risk premia and riskfree rate in the model

The figure shows how the riskfree rate and risk premium vary with the size of the consumption decline in a disaster for the production models with and without inventory. The moments plotted are: the equity premium in the models with and without inventory, rp^{inv} and rp^{*}; the riskfree rates in the models with and without inventory, r_f^{inv} and r_f^{*}; and the expected return on capital, $\mathbb{E}[r_K]$. The equity premium is defined as the log expected return on the output claim minus the log riskfree rate. Risk aversion $\gamma = 6$, the EIS $\psi = 1$, the patience parameter $\beta = 0.964$, depreciation $\delta = 0.057$, the probability of disaster $p = 0.0343$, and the marginal product of capital $A = 0.12$. The dotted black line represents the point at which the net riskfree rate is equal to 0 in the model without inventory.

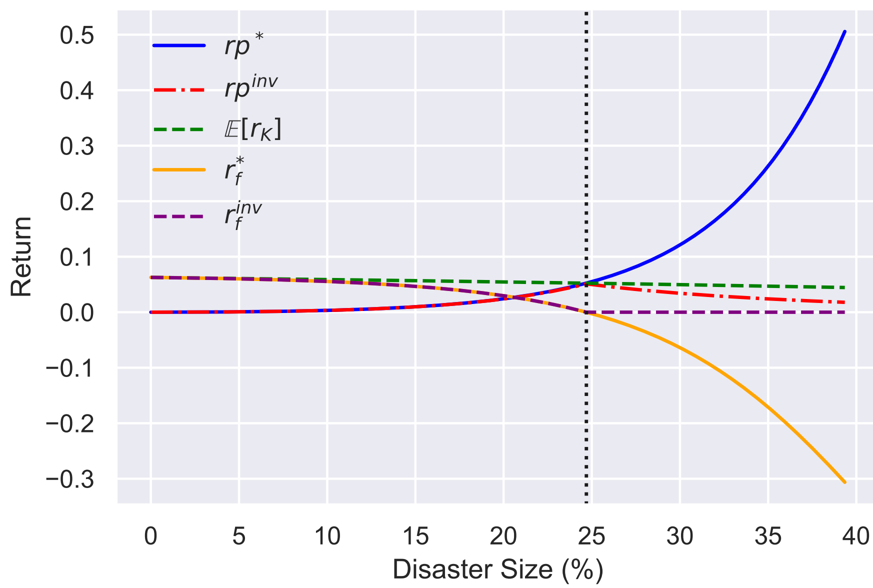

Finally, the inflation-adjusted Treasury yield is equal to

$$y_b = \log\left(p(1+\alpha r_{K,\eta})^{1-\gamma} + (1-p)(1+\alpha r_{K,0})^{1-\gamma}\right)$$
$$- \log\left(p(1+\alpha r_{K,\eta})^{-\gamma}(1-\lambda\eta) + (1-p)(1+\alpha r_{K,0})^{-\gamma}\right). \quad (43)$$

Notice that, while the true riskfree rate cannot go below zero, the yield and expected return on the defaultable claim could be positive or negative, depending on the direction of the default risk premium.

The model is calibrated to match the real interest rate, price-dividend ratio, and GDP growth in the U.S., as in the sections above. Calibrating to match these data requires solving a system of three equations in three unknowns, where the unknowns are the parameters β, λ, and δ and the three equation are Equations (41), (43) and

$$\frac{Y_{t+1}(0)}{Y_t(0)} = \beta\left(\alpha(1-\delta+A)+(1-\alpha)\right) \tag{44}$$

which is GDP growth when the disaster does not occur. Indeed, each of the moments to which we calibrate parameters is the value in the no-disaster state ($\chi_t = 0$), consistent with the fact that we do not observe any disasters in our sample. We then find the values of the parameters of interest that make it such that the data moments match their corresponding model moments.

The results from the calibration are displayed in Table 3. The model with inventory is able to match the data moments with a quantitatively reasonable calibration of β, λ, and δ. The slight increase in β matches the modest rise in the CAPE ratio; similarly, higher depreciation matches the lower growth in the second sample. The calibrated inflation default premium is in line with the empirical estimates for λ in Section 3. One can interpret the estimated values of $\lambda\eta$ as the percent inflation that would occur in a disaster. In the model with inventory, the calibration suggests an 11% disaster inflation in the first sample and a 2% deflation in the second. Finally, note that the presence of inventory allows the model to match the low growth from 2001 to 2016 with a smaller rise in depreciation than the model without inventory.

The smaller rise in depreciation comes from the crowding out effect, whereby inventory substitutes for investment in productive capital. This crowding out is shown in Table 4 by the lower investment-capital ratio from 2001–2016 in the model with inventory. In the model without inventory, the investment-capital ratio increases from the first half to the second half of the sample, due to the slightly higher β and δ. In the model with inventory, however, the investment-capital ratio is lower in the second half of the sample in the state where a disaster has not occurred, as the agent diverts investment from capital to inventory in the non-disaster states, anticipating that a disaster will occur in the future. When the disaster finally comes, the agent invests heavily, as seen in Figure 7. The crowding out effect of inventory in normal states allows the model to match the reduction in the investment-capital ratio in the data *without targeting it*.

Table 3: Inventory and inflationary default in a model with production

The model is solved with risk aversion $\gamma = 6$ and EIS $\psi = 1$. Consumption declines 30% in a disaster ($\eta = 0.30$), the probability of disaster $p = 3.43\%$, and the marginal product of capital $A = 0.12$.

		Values	
	Parameter	1984–2000	2001–2016
Panel A: Moments in the data			
US CAPE ratio	κ	25.97	26.73
Inflation-adjusted Treasury yield	y_b	0.0279	-0.0035
US GDP growth	$\frac{\Delta Y}{Y}$	0.0368	0.0191
Panel B: Without inventory, $\gamma = 6$, EIS $= 1$, $\eta = 0.30$			
Discount factor	β	0.963	0.964
Fraction of bond value lost	$\lambda\eta$	0.108	0.055
Capital depreciation	δ	0.043	0.063
Panel C: With inventory, $\gamma = 6$, EIS $= 1$, $\eta = 0.30$			
Discount factor	β	0.963	0.964
Fraction of bond value lost	$\lambda\eta$	0.108	-0.018
Capital depreciation	δ	0.043	0.057

Table 4 also displays the endogenous zero lower bound that inventory creates in the second half of the sample. In the sample from 2001–2016, the unconstrained riskfree rate in the model with inventory is -1.1 percentage points (meaning that, were it not for the existence of inventory, the riskfree rate would be 1.1 percentage points lower). To achieve equilibrium in the model, the agent increases his or her holdings of inventory until the riskfree rate in the economy is zero. In Table 4, this happens when roughly 9% of the portfolio is invested in inventory. In the first half of the sample, however, the unconstrained riskfree rate is slightly greater than zero (around 20 basis points) meaning that the representative agent has no incentive to hold inventory.

Table 4: Inventory and inflationary default with production: untargeted moments

The model is solved with risk aversion $\gamma = 6$ and EIS $\psi = 1$. Consumption declines 30% in a disaster ($\eta = 0.30$), the probability of disaster $p = 3.43\%$, and the marginal product of capital $A = 0.12$. The calibrated parameters from Table 3 are used for β, λ, and δ.

		Values	
	Parameter	1984–2000	2001–2016
Panel A: Without inventory, $\gamma = 6$, EIS $= 1$, $\eta = 0.30$			
Risky capital share	α	1.000	1.000
Investment-capital ratio	$\frac{X}{K}$	0.080	0.082
Unconstrained riskfree rate	r_f^*	0.002	-0.016
Panel B: With inventory, $\gamma = 6$, EIS $= 1$, $\eta = 0.30$			
Risky capital share	α	1.000	0.912
Investment-capital ratio	$\frac{X}{K}$	0.080	0.077
Unconstrained riskfree rate	r_f^*	0.002	-0.011

5 Concluding remarks

The puzzle of low interest rates is a puzzle not only from the point of view of the last quarter century, but over a much longer horizon. It is also a joint puzzle: why have low interest rates not been accompanied by higher valuation ratios?

The purpose of this article is to argue that the most natural explanation is not an increased demand for savings, which would lower interest rates and raise valuation ratios; nor a decrease in growth, which is hardly enough on its own to account for the observed change; nor an increase in the risk premium, as there is no evidence that risk has increased by nearly the required amount. These joint phenomena have a simple explanation, which is that *the true riskfree rate has hardly changed at all*. Short-term debt claims are defaultable, and investors have come to require a lower premium for this risk of default.

Because our explanation implies that the true riskfree rate has remained roughly constant, we require a framework that allows for, first, a riskfree rate that is sufficiently low to explain nominal debt yields at zero and, second, an explanation that survives the existence of a zero lower bound. We accomplish the former using a model with a risk of rare disasters. In a rare disaster model, investors' precautionary savings pushes the riskfree rate below zero. We accomplish the latter by introducing a costless storage technology into a production economy. When parameters are such that the true riskfree rate is below zero, agents choose to save into inventory until markets clear at a riskfree rate of zero.

What we do not model is the cause for the decline in investor expectations of sovereign default. Evidence suggests that this decline both has a relatively short-term component based on the history of the last 30 years and a long-term component spanning centuries, based on a growing faith over time in the stability of sovereigns. The forces determining this shift in expectations are an interesting topic for further research.

References

Abel, Andrew B., "Risk premia and term premia in general equilibrium," *Journal of Monetary Economics*, 1999, *43* (1), 3 – 33.

Avdis, Efstathios and Jessica A. Wachter, "Maximum likelihood estimation of the equity premium," *Journal of Financial Economics*, 2017, *125* (3), 589–609.

Barro, Robert J., "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 2009, *99* (1), 243–264.

Baumol, William J., "The Transactions Demand for Cash: An Inventory Theoretic Approach," *Quarterly Journal of Economics*, 1952, *66* (2).

Bernanke, Ben, "The global savings glut and the U.S. current account deficit," March 2005. Sandridge Lecture, Virginia Association of Economics, Richmond, VA.

Binsbergen, Jules H. van, "Duration-Based Stock Valuation," 2020. National Bureau of Economic Research Working Paper No. 27367.

_ and Ralph S. J. Koijen, "Likelihood-Based Estimation of Exactly-Solved Present-Value Models," *Journal of Finance*, 2010, *65* (4), 1439–1471.

_ and _ , "Likelihood-Based Estimation of Exactly-Solved Present-Value Models," 2011. Working paper, University of Pennsylvania, University of Chicago.

Blanchard, Olivier, "Public Debt and Low Interest Rates," *American Economic Review*, 2019, *109* (4), 1197–1229.

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts, "On the importance of measuring payout yield: Implications for empirical asset pricing," *Journal of Finance*, 2007, *62* (2), 877–915.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre Olivier Gourinchas, "An equilibrium model of global imbalances and low interest rates," *American Economic Review*, 2008, *98* (1), 358–393.

Campbell, John Y., "Consumption-based asset pricing," in "Handbook of the Economics of Finance," Vol. 1, Elsevier, 2003, pp. 803 − 887.

_ and Robert J. Shiller, "The dividend-price ratio and expectations of future dividends and discount factors," *Review of Financial Studies*, Fall 1988, *1* (3), 195–228.

_ , Carolin Pflueger, and Luis M. Viceira, "Macroeconomic Drivers of Bond and Equity Risks," *Journal of Political Economy*, 2020, *Forthcoming.*

Chernov, Mikhail, Lukas Schmid, and Andres Schneider, "A Macrofinance View of U.S. Sovereign CDS Premiums," *Journal of Finance*, 2020, *Forthcoming.*

Cochrane, John H., *Asset Pricing*, Princeton, NJ: Princeton University Press, 2001.

Cochrane, John H, "The Fiscal Theory of the Price Level," April 2021. Unpublished manuscript, Stanford University.

Dormans, Eduard H.M, "Het tekort. Staatsschuld in de tijd der Republiek." PhD dissertation, Tilburg University 1991.

Durand, David and Willis J. Winn, *Basic Yields of Bonds, 1926-1947: Their Measurement and Patterns*, Vol. 6, Cambridge, MA: National Bureau of Economic Research, 1947.

Epstein, Larry and Stan Zin, "Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, 1989, *57*, 937–969.

Fama, Eugene F and Kenneth R French, "Disappearing dividends: changing firm characteristics or lower propensity to pay?," *Journal of Financial economics*, 2001, *60* (1), 3–43.

Fama, Eugene F. and Kenneth R. French, "The Equity Premium," *Journal of Finance*, 2002, *57* (2), 637–659.

Farhi, Emmanuel and François Gourio, "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia," *Brookings Papers on Economic Activity*, Fall 2018, pp. 147–250.

Gabaix, Xavier, "Disasterization: A simple way to fix the asset pricing properties of macroeconomic models," *American Economic Review: Papers & Proceedings*, 2011, *101* (3), 406–409.

Gomes, João F, Marco Grotteria, and Jessica A Wachter, "Cyclical Dispersion in Expected Defaults," *Review of Financial Studies*, April 2019, *32* (4), 1275–1308.

Gomme, Paul, B. Ravikumar, and Peter Rupert, "The Return to Capital and the Business Cycle," *Review of Economic Dynamics*, 2011, *14* (2).

Gordon, Robert J, "Secular Stagnation: A Supply-Side View," *American Economic Review: Papers & Proceedings*, 2015, *105* (5), 54–59.

Gourio, François, "Disaster Risk and Business Cycles," *American Economic Review*, 2012, *102* (6), 2734–2766.

Grossman, Richard S., "New Indices of British Equity Prices," *Journal of Economic History*, 2002, *62* (1).

Hansen, Alvin H, "Economic Progress and Declining Population Growth," *American Economic Review*, 1939, *29* (1), 1–15.

Homer, Sidney and Richard E. Sylla, *A History of Interest Rates*, New Brunswick, NJ: Rutgers University Press, 2005.

Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, "The declining equity premium: What role does macroeconomic risk play?," *Review of Financial Studies*, 2008, *21*, 1653–1687.

Newey, Whitney K. and Kenneth D. West, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, *55* (3).

Òscar Jordà, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor, "The Rate of Return on Everything, 1870–2015," *Quarterly Journal of Economics*, 2019, *134* (3), 1225–1298.

Reis, Ricardo, "The anchoring of long-run inflation expectations today," March 2020. Inflation dynamics in Asian and the Pacific, BIS Papers 111.

Samuelson, Paul A., "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 1958, *66* (6).

Schmelzing, Paul, "Eight centuries of global real interest rates, $R-G$, and the 'suprasecular' decline, 1311–2018," *Bank of England Staff Working Paper No. 845*, 2020, pp. 1–110.

Seo, Sang Byung and Jessica A. Wachter, "Do Rare Events Explain CDX Tranche Spreads?," *The Journal of Finance*, August 2018, *73* (5), 2343–2383.

— and — , "Option Prices in a Model with Stochastic Disaster Risk," *Management Science*, 2019, *65* (8), 3449–3469.

Shiller, Robert J., *Irrational Exuberance*, Princeton University Press, 2000.

Siriwardane, Emil, "The Probability of Rare Disasters: Estimation and Implications," October 2015. Harvard Business School Working Paper, No. 16-061.

Summers, Lawrence H, "Demand Side Secular Stagnation," *American Economic Review: Papers & Proceedings*, 2015, *105* (5), 60–65.

Weeveringh, J.J., *Handleiding tot de Geschiedenis der Staatsschulden*, Amsterdam: Gebroeders Kraay, 1852.

Weil, Philippe, "Nonexpected utility in macroeconomics," *Quarterly Journal of Economics*, February 1990, *105*, 29–42.

Zimmermann, Kaspar, "Breaking Banks? Monetary Policy and Bank Profitability," September 2017. Working paper, University of Bonn.

A Data appendix

We use various series to illustrate the secular decline in interest rates in the short- and long-run. To obtain interest rates from 1311–2018, we rely on data from Schmelzing (2020). The dataset contains nominal interest rate and inflation time series for several developed economies over the last eight centuries. Specifically, the data include long-term sovereign borrowing rates with an average maturity that hovers around 10 years; however, this varies over time and across countries. From these data, we plot the nominal sovereign borrowing yields for the United Kingdom, Holland, Germany, Italy, and the United States in Panel A of Figure 1. The data are collected from a variety of sources, outlined in detail in the paper and online appendix. The U.K. borrowing rates come from the Calendar of State Papers and the Bank of England. Data before 1694 for the U.K. (before the founding of the Bank of England) are not used, since the data are incomplete. Data for the Netherlands come from Dormans (1991), Weeveringh (1852), the European Central Bank, and various sources from Leiden, Haarlem, Utrecht, Schiedam, and Amsterdam. German data come from various sources from several German principalities. U.S. data come from Durand and Winn (1947), Homer and Sylla (2005), the NBER Macrohistory database, and Federal Reserve Economic Data (FRED) from the Federal Reserve Bank of St. Louis.

We also report the Bank of England (BoE) short-term lending rate (series BOERUKM) from FRED. From 1694 to 1971, the "bank rate" is used; from 1972 to 1981, the minimum lending rate is used; from 1981 to 1997, the BoE base rate is used; and from 1997 to the present, the BoE Operational interest rate is used. For more information see the Bank of England research datasets webpage.

Data for U.S. interest rates from 1984 to 2016 come from FRED. Our main measure for nominal interest rates in the U.S. is the effective Federal Funds Rate (series FEDFUNDS), the rate corresponding to the median volume of overnight unsecured loans between depository institutions. This is plotted in Panel A of Figure 2. In our calibration exercises, for comparability with Farhi and Gourio (2018) we use the one-year constant maturity Treasury rate, less current inflation. Data for U.K. interest rates from 1984 to 2016 come from Jordà et al. (2019), who in turn use data from Zimmermann (2017) and the Bank of England, who use the average rate on 3-month Treasury bills.

Data on U.S. inflation expectations come from FRED and the Survey of Professional Forecasters. From FRED, we use the inflation expectations from the Surveys of Consumers of University of Michigan (series MICH), which covers short-term inflation expectations, and the expected 10-year-ahead inflation implied from Treasury Inflation-Indexed Constant Maturity Securities (series T10YIE). From the Survey of Professional Forecasters, we use the 10-year ahead inflation expectations. These data are shown in Figure 5. Further, we use median one-year-ahead expected inflation from the Survey of Professional Forecasters to construct the deviation of expected inflation from realized inflation, shown in Figure 6.

Growth data come from different sources. In Tables 1–2, the U.S. growth parameter μ is set to the same values as in Farhi and Gourio (2018)[15], who use a composite growth parameter obtained by combining the growth in population, investment, total factor productivity, and the Cobb-Douglas production function parameter α (which they estimate). In Figure 2 and Table 3, we use real GDP growth rates from FRED (series GDPC1) as the growth rate for the U.S. When calibrating to the U.K. data, we use the real GDP growth series from Jordà et al. (2019).

Data on investment and capital stock come from the Bureau of Economic Analysis (BEA) Fixed Assets Accounts Tables. Investment data come from Table

[15]In particular, we set μ equal to the g_T parameter in Farhi and Gourio.

1.5, Line 2 and capital stock data come from Table 1.1, Line 2. In these data, investment as a fraction of capital averaged 7.7% from 1984–2000 and 6.9% from 2001–2016.

Price-dividend ratio data for the U.S. from 1984 to 2016 are from the Center for Research in Security Prices (CRSP). Specifically, we use cum-dividend returns (series VWRETD) and ex-dividend returns (series VWRETX). To calculate the price-dividend ratio, we back out prices and dividends from cum- and ex-dividend returns. This series is plotted in Panel B of Figure 2. We use this procedure to calculate our price-dividend ratio moments for the calibrations in Tables 1 and 2.

For the longer U.S. valuation data, we use prices and dividends on the S&P 500 from Shiller (2000). We also form the cyclically-adjusted price-earnings ratio (CAPE): the price divided by the average inflation-adjusted earnings from the previous 10 years. See (Shiller, 2000) and online data description. For the U.K. valuation data, we use data from Jordà et al. (2019). Jordà et al. aggregate total returns data from Grossman (2002) and from Barclays Equity Gilt Study.

Finally, we obtain the Volatility Index (VIX) series from the Chicago Board Options Exchange (CBOE). The CBOE calculates the risk-neutral expected 30-day quadratic variation using option prices. There are small differences in the calculation methodology over the years; see CBOE white paper.

B Derivations for Section 3: Endowment economy with rare disasters

B.1 Price-consumption ratio

Given the SDF (3), the Euler equation with respect to the consumption claim is

$$1 = \mathbb{E}_t\left[\beta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}} R_{W,t+1}^\theta\right]. \tag{B.1}$$

Conjecture a constant price-consumption ratio

$$\kappa \equiv (W_t - C_t)/C_t. \tag{B.2}$$

Substituting (B.2) into (B.1) and using $R_{W,t+1} = W_{t+1}/(W_t - C_t)$ implies

$$1 = \beta^\theta \mathbb{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{\theta(1-\frac{1}{\psi})} \left(\frac{\kappa+1}{\kappa}\right)^\theta\right]. \tag{B.3}$$

Given (2–1),

$$\frac{\kappa}{\kappa+1} = \beta e^{(1-\frac{1}{\psi})\mu}\left[1 + p((1-\eta)^{1-\gamma} - 1)\right]^{\frac{1}{\theta}}. \tag{B.4}$$

A solution exists provided that the right hand side of (B.4) is less than one. We restrict attention to parameter combinations satisfying this restriction. Finally,

$$\kappa = \frac{\beta e^{(1-\frac{1}{\psi})\mu}\left[1 + p((1-\eta)^{1-\gamma} - 1)\right]^{\frac{1}{\theta}}}{1 - \beta e^{(1-\frac{1}{\psi})\mu}\left[1 + p((1-\eta)^{1-\gamma} - 1)\right]^{\frac{1}{\theta}}}, \tag{B.5}$$

verifying the conjecture.

B.2 Riskfree rate

The riskfree rate is given by the Euler equation for the riskfree asset

$$R_f = \mathbb{E}_t\left[\beta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}} R_{W,t+1}^{\theta-1}\right]^{-1}. \tag{B.6}$$

This simplifies to

$$R_f = \mathbb{E}_t \left[\beta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(\frac{\kappa}{\kappa+1} \right)^{1-\theta} \right]^{-1}. \tag{B.7}$$

where $\kappa/(\kappa+1)$ is given by (B.4). Solving this yields the expression for the gross riskfree rate

$$R_f = \beta^{-1} e^{\frac{1}{\psi}\mu} \Big[1 + p((1-\eta)^{-\gamma} - 1) \Big]^{-1} \Big[1 + p((1-\eta)^{1-\gamma} - 1) \Big]^{\frac{\theta-1}{\theta}} \tag{B.8}$$

which implies that the log riskfree rate is given by

$$\begin{aligned} \log R_f = -\log\beta + \frac{1}{\psi}\mu - \log(1 + p((1-\eta)^{-\gamma} - 1)) \\ + \left(\frac{\theta-1}{\theta} \right) \log(1 + p((1-\eta)^{1-\gamma} - 1)). \end{aligned} \tag{B.9}$$

B.3 Interpretation of rising disaster probability

In Section 3.1, three parameters change across the two samples to affect the price-dividend ratio and the riskfree rate: the patience parameter β, the drift term in growth μ, and the probability of a disaster p. We now drill down to find out the role of each of the three. Table B.1 reports the results. The effect of the parameters on the price-dividend ratio can be decomposed into a term that depends only on the riskfree rate, a term that depends on the risk premium, and a term that depends on expected growth, analogous to the decomposition in Campbell and Shiller (1988):

$$\begin{aligned} \log \frac{\kappa}{\kappa+1} \approx -\underbrace{\Big(-\log\beta + \frac{1}{\psi}\mu - p((1-\eta)^{-\gamma} - 1) + \left(\frac{\theta-1}{\theta} \right) p((1-\eta)^{1-\gamma} - 1) \Big)}_{\text{riskfree rate effect}} \\ - \underbrace{\big(p\eta((1-\eta)^{-\gamma} - 1) \big)}_{\text{risk premium effect}} + \underbrace{(\mu - p\eta)}_{\text{cash flow effect}} \end{aligned} \tag{B.10}$$

First note that β only affects the price-dividend ratio through the riskfree rate. Greater patience lowers the riskfree rate, and thus the rate at which investors discount all future cash flows, raising the price-dividend ratio. Because the price-dividend ratio is a convex function of β, apparently small shifts in β cause massive changes.[16] One way to understand this result is through duration: when the price-dividend ratio discounts cash flows in the distant future, their valuation is more sensitive to small changes in rates than short term cash flows. Just to send the riskfree rate even half of the distance between the two samples would, therefore, send the price-dividend ratio soaring to nearly 100. Such a stark rise in valuation ratios reveals the fundamental problem with basing the decline in interest rates on an increased desire to save. Accounting for the decline in growth rates helps to lower the price-dividend ratio, provided that the EIS is greater than one.

[16]Specifically, the price-dividend ratio is a convex function of $\kappa/(\kappa+1)$, meaning that, at values of $\kappa/(\kappa+1)$ close to 1 (for which the price-dividend ratio is high), a small increase in β implies a very large increase in the price-dividend ratio.

Table B.1: Contribution of each parameter

This table starts with the model of Panel A in Table 1, calibrated to 1984–2000 and changes first the discount factor β, then average consumption growth μ, then the disaster probability p to obtain the calibration that matches 2001–2016. Risk aversion $\gamma = 12$ and consumption declines 15% in a disaster. In Panel A, the elasticity of intertemporal substitution (EIS) equals 2; in Panel B, the EIS equals 0.5. Parameters and yields are in annual terms.

Panel A: $\gamma = 12$, EIS $= 2$, $\eta = 0.15$					
	Parameter values			Targeted moments	
	β	μ	p	PD ratio	r_f
Baseline calibration (1984–2000)	0.967	0.0350	0.0343	42.34	0.0279
Higher β	0.979	0.0350	0.0343	94.74	0.0151
Higher β & lower μ	0.979	0.0282	0.0343	71.44	0.0117
Baseline calibration (2001–2016)	0.979	0.0282	0.0667	50.11	-0.0035
Panel B: $\gamma = 12$, EIS $= 0.5$, $\eta = 0.15$					
	Parameter values			Targeted moments	
	β	μ	p	PD ratio	r_f
Baseline calibration (1984–2000)	0.997	0.0350	0.0343	42.34	0.0279
Higher β	0.983	0.0350	0.0343	25.63	0.0428
Higher β & lower μ	0.983	0.0282	0.0343	31.27	0.0292
Baseline calibration (2001–2016)	0.983	0.0282	0.0667	50.11	-0.0035

In contrast, the growth rate μ enters into the price-dividend ratio in two ways, once multiplied by the multiplicative inverse of the EIS, representing its effect on the riskfree rate, and once multiplied by unity, representing its effect on future cash flows. A decrease in μ lowers the riskfree rate, following the usual consumption Euler equation intuition: the lower is expected growth, the greater the desire to save for the future, and hence the lower the riskfree rate must be. Or, in a production economy, the lower is growth, the lower the demand for borrowing, and the lower the riskfree rate. Either way, low interest rates and low growth clearly go together. However, when the EIS is above one, the effect of growth on the interest rate is small: unlike patience, the fall in growth lowers the riskfree rate, but raises the price-dividend ratio, so the cash flow effect dominates the riskfree rate effect.

Panel A of Table B.1 shows that accounting for both an increase in β and a decline in growth leads to a price-dividend ratio of about 70, not 50 as the data require. The remainder must be filled in by an increase in the risk premium (and a further decrease in expected future cash flows) through the disaster probability. That the increase in the disaster probability is concomitant with a decline in the riskfree rate helps the model even further.

When the EIS is below one, any decrease in the growth rate μ will lead to an increase in the price-dividend ratio, as the riskfree rate effect will dominate the cash flow effect. It will also, through the channel described above, lead to a decline in the riskfree rate, and the decline should be larger than that in the case of EIS greater than one. Suppose one starts with the demand-side view that investors have become more patient: β has risen. As Panel A shows, any increase in β sends the price-dividend ratio soaring (this effect is not mediated by the EIS, and so is present regardless of what side of unity the EIS is on). One then needs to change growth μ and/or the disaster probability p to match the fact that the price-dividend ratio did *not* soar. When the EIS is below one, however, a decrease in μ makes the problem worse in that it raises the price-dividend ratio still further. It is then impossible to match the data with p because, again, an EIS less than one means that increasing p lowers the riskfree rate and *raises* the price-dividend ratio. If one tries to bring the price-dividend ratio down with lower p, the riskfree rate rises. For this reason, matching the price-dividend ratio and the riskfree rate requires that investors be less, not more, patient.

To summarize: when the EIS is greater than one, one can reconcile the demand and supply intuition for lower interest rates, at the cost of assuming a higher probability of disaster. However, this reconciliation is fragile: it falls apart with an EIS less than one. Thus, accepting greater patience and lower growth as an explanation for the decline in interest rates requires accepting also that there is a higher probability of disaster and an EIS greater than one. While many models do assume an EIS greater than one, it is unsettling to have qualitative predictions of the model depend so heavily on a parameter falling within a certain range, for which there is little direct intuition or outside data support.

B.4 Calibrating to alternative moments

In this section, we note that the result of an increase in the the disaster probability from 3.43% to 6.67% from the first sample to the second sample is dependent on calibrating to the price-dividend ratio in the United States, which may be inflated by changes in the tendency of U.S. companies to pay dividends (Fama and French, 2001). If one calibrates the model to alternative moments, such as the cyclically-adjusted price-earnings ratio in the U.S., or to the U.K. price-dividend ratio, that do not have this issue, one needs a much greater increase in the disaster probability to match valuation ratios (Table B.2). In particular, calibrating to the CAPE ratio (U.K. price-dividend ratio) requires a rise in the disaster probability from 5.56% (1.34%) in the first sample to a quite large 10.1% (5.33%) disaster probability in the second sample.

Table B.2: Accounting for the data: alternative measures for valuations and rates

This table repeats the exercise of Panel B, Table 1, except that parameters are calibrated to the CAPE ratio rather than the price-earnings ratio (Panel A) and to U.K. data rather than U.S. data (Panel B).

		Values	
	Parameter	1984–2000	2001–2016
Panel A: Calibration to US CAPE ratio			
CAPE ratio	κ	25.97	26.73
Inflation-adjusted Treasury yield	y_b	0.0279	-0.0035
Average consumption growth	μ	0.0350	0.0282
Discount factor	β	0.957	0.968
Probability of disaster	p	0.0556	0.101
Panel B: Calibration to UK moments			
UK price-dividend ratio	κ	27.78	30.86
Inflation-adjusted Treasury bill yield	y_b	0.0500	0.0040
Average consumption growth	μ	0.0278	0.0156
Discount factor	β	0.955	0.971
Probability of disaster	p	0.0134	0.0533

B.5 Yield and expected return with sovereign default risk

Consider the defaultable short-term government bond paying $(1-L_{t+1})$ dollars—that is, 1 dollar in the case of no default and $1-\lambda\eta$ dollars in the case of default. The price of this claim is obtained by solving the Euler equation

$$Q_t = \mathbb{E}_t\left[\beta^\theta\left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}} R_{W,t+1}^{\theta-1}(1-L_{t+1})\right], \tag{B.11}$$

which simplifies to

$$Q_t = \mathbb{E}_t\left[\beta^\theta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\frac{\kappa}{\kappa+1}\right)^{1-\theta}(1-L_{t+1})\right], \tag{B.12}$$

where $\kappa/(\kappa+1)$ is given by (B.4). This gives the price of the defaultable claim as

$$Q_t = \beta e^{-\frac{1}{\psi}\mu}\left[1+p((1-\eta)^{1-\gamma}-1)\right]^{\frac{1-\theta}{\theta}}\left[1+p((1-\lambda\eta)(1-\eta)^{-\gamma}-1)\right]. \quad \text{(B.13)}$$

The yield on the defaultable claim is defined as $y_{b,t} \equiv -\log Q_t$, and is thus equal to the constant

$$y_b = \log R_f + \log\left(1+p((1-\eta)^{-\gamma}-1)\right) - \log\left(1+p((1-\lambda\eta)(1-\eta)^{-\gamma}-1)\right), \quad \text{(B.14)}$$

where $\log R_f$ is given by (B.9). The expected excess return on the bond is the expected payoff divided by the price, less the log riskfree rate, and therefore equals

$$\begin{aligned}\log \mathbb{E}_t\left[R_{b,t+1}\right] - r_f &= \log\left(1+p((1-\lambda\eta)-1)\right) \\ &\quad + \log\left(1+p((1-\eta)^{-\gamma}-1)\right) - \log\left(1+p((1-\lambda\eta)(1-\eta)^{-\gamma}-1)\right). \quad \text{(B.15)}\end{aligned}$$

Suppose instead of being subject to outright default, the bond is a nominally riskfree asset and so the government partially defaults through inflation. Assume inflation is given by the process (11). The price of this defaultable claim is obtained by solving the Euler equation

$$Q_t^{\$} = \mathbb{E}_t\left[\beta^{\theta}\left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}} R_{W,t+1}^{\theta-1}\frac{\Pi_t}{\Pi_{t+1}}\right], \quad \text{(B.16)}$$

which simplifies to $Q_t^{\$} = Q_t e^{-\mu_{\pi,t}}$ for the price Q_t given by (B.13). Subsequent results in the main text then follow straightforwardly.

C Volatility Index in a disaster economy

For tractability, we adapt the simple disaster model to continuous time, following Seo and Wachter (2019). Suppose consumption follows the jump-diffusion process

$$\frac{dC_t}{C_{t^-}} = \mu dt + \sigma dB_t + (e^{-z_t} - 1)dN_t, \tag{C.1}$$

where B_t is a standard Brownian motion, N_t is a Poisson process with constant intensity λ, and z_t has time-invariant distribution υ. As in Abel (1999) and Campbell (2003), we model dividends as levered consumption: $D_t = C_t^{\phi}$. Under both power utility and recursive preferences, it follows that the price of the claim to the dividend stream follows the process

$$\frac{dS_t}{S_{t^-}} = \mu_S dt + \phi\sigma dB_t + (e^{-\phi z_t} - 1)dN_t. \tag{C.2}$$

The quadratic variation is then given by

$$QV_{t,t+\tau} \equiv \int_t^{t+\tau} d[\log S, \log S]_s = \phi^2\sigma^2\tau + \int_t^{t+\tau} \phi^2 z_s^2 dN_s. \tag{C.3}$$

For risk-neutral measure Q, the VIX is then given by

$$\text{VIX}_t^2 \equiv \mathbb{E}_t^Q[QV_{t,t+\tau}] = \phi^2\left(\sigma^2 + \lambda\mathbb{E}_\upsilon\left[e^{\gamma z_t} z_t^2\right]\right)\tau, \tag{C.4}$$

where the last term follows from Girsanov's theorem:

$$\mathbb{E}_{t^-}^Q\left[\phi^2 z_s^2 dN_s\right] = \mathbb{E}_{t^-}\left[\frac{\pi_t}{\pi_{t^-}}\phi^2 z_s^2 dN_s\right] = \lambda\phi^2\mathbb{E}_\upsilon\left[e^{\gamma z_t} z_t^2\right]. \tag{C.5}$$

Note that these formulas hold for both time-additive utility and recursive preferences.

To calculate the implied VIX in the model, we choose parameters according to our calibration in Table 1: disaster size $z = -\log 0.85$, relative risk aversion coefficient $\gamma = 12$, consumption volatility $\sigma^2 = 0.02$, first sample disaster intensity $\lambda_1 = 0.03$, and second sample disaster intensity $\lambda_2 = 0.07$. These are annualized parameters, so $\tau = 1/12$ matches the time interval used to calculate the VIX. We then choose ϕ^2 such that (C.4) with λ_1 is equal to the empirically observed value 0.2056^2 in the first sample. Given this calibration—which implies $\phi^2 = 19.8$—we calculate that the implied VIX with $\lambda_2 = 0.07$, using this value of ϕ^2, is 23.36 compared to the empirical average of 20.66. Using Newey-West standard errors with two lags on the monthly VIX, the t-statistic on this test is 2.66.

D Production model

D.1 Solution to the no-inventory case

Consider the model in Section 4.1. The agent maximizes (22), subject to (21). Conjecture that

$$V(W_t) = \nu W_t, \tag{D.1}$$

for some constant $\nu > 0$. Substituting this conjecture into (22), with $R_{W,t+1} \equiv R_{f,t+1} + \alpha_t(R_{K,t+1} - R_{f,t+1})$ implies

$$(1-\beta)\log\nu + \log W_t = \max_{C_t,\alpha_t}\left\{(1-\beta)\log C_t + \beta\log\left(W_t - C_t\right) + \frac{\beta}{1-\gamma}\log\left(\mathbb{E}_t\left[R_{W,t+1}^{1-\gamma}\right]\right)\right\}. \tag{D.2}$$

At the optimum. the derivative of the right-hand side with respect to C_t equals zero. Thus:

$$\frac{1-\beta}{C_t} - \frac{\beta}{W_t - C_t} = 0$$

yielding the result $C_t/W_t = 1-\beta$. Setting the derivative of the right hand side with respect to α equal to zero yields (28).

D.2 Solution to the general case

The agent can invest in an inventory asset with net return $r_I = 0$, a riskfree bond with net return $r_{f,t+1}$, and a risky capital asset with net return $r_{K,t+1}$. Let $r_{j,t+1}$, $j \in \mathcal{J} = \{I, f, K\}$, represent net returns, and let $\alpha_{j,t}$ denote the percent allocation of savings to asset j. Note that, in our setting with a binary shock χ_{t+1}, markets are complete, so the agent will be able to construct any state-contingent portfolio return $r_{i,t+1}$. Inventory and capital are the only securities in positive net supply; furthermore, we restrict inventory to be in non-negative supply ($I_t \geq 0$). It follows from this setup that the return on wealth $R_{W,t+1} = \sum_{j\in\mathcal{J}} \alpha_{j,t}(1+r_{j,t+1})$, where $\sum_{j\in\mathcal{J}} \alpha_{j,t} = 1$.

Suppose that the agent has Epstein-Zin utility with unit EIS. The agent's optimization problem is therefore

$$\max_{C_t,\{\alpha_{j,t}\}_{j\in\mathcal{J}}} \left(C_t^{1-\beta} \left(\mathbb{E}_t \left[V(W_{t+1})^{1-\gamma} \right] \right)^{\frac{\beta}{1-\gamma}} \right), \tag{D.3}$$

subject to the dynamic budget constraint

$$W_{t+1} = (W_t - C_t)R_{W,t+1} = (W_t - C_t) \sum_{j\in\mathcal{J}} \alpha_{j,t}(1 + r_{j,t+1}), \tag{D.4}$$

the portfolio weight restriction

$$\sum_{j\in\mathcal{J}} \alpha_{j,t} = 1, \tag{D.5}$$

and the inventory non-negativity constraint

$$\alpha_{I,t} \geq 0. \tag{D.6}$$

Let ζ_t and ξ_t denote the Lagrange multipliers on the constraints (D.5) and (D.6), respectively.

Substituting (D.1) and the budget constraint (D.4) into (D.3), then taking logs, we again obtain (D.2) and the identical first-order condition for consumption as above. The first-order condition with respect to asset allocation $\alpha_{j,t}$, $j \neq I$, is

$$\beta \mathbb{E}_t \left[R_{W,t+1}^{1-\gamma} \right]^{-1} \mathbb{E}_t \left[R_{W,t+1}^{-\gamma}(1 + r_{j,t+1}) \right] = \zeta_t, \tag{D.7}$$

and the first-order condition with respect to the inventory allocation $\alpha_{I,t}$ is

$$\beta \mathbb{E}_t \left[R_{W,t+1}^{1-\gamma} \right]^{-1} \mathbb{E}_t \left[R_{W,t+1}^{-\gamma} \right] + \xi_t = \zeta_t. \tag{D.8}$$

Multiply both sides of (D.7) by $\alpha_{j,t}$, take the sum over $j \in \mathcal{J} \setminus \{I\}$, and substitute in (D.8) to see that

$$\zeta_t = \beta + \xi_t \alpha_{I,t} = \beta, \tag{D.9}$$

by complementary slackness. This implies the Euler equation for gross returns

$$\mathbb{E}_t \left[R_{W,t+1}^{1-\gamma}\right]^{-1} \mathbb{E}_t \left[R_{W,t+1}^{-\gamma} R_{j,t+1}\right] = 1 \tag{D.10}$$

and the Euler equation for inventory

$$\mathbb{E}_t \left[R_{W,t+1}^{1-\gamma}\right]^{-1} \mathbb{E}_t \left[R_{W,t+1}^{-\gamma}\right] + \frac{\xi_t}{\beta} = 1. \tag{D.11}$$

Note the market clearing condition $\alpha_{I,t} = 1 - \alpha_{K,t}$, where $\alpha_{K,t}$ is simply denoted α_t in our setup in the main text. We thus have that $\xi_t > 0$ if and only if $\alpha_t < 1$.

We now show formally that inventory imposes a zero lower bound. Throughout, we assume that the bond is in zero net supply.

Lemma 1. *If $\alpha_t < 1$, then the gross real riskfree rate $R_{f,t+1} = 1$. If $\alpha_t = 1$, then $R_{f,t+1} \geq 1$ and is equal to the real riskfree rate in a no-inventory economy $R^*_{f,t+1}$.*

Proof. If $\alpha_{I,t} > 0$, then $\xi_t = 0$ and (D.10) and (D.8) combine to give us $R_{f,t+1} = 1$. If $\alpha_{I,t} = 0$, then $\xi_t \geq 0$ and

$$R_{f,t+1} = \frac{\beta}{\beta - \xi_t}, \tag{D.12}$$

which is greater than or equal to 1. Moreover, if $\alpha_{I,t} = 0$, then market clearing implies $R_{W,t+1} = R_{K,t+1}$ and the Euler equation (D.10) yields

$$R_{f,t+1} = \mathbb{E}_t \left[R_{K,t+1}^{1-\gamma}\right] \mathbb{E}_t \left[R_{K,t+1}^{-\gamma}\right]^{-1}, \tag{D.13}$$

which is the same as the riskfree rate $R^*_{f,t+1}$ in the no-inventory economy. ∎

We next show that the unconstrained riskfree rate determines α.

Theorem 1. *If the unconstrained gross riskfree rate* $R^*_{f,t+1} < 1$*, then* $\alpha_t < 1$ *and the constrained riskfree rate* $R_{f,t+1} = 1$*. If* $R^*_{f,t+1} \geq 1$*, then* $\alpha_t = 1$ *and the equilibrium is as in a standard no-inventory production economy with* $R_{f,t+1} = R^*_{f,t+1}$*.*

Proof. We will prove the theorem by contradiction using Lemma 1.

Suppose $R^*_{f,t+1} < 1$ and $\alpha_{I,t} = 0$. Then $R_{f,t+1} = R^*_{f,t+1} < 1$, which contradicts Lemma 1. It must therefore be the case that $R^*_{f,t+1} < 1$ implies $\alpha_{I,t} > 0$, which implies $R_{f,t+1} = 1$.

Now suppose $R^*_{f,t+1} > 1$ and $\alpha_{I,t} > 0$. Then $R_{f,t+1} = 1 < R^*_{f,t+1}$, which contradicts Lemma 1. Moreover, in the knife-edge case $R^*_{f,t+1} = 1$, the equilibrium conditions (D.10) and (D.8) imply $\xi_t = 0$, which implies that $\alpha_{I,t} = 0$ and $R_{f,t+1} = R^*_{f,t+1} = 1$. Thus, it must be that $R^*_{f,t+1} \geq 1$ implies $\alpha_{I,t} = 0$, which implies $R_{f,t+1} = R^*_{f,t+1} \geq 1$. ■

We conjecture that the price-dividend ratio depends only on the current state χ_t (i.e., whether the disaster occurred or not). The intuition for this is that output growth Y_{t+1}/Y_t is a function of χ_t. Thus,

$$1 = \mathbb{E}_t \left[R^{1-\gamma}_{W,t+1}\right]^{-1} \mathbb{E}_t \left[R^{-\gamma}_{W,t+1} \left(\frac{\kappa^Y(\chi_{t+1}) + 1}{\kappa^Y(\chi_t)}\right)\left(\frac{Y_{t+1}}{Y_t}\right)\right]. \tag{D.14}$$

This implies that we have two equations, one for the non-disaster state,

$$\begin{aligned}\kappa^Y(0) = \hat{\beta}\Big[&(1-p)(1+\alpha r_{K,0})^{1-\gamma}(\kappa^Y(0)+1) \\ &+ p(1+\alpha r_{K,\eta})^{-\gamma}(\kappa^Y(\eta)+1)(1-\eta)(1+\alpha r_{K,0})\Big],\end{aligned} \tag{D.15}$$

and one for the disaster state,

$$\begin{aligned}\kappa^Y(\eta) = \hat{\beta}\Big[&(1-p)(1+\alpha r_{K,0})^{-\gamma}(\kappa^Y(0)+1)(1-\eta)^{-1}(1+\alpha r_{K,\eta}) \\ &+ p(1+\alpha r_{K,\eta})^{1-\gamma}(\kappa^Y(\eta)+1)\Big].\end{aligned} \tag{D.16}$$

In these equations, $\hat{\beta} \equiv \beta\left[(1-p)(1+\alpha r_{K,0})^{1-\gamma} + p(1+\alpha r_{K,\eta})^{1-\gamma}\right]^{-1}$, $r_{K,0} \equiv (1-\delta+A)-1$, and $r_{K,\eta} \equiv (1-\delta+A)(1-\eta)-1$. The solution to this system is

$$\kappa^Y(0) = \frac{\hat{\beta}}{1-\beta}\left[(1-p)(1+\alpha r_{K,0})^{1-\gamma} + p(1+\alpha r_{K,\eta})^{-\gamma}(1-\eta)(1+\alpha r_{K,0})\right], \tag{D.17}$$

$$\kappa^Y(\eta) = \kappa^Y(0)(1-\eta)^{-1}\frac{1+\alpha r_{K,\eta}}{1+\alpha r_{K,0}}. \tag{D.18}$$

Although the price-dividend ratio is state-dependent when the agent chooses to hold inventory, the risk premium is not. The risk premium at time t when the agent holds inventory is given by $\log \mathbb{E}_t[R^Y_{t+1}] - \log R_f$ where the expected return on the output claim is given by

$$\mathbb{E}_t[R_{Y,t+1}] = \mathbb{E}_t\left[\left(\frac{\kappa^Y(\chi_{t+1})+1}{\kappa^Y(\chi_t)}\right)\left(\frac{Y_{t+1}}{Y_t}\right)\right]. \tag{D.19}$$

If the expected return on the output claim is the same across states, than so is the risk premium, so we focus on the expected return on the output claim here. In the no disaster state, the expected return on the output claim is

$$\mathbb{E}_t[R_{Y,t+1}|\chi_t = 0] = \left(\frac{(1-p)\kappa^Y(0) + p\kappa^Y(\eta)+1}{\kappa^Y(0)}\right)\times \left(\beta(1-p\eta)\left(\alpha(1-\delta+A)+1-\alpha\right)\right) \tag{D.20}$$

and in the disaster state by

$$\mathbb{E}_t[R_{Y,t+1}|\chi_t = \eta] = \left(\frac{(1-p)\kappa^Y(0) + p\kappa^Y(\eta)+1}{\kappa^Y(\eta)}\right)\times \left(\beta(1-p\eta)\left(\alpha(1-\delta+A)+\left(\frac{1-\alpha}{1-\eta}\right)\right)\right). \tag{D.21}$$

Examining the two expressions, we see that the expected return in both states are the same provided that

$$\kappa^Y(\eta)(1-\eta)\Big(\alpha(1-\delta+A)+1-\alpha\Big) = \kappa^Y(0)\Big(\alpha(1-\delta+A)(1-\eta)+1-\alpha\Big).$$

Note that the terms inside the parentheses can be written as

$$\kappa^Y(\eta)(1-\eta)(1+\alpha r_{K,0}) = \kappa^Y(0)(1+\alpha r_{K,\eta})$$

which, after rearranging is identical to Equation (D.18). This implies that while the price-dividend ratio is time-varying, the risk premium is not.