Growth-Indexed Securities

Stavros Panageas
UCLA Anderson School of Management and NBER*

August 2025

Abstract

Several key issues in macroeconomics and finance center on the assumption that the risk-adjusted growth rate—computed under the "risk-neutral" measure—is smaller than the interest rate. ($E^Q(g) < r$). By studying a historical episode where government bonds were indexed to aggregate growth, I obtain a direct measurement of $E^Q(g)$ from the prices of these bonds. Two findings stand out: First, the risk compensation for bearing aggregate risk is moderate, and does not appear to be the driving force for the large equity premium observed in the data. Second, the risk compensation for bearing GDP risk is not large enough for the risk-neutral growth rate to exceed the interest rate. While these findings are problematic from the perspective of representative-agent models, they are consistent with models that emphasize agent heterogeneity.

Keywords: Growth-indexed Securities, GDP-linked bonds, Long-maturity dividend strips, Bubbles, Dynamic efficiency

*I thank Ravi Bansal, Lorenzo Garlappi, Martin Lettau, Erik Loualiche, Rodolfo Prieto, Dejanir Silva, Mindy Xiaolan and conference and seminar participants at Arizona State University, Cleveland Fed, Colorado Finance Summit, European Financial Management Association, Insead, London Business School, Oslo Center for Advanced Studies, SFS Cavalcade, University of British Columbia, University of California Los Angeles, Utah Winter Finance Conference, and Western Finance Association for useful comments. A previous version of the paper was circulated under the title "Of Trees and Beanstalks: A Paradox of Co-integration and the Sharpe ratio of Growth-indexed Bonds".

1 Introduction

Macro-finance models featuring an infinitely-lived, representative agent typically imply that (a) The equity premium reflects compensation for aggregate-endowment risk, and (b) The long-run, risk-adjusted growth rate of consumption is smaller than the risk-free rate, so that the present discounted value of aggregate consumption is finite (the "transversality condition"). The literature has shown that this second implications, namely the finiteness of the present value of aggregate consumption, is a sufficient condition for excluding bubbles on positive-supply assets.[1] If, on the other hand, the present discounted value of aggregate consumption is infinite, as it is in some seminal overlapping-generations models, then the value of government debt may not correspond to the present value of primary surpluses. More broadly, policies that result in inter-generational transfers (such as government debt issuance) may lead to Pareto improvements.[2]

At first pass, it seems obvious that the present value of aggregate consumption ought to be finite, because there exist finitely-valued assets, such as real estate, with dividends that are cointegrated with aggregate consumption. Because of cointegration, one might think that long-term "strips" to either consumption or the dividends of such assets should command similar risk-adjustment and discount rates; therefore if the value of land is finite, so should be the value of the consumption claim. Appealing as it may be, this argument is wrong. This paper starts by showing that a cash-flow process C_t may be co-integrated with a process X_t, and yet the price of strips to the first process diverge to infinity, while the strips to the second process converge to zero and are summable as the maturity of the strips grows. This somewhat counterintuitive result shows that one cannot rely on some basic and broad theoretical argument to conclude that the finiteness of the value of land implies the finiteness of a claim to aggregate consumption.

A more convincing approach to examine if the risk-adjusted growth rate of the economy exceeds the interest rate would be to directly examine the yields on securities that are explicitly indexed to aggregate growth, such as GDP-indexed bonds. While such securities

[1] See, e.g. Santos and Woodford (1997a).

[2] See e.g., Abel and Panageas (2022) and Bloise and Reichlin (2023).

were never traded in the US, the paper exploits a historical episode where such bonds were traded in France. The main finding is that the risk-adjustment for bearing GDP risk is modest; the implied Sharpe ratio of bearing "pure" GDP risk is around 0.4.[3] With such a modest value for the GDP risk premium, the risk-adjusted growth rate of the economy slightly exceeds the interest rate. In addition, combining the low compensation for GDP risk with the relatively low correlation between GDP growth and the stock market implies that only a modest fraction of the equity premium (about 40%) can be attributed to compensation for aggregate risk. The rest of the equity premium reflects compensation for risks orthogonal to GDP fluctuations, an observation that is also at odds with most representative agent models.

Having summarized the paper, I next provide a more detailed explanation of the points made above starting with a short overview of some theoretical issues. The main reason for outlining these theoretical issues is to illustrate that the finiteness of present value of aggregate consumption is not a trivial issue that can be derived from some high-level theoretical argument.

In representative-agent models, the transversality condition implies that the present value of aggregate consumption must be finite. By contrast, in overlapping generations economies, the transversality condition implies only that aggregate wealth should equal the present value of consumption of the *currently alive* generations. The present value of *aggregate consumption* from the present time to infinity may well be infinite, an outcome that is not uncommon in parameterizations of these models.[4] As shown by Santos and Woodford (1997b), when the present value of aggregate consumption is infinite, it is no longer necessary for the price of a positive supply asset to equal the present value of its dividends.

It is tempting to argue that the value of the consumption claim must surely be finite by using an argument based on co-integration: Take for instance a non-depreciating asset, such as land. Suppose that rental payments and aggregate consumption share the same

[3]This low value of the GDP-risk premium is consistent with the low "shadow valuation" of GDP-risk in US data, in a sense that is made more precise later in the paper.

[4]If the consumption of the current generations is a progressively declining fraction of future aggregate consumption, then it may well be that the present value of consumption of the currently alive agents is finite, while the present value of aggregate consumption isn't.

stochastic trend (this is an implication of co-integration). But then, it should be the case that the long-run "strips" to consumption are discounted with the same discount rates as the aggregate dividends of land. Because the value of land is finite, the strips to the dividend payments of land must be summable; and since consumption strips must exhibit the same asymptotic behavior as the strips of land, consumption strips must be summable and the value of the consumption claim must be finite.

Unfortunately, this simple argument is erroneous: The paper shows that two cash-flow processes C_t and X_t may well be co-integrated, but the price of a strip to the cash flow C_T grows without bound as $T \to \infty$, while the price of a strip to the cash flow X_T goes to zero. Specifically, if there exists some portfolio of assets whose expected logarithmic rate of return exceeds the average logarithmic growth rate of C_t, then it is always possible to start a fund (with finite wealth) that pays out a cash flow X_t, with the property that X_t and C_t are co-integrated, and yet, the present value of C_t may well be infinite. An implication of this construction is that the time 0 value of a "strip" to the cash flow C_T diverges to infinity as $T \to \infty$, while the time 0 value of a "strip" to its co-integrated cash flow X_T converges to zero as $T \to \infty$. Such a payout process X_t always exists, whenever the available assets reward some risk(s) that are orthogonal to the risks impacting C_t, or, phrased differently, whenever the shocks affecting C_t are not the only priced risks in the economy.[5]

To summarize, it is not straightforward to use some high level argument to guarantee that the value of the aggregate endowment (and by implication aggregate consumption) is finite. Relatedly, one cannot simply use the unlevered equity premium to risk-adjust long run consumption strips, since the stock market return may well reward risks orthogonal to GDP (this is exactly the type of situation that can occur in OLG or heterogeneous-agents models where distributional risks either between, or within cohorts are priced.)

Because of these limitation of stock-market data for discounting claims to the aggregate endowment, one would ideally like to infer the discount rates for GDP-growth strips from

[5]Another way to argue the finiteness of the aggregate endowment and preclude bubbles comes from the macro-economics literature. Abel et al. (1989) argue that the US economy appears "dynamically efficient," i.e., the expected logarithmic return on capital exceeds the logarithmic growth rate of the economy. However, as Abel and Panageas (2022) show, ins a stochastic economy it is well possible that the economy is dynamically efficient, and yet the present value of consumption is infinite.

contracts whose payoffs are contingent on aggregate gdp growth (such as GDP futures or GDP-linked bonds.)

While there have been many proposals to introduce such bonds into financial markets for developed economies, no active markets developed for trading such claims. In the few cases where such bonds were introduced, it was in the aftermath of default in developing economies. Scarred by the recent experience of default in those countries, market participants did not embrace these bonds.

One important exception, however, was the historical experience of France in the late 1950s and early 1960s. In 1956 France introduced industrial-production-contingent bonds. The index of industrial production was essentially perfectly correlated with the real GDP growth of France, so in essence these bonds were GDP-linked bonds. They had a maturity of 15 years and their coupons were contingent on the evolution of the index of industrial production in France. The historical financial database of France provides biweekly prices for these bonds.

By observing the prices of these bonds one can estimate the implied growth rate of GDP under the risk-neutral measure. By comparing this number to the average (logarithmic) growth rate of GDP and dividing by its volatility, one can also obtain the associated Sharpe ratio for "pure GDP" risk. The main conclusion of this empirical exercise is that the market price for pure GDP risk is not particularly large; the baseline estimate is 0.39 (excess return per unit of standard deviation). This value is remarkably stable across various splits of the sample and various volatility specifications.[6]

There are two conclusions that can be drawn from this implied market risk of GDP growth. First, if one accepts that the baseline estimate of 0.39 is indicative of the risk com-

[6]As a plausibility check for this finding, I also use US data and perform the following thought experiment: The fact that GDP-contingent bonds were many times proposed for the US, but never broadly embraced, suggests that market participants must have suspected that these claims would not have mattered for investors; they would not have materially impacted their marginal utility, and thus would not have changed the asset prices of existing securities. Under this supposition, it becomes possible to infer the "shadow" market price of risk for GDP risk from the covariance of GDP growth with the mean-variance efficient portfolio. Using various combinations of test assets to build variations of the mean-variance portfolio, I arrive at the same conclusion as with the direct estimation of GDP risk from the French data: The market price of risk for GDP risk is not particularly large. Even taking the estimation errors into account (by performing a bootstrap) the market price of GDP risk has a mean value around 0.2 and rarely exceeds 0.4 in the bootstrapped samples.

pensation for pure GDP risk, then it appears that a good fraction of the equity premium does not reflect compensation for aggregate GDP risk, especially after taking into consideration the relatively low correlation between stock returns and gdp growth. This can be seen by the following "back-of-the-envelope" exercise: Suppose that GDP futures (normalized to have unit standard deviation) were traded and the Sharpe ratio of those futures were 0.39. Suppose that one used GDP futures to create a portfolio that invests in the stock market and shorts an appropriate amount of GDP-futures to neutralize the exposure of the portfolio to GDP risk. If the volatility of the stock market is 16 percent annually, the correlation between the stock market and GDP growth is 0.4, the Sharpe ratio of GDP risk is 0.39, and the equity premium is 6%, then the excess return of the GDP-risk-hedged portfolio would be 6% - $0.4 \times 0.16 \times 0.39 = 3.5\%$. This suggests that other risks (capital depreciation risks, investment specific shocks, displacement risks, distributional risks) play an important role in accounting for the equity premium. This finding implies that the stock market return cannot be viewed as a levered investment in the aggregate endowment. If that were true, then the GDP-hedged return on the stock market should be zero (and GDP growth would be perfectly correlated with the stock market return). The finding also implies that the return on capital cannot be used to ascertain whether the value of the aggregate endowment is likely to be finite, because the stock market rewards risks that are orthogonal to aggregate-growth risks.

Second, the empirical plausibility of the assumption that the present value of the aggregate endowment is finite is far from obvious: In the historical experiment that we examine, the expected growth rate under the risk-neutral measure ("risk adjusted growth rate") exceeded the interest rate over that sample. Part of the reason was the aforementioned observation that the compensation for bearing GDP risk was not particularly large. If one were to assume that this low compensation for GDP risk is indicative of the magntiude of risk compensation that would be encountered if the US were to issue GDP-indexed bonds, then it is tempting to perform the following thought experiment: In post WW-II data the US annual real GDP growth rate has been approximately 3% with a volatility of approximately 2.3%. If the Sharpe ratio for GDP-growth risk is 0.39, the risk-adjusted growth rate is 3%

-0.39×2.3% = 2.1%. By comparison, the (par) yields on either 20- or 30-year TIPS have only rarely exceeded 2%. Therefore, if the French experience with GDP-linked bonds is indicative of the order of magnitude of "pure" GDP risk compensation, it is far from clear whether the conditions that ensure the finiteness of the aggregate endowment are satisfied or not.

Aside from the theoretical issues that hinge on the gap between the risk-adjusted growth rate and the interest rate, there is a simple and practical dimension of the paper's findings. Assuming a 100% debt to GDP ratio, the gap between the risk-adjusted growth rate and the interest rate is equal to the primary deficit (as a fraction of GDP) that will *deterministically* ensure no further rise of the debt-to-GDP ratio if the government were to replace nominal with GDP-indexed bonds. The reason for using the word "deterministically" is that while GDP-contingent bods make the coupon and principal payments uncertain in units of dollars, they make the payments certain when GDP is used as the "numeraire". In turn, this means that for any given assumption on the deficit as a fraction of GDP, it becomes possible to exactly map out the evolution of debt as a fraction of GDP.

1.1 Literature review

The notion that $r < g$ is associated with the possibility that government debt does not have to equal the present value of surpluses dates back to the deterministic model of Diamond (1965). Cass (1972) analyzes the notion of "dynamic efficiency." An economy is dynamically efficient if it is impossible to raise aggregate consumption today (and reduce capital accumulation) without lowering aggregate consumption in some future period. The main result in Cass (1972) is that an economy (that possesses a steady state) is dynamically efficient if the logarithmic return on capital exceeds the logarithmic growth rate of capital (or equivalently the growth rate of output, since the capital-to-output ratio is constant in steady state.) In a later section, I provide a brief summary of the arguments used in Cass (1972) to arrive at his main conclusion. Abel et al. (1989), Zilcha (1990) and Zilcha (1991) extend Cass's analysis to stochastic environments. Abel et al. (1989) shows that if the profit share in an economy exceeds the investment share in every period, then the economy is dynamically efficient. Zilcha (1991) shows that the Cass criterion (that the logarithmic rate of return

on capital exceed the logarithmic rate of growth) applies also in a dynamic setup, except that now the comparison applies to the expected logarithmic return on capital and aggregate growth respectively.[7]

In a deterministic economy, the issue of dynamic efficiency and the possibility of bubbles are tightly linked, mainly because the rate of return on capital and the risk-free rate of return are the same. The centrality of the "$r - g$" comparison for the possibility of bubbles dates back at least to Tirole (1985). Santos and Woodford (1997b) show that the key issue is not whether the risk free rate exceeds the growth rate, but rather whether the present value of the aggregate endowment is finite. (In a later section I provide a short, self-contained argument for why the finiteness of the present value of the aggregate endowment plays a key role.) Clearly, in a deterministic setup the rate of return on capital is equal to the risk free rate, which is also the appropriate rate to discount the aggregate endowment. This is the reason why both the dynamic efficiency conditions and the finiteness of the value of the aggregate endowment are equivalent to $r - g > 0$. Of course, in a stochastic economy different cash flows command different discount rates. Therefore there is no longer an equivalence between dynamic efficiency on the one hand and the possibility of bubbles (and the closely related notion of interim Pareto inefficiency) on the other. (Bertocchi (1991), Binswanger (2005), Barbie et al. (2007), Bloise and Reichlin (2023), Abel and Panageas (2022).) Intuitively, the reason is that an economy may be Pareto inefficient due to the lack of efficient intergenerational risk-sharing, not capital over-accumulation. Specifically, even if it is impossible to find interventions that (weakly) increase aggregate consumption in every period, it may still be possible to find interventions that increase every generation's welfare, simply by improving the sharing of aggregate consumption between generations.[8,9]

The possibility of rolling debt forever without primary surpluses is fundamentally the

[7]Rangazas and Russell (2005) and Barbie and Kaul (2009) criticize some results in Zilcha (1991), which require stronger assumptions than those used by Zilcha (1991). However, the conclusion that the economy is dynamically efficient when the expected logarithmic return on capital exceeds the logarithmic growth rate of the economy, which is the main result used in this paper, remains valid.

[8]There are other approaches to drive a wedge between rates of return and thus allow for the possibility of bubbles, as in Farhi and Tirole (2011), Martin and Ventura (2012), Ball and Mankiw (2023).

[9]One important aspect of this paper is the usage of risk-adjusted growth rates. Lucas (2012) highlights the importance of using risk-adjustments in public finance.

same issue as a bubble on government debt. It should come as no surprise that the secular decline in interest rates and the rising levels of deficits has brought the possibility of such debt rollover to the fore. Blanchard and Weil (2001), Blanchard (2019), Barro (2023), Cochrane (2021), Reis (2021), Hellwig (2021), Kocherlakota (2023a), Kocherlakota (2023b), Jiang et al. (2019), Brunnermeier et al. (2021), Abel and Panageas (2022), Aguiar et al. (2021), Amol and Luttmer (2022) are recent examples of this growing literature.[10] A simple way to see the connection between the (somewhat more theoretical) literature on bubbles and the (somewhat more applied) literature on government debt is as follows: If the risk-adjusted growth rate of the economy (the growth of the economy under the risk-neutral measure) is above the risk-free rate on debt, then (a) the value of the aggregate endowment is infinite and (b) the yield on a (fictitious) GDP-growth-contingent bond is negative as well. In other words, the government could issue a liability in the form of a growth-contingent bond, and be certain that the value of this liability as a fraction of GDP would be lower in the future (due to its negative yield.) This deterministic decline in the value of the liability between two periods of time opens up the possibility of permanent rollover.

1.2 Outline

Section 2 introduces the notion of a co-integrated, cash-flow mimicking fund. In particular, this sections shows that as long as the logarithmic return of a portfolio of assets exceeds the logarithmic growth rate of some target process C_t, then — starting with finite initial wealth — it is always possible to finance a cash flow X_t that is co-integrated with C_t. Section 3 focuses on a complete market, so that there is no ambiguity on how to "price" cash flows and shows that the yields on "dividend" strips to C_t and X_t may be entirely different and not even converge to the same number asymptotically. Section 4 discusses the implications of Section 3 for dynamic efficiency and the finiteness of the aggregate endowment. Section 5 discusses the French experience with growth-contingent bonds. All proofs are contained in the appendix with the exception of one proof that is instructive.

[10] Indicative earlier contributions to this important topic include Aiyagari and McGrattan (1998), Blanchard and Weil (2001), Ball et al. (1998), Bohn (1995).

2 Co-integrated cash-flow mimicking funds

This section introduces the concept of a cash-flow tracking fund. To start, consider a given ("target" cash flow) process C_t with log-normal dynamics

$$\frac{dC_t}{C_t} = gdt + \sigma_C dB_t^C, \tag{1}$$

where g is a constant, and B_t^C is a standard Brownian motion.

Suppose that there exists some fund that can invest in N assets with a vector of rates of return equal to

$$\underbrace{dR_t}_{N\times 1} = \underbrace{\mu}_{N\times 1} dt + \underbrace{\sigma}_{N\times M}\underbrace{dB_t}_{M\times 1}, \text{ where } dB_t \equiv \begin{bmatrix} dB_t^C \\ d\overrightarrow{B}_t \end{bmatrix}, \tag{2}$$

where μ is an $N \times 1$ vector of expected returns, and σ is and $N \times M$ matrix of exposures to the M Brownian motions driving the returns. Assume that $N \leq M$ and that σ has rank N. Accordingly, the market could be either complete ($N = M$) or incomplete ($N < M$.)

The fund can make distributions equal to X_t per unit of time dt, so that the dynamic evolution of the fund's value,W_t, is

$$\frac{dW_t}{W_t} = \left(r + w'\mu^e - \frac{X_t}{W_t}\right) dt + w'\sigma dB_t \tag{3}$$

where r is the risk-free rate, $\mu^e \equiv \mu - r1_N$ is is an $N \times 1$ column vector of expected *excess* returns, and w is an $N \times 1$ vector of the fund's portfolio holdings.

Let $z_t \equiv \log W_t - \log C_t$ denote the fund-value-to-cash-flow ratio. Ito's Lemma implies that

$$dz_t = \left(r + w'\mu^e - g - \frac{w'\sigma\sigma'w - \sigma_C^2}{2} - \frac{X_t}{W_t}\right) dt + w'\sigma dB_t - \sigma_C dB_t^C. \tag{4}$$

Suppose that the fund makes distributions according to the following rule:

$$\frac{X_t}{W_t} \equiv \max(\alpha + \beta z_t, \varepsilon), \text{ where } \beta > 0 \text{ and } \varepsilon > 0. \tag{5}$$

In other words, the fund chooses a higher distribution rate, $\frac{X_t}{W_t}$, when z_t, is high, and a lower distribution rate when z_t is low. The truncation at some $\varepsilon > 0$ ensures that the fund never makes negative distributions. Substituting (5) into (4) results in the dynamics

$$dz_t = \left(r + w'\mu^e - g - \frac{w'\sigma\sigma'w - \sigma_C^2}{2} - \max(\alpha + \beta z_t, \varepsilon) \right) dt + w'\sigma dB_t - \sigma_C dB_t^C. \tag{6}$$

Next define the following quantity, which is the difference between the logarithmic rate of rate of return on the portfolio and the logarithmic growth rate of C_t:

$$A = A(w) \equiv r - g + w'\mu^e - \frac{w'\sigma\sigma'w}{2} + \frac{\sigma_C^2}{2}. \tag{7}$$

The following Proposition shows that the when $A(w) > \epsilon > 0$, z_t has a stationary distribution.

Proposition 1 *Define* $\phi \equiv w'\sigma - [\sigma_C, 0, .., 0]$ *and let* $\sigma_Z \equiv \sqrt{\phi'\phi}$ *and* $z^* \equiv \frac{\varepsilon - \alpha}{\beta}$. *Assume that* $\sigma_Z > 0$. *If* $A(w) > \varepsilon$, *then* z_t *has a stationary distribution* $f(z)$ *given by*

$$f(z) = \frac{1_{z \leq z^*} e^{\frac{2}{\sigma_z^2}(A-\varepsilon)z} + 1_{z > z^*} e^{\frac{2}{\sigma_z^2}\left[-\beta\frac{z^2-(z^*)^2}{2} + (A-\alpha)z - (\varepsilon-\alpha)z^*\right]}}{\int_{-\infty}^{+\infty} 1_{z \leq z^*} e^{\frac{2}{\sigma_z^2}(A-\varepsilon)z} + 1_{z > z^*} e^{\frac{2}{\sigma_z^2}\left[-\beta\frac{z^2-(z^*)^2}{2} + (A-\alpha)z - (\varepsilon-\alpha)z^*\right]} dz}. \tag{8}$$

Figure 1 illustrates $f(z)$. It is a concatenation of an exponential distribution (decaying at the rate $A(w) > 0$) for values of $z \leq z^*$ and a normal distribution for $z > z^*$.

The two assumptions of Proposition 1 are that (a) $\sigma_Z > 0$ and (b) $A(w) > \varepsilon > 0$. Assumption (a) implies that the return of the portfolio w is not perfectly correlated with the increments $\frac{dC}{C}$. Such a portfolio doesn't exist if the market is incomplete ($N < M$) since the matrix σ has rank N. But even if the market is complete ($N = M$), the proposition focuses on portfolios that are not perfectly correlated with $\frac{dC}{C}$, so that the logarithmic difference

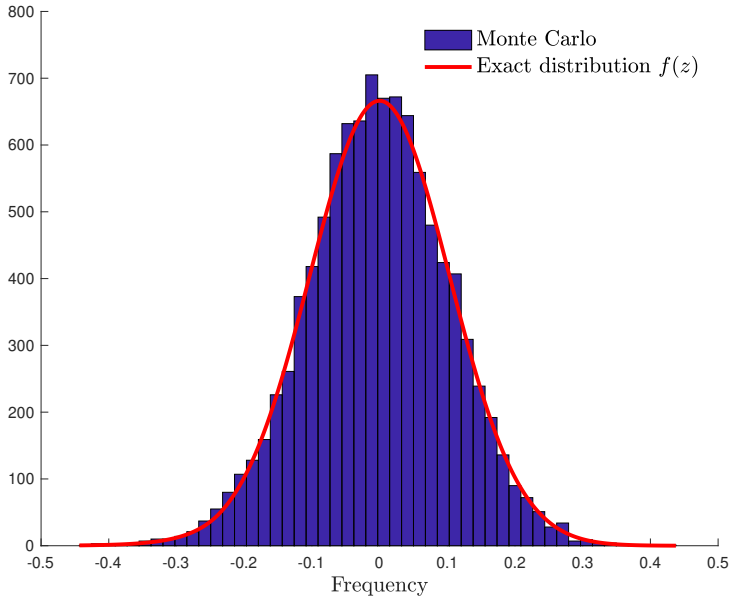


Figure 1: Simulated and exact distribution of the log fund-to-cash-flow value, z_t. Parameters are $A(w) = 0.08, \varepsilon = 0.01, \alpha = 0.08, \beta = 0.3, \sigma_Z = 0.08$.

between the fund and the cash flow (the "tracking error") has non-zero standard deviation, σ_Z. Assumption (b) is the more substantive assumption. In effect it requires that $A(w) > 0$. (The reason is that ε is arbitrary, and therefore it is always possible to find some $\varepsilon \in (0, A)$ as long as $A(w) > 0$.) Proposition 1 asserts that under conditions (a) and (b), the process z_t is stationary.

The intuition behind Proposition 1 is that whenever $z \leq z^*$, the logarithmic growth rate of the fund becomes $r + w'\mu^e - \frac{w'\sigma\sigma'w}{2} - \varepsilon$, whereas the logarithmic growth rate of the cash flow C_t, is $g - \frac{\sigma_C^2}{2}$. The assumption that $A(w) > \varepsilon$ ensures that the logarithmic growth rate of the fund exceeds that of the cash flow C_t. As a result the drift of z_t is positive. Vice vera, when z_t exceeds z^*, the drift of z_t becomes negative for high enough z_t : the payout, $\frac{X_t}{W_t}$, is linear in z_t and eventually exceeds the logarithmic return of the fund, $r + w'\mu^e - \frac{w'\sigma\sigma'w}{2}$ for high enough values of z_t. The positive drift for low z_t, which turns negative for high z_t is the reason for the mean reversion in z_t.

Note that if $A(w)$ were smaller than zero, then it would be impossible to find a non-negative payout X_t, such that z_t would be stationary. Indeed, in that case $\lim_{T\to\infty} z_t = -\infty$

a.s. for any payout policy $X_t \geq 0$.[11]

Maintaining the assumption that $A(w) > \varepsilon > 0$, then the (log) distributions made by the fund, X_t are co-integrated with the (log) cash-flow process C_t, since

$$\log X_t - \log C_t = \log \frac{X_t}{W_t} - \log \left(\frac{C_t}{W_t} \right) = \log \left(\max(\alpha + \beta z_t, \varepsilon) \right) + z_t, \tag{9}$$

where the last equation follows from (5) and the definition of z_t. Since z_t is stationary, equation (9) implies that $\log X_t - \log C_t = \log \left(\max(\alpha + \beta z_t, \varepsilon) \right) + z_t$ is stationary, or equivalently that the cash flow processes $\log X_t$ and $\log C_t$ are co-integrated.

3 Long-dated strips

The key takeaway from the previous section is that it is always possible to find a payout strategy, X_t, that is co-integrated with C_t, as long as there exists a portfolio w such that $A(w) > 0$. It is important to note that the portfolio w that finances the cash-flow X_t is not a "replicating" portfolio. The assumption $\sigma_Z > 0$ of Proposition 1 implies that $\log X_t$ and $\log C_t$ are co-integrated, but not perfectly correlated, irrespective of whether the market is complete $(N = M)$ or incomplete $(N < M)$.

This section shows that even though the cash-flow processes $\log X_t$ and $\log C_t$ are co-integrated, the value of long-dated divided "strip" to the cash flow X_T at time T will converge to zero as $T \to \infty$, even though the value of of a dividend-strip to C_T may diverge to infinity. Moreover, even if the value of both strips converges to zero as $T \to \infty$, the yields of these strips will in general not converge to the same number as $T \to \infty$. The remainder of this section formalizes these claims.

To remove any ambiguity about how to compute the value of dividends, the remainder of this section focuses on the case $N = M$, so that the market is complete. Define the unique

[11] The easiest way to see this is to let $z_t^{(0)}$ denote the process for z_t assuming that $X_t = 0$ and note that for any $X_t \geq 0$ we have $z_t^{(0)} \geq z_t$. The properties of the Brownian motion with negative drift impliy that $-\infty = \lim_{T \to \infty} z_t^{(0)} \geq \lim_{T \to \infty} z_t$.

stochastic discount factor as

$$\frac{dm_t}{m_t} = -rdt - \underbrace{\kappa}_{N\times 1} \left[\begin{array}{c} dB_t^c \\ d\overrightarrow{B}_t \end{array} \right], \tag{10}$$

where r is the prevailing interest rate and the vector $\kappa \equiv \sigma^{-1}\mu^e$ is the vector of the "market prices of risk."

The price $P_t^{(X_T)}$ of a dividend strip to the cash flow X_T,divided by the current level of X_t is defined as

$$\frac{P_t^{(X_T)}}{X_t} \equiv E_t\left(\frac{m_T}{m_t}\frac{X_T}{X_t}\right) = E_t^Q\left(e^{-r(T-t)}\frac{X_T}{X_t}\right), \tag{11}$$

where Q is the "risk-neutral" measure associated with the SDF (10). Because the integral $\int_t^\infty P_t^{(X_u)}du$ is bounded above by the time-t value of the fund, W_t, it must be the case that $\lim_{T\to\infty} P_t^{(X_T)} = 0$.[12] Similarly, the value of a long-dated dividend-strip to the cash flow C_t is given by

$$\frac{P_t^{(C_T)}}{C_t} \equiv E_t\left(\frac{m_T}{m_t}\frac{C_T}{C_t}\right) = E_t^Q\left(e^{-r(T-t)}\frac{C_T}{C_t}\right) = e^{-\left(r-g+\kappa^{(1)}\sigma_C\right)(T-t)}, \tag{13}$$

where $\kappa^{(1)}$ is the first element of the vector κ (i.e, the price of risk associated with the

[12]Note that under the probability measure Q, equation (3) becomes

$$dW_t = (rW_t - X_t)\,dt + \sigma W_t dB_t^Q. \tag{12}$$

Multiplying both sides of (12) with e^{-rt} and taking expectations gives the following present-value relation:

$$W_t = E_t^Q\int_t^T e^{-r(u-t)}X_u du + E_t^Q e^{-r(T-t)}W_T.$$

Since $W_T > 0$ for all T, it follows that the initial value of the fund, W_t, is an upper bound to $E^Q\int_t^T e^{-r(u-t)}X_u du$ for any T; in particular, as $T\to\infty$

$$W_t \geq E_t^Q\int_t^\infty e^{-r(u-t)}X_u du.$$

Since X_t is non-negative, the fact that $E_t^Q\int_t^\infty e^{-r(u-t)}X_u du$ is bounded above by W_t leads to the conclusion that $\lim_{T\to\infty} E_t^Q e^{-r(T-t)}X_T = 0$.

Brownian motion B_t^C). The quantity $g - \kappa^{(1)}\sigma_C$ can be thought of as the "risk-adjusted" growth rate of C_t. In particular, if $r < g + \kappa^{(1)}\sigma_C$, then $\lim_{T\to\infty} \frac{P^{(C_T)}}{C_t} = \infty$. The conditions of Proposition 1 may hold, even though $r < g - \kappa^{(1)}\sigma_C$, as the next proposition asserts.

Proposition 2 *There exists (an open set of) parameters such that (a)* $r - g + \kappa^{(1)}\sigma_C < 0$ *and (b) the conditions of Proposition 1 both hold. Therefore,* $\log \frac{X_T}{C_T}$ *is stationary, and yet* $\lim_{T\to\infty} P_t^{(C_T)} = \infty$, *while* $\lim_{T\to\infty} P_t^{(X_T)} = 0$.

Proposition 2 presents a conundrum. One would expect that the common stochastic trend of X_t and C_t would dominate the variation of both claims in the long run. As a result, one would expect the same yields on long term strips to either cash flow. Yet, Proposition 2 shows that the price of a dividend strip to C_T diverges to infinity as T grows, whereas the strip to X_T goes to zero. How is it possible that the dividend strips of two co-integrated cash flows can differ so dramatically in the long run?

To provide some perspective on the conundrum, define

$$f(z_t) \equiv \exp\left\{\log\left(\max(\alpha + \beta z_t, \varepsilon)\right) + z_t\right\}. \tag{14}$$

and note that equations (11) and (9) and (13) imply that

$$\begin{aligned}
\frac{1}{T-t}\log\left(\frac{P_t^{(X_T)}}{X_t}\right) &= \frac{1}{T-t}\log E_t^Q\left(e^{-r(T-t)}\frac{X_T}{X_t}\right) = \\
&= \frac{1}{T-t}\log\left\{e^{-r(T-t)}E_t^Q\left(\frac{C_T}{C_t}\right) \times E_t^Q\left(\frac{C_T}{E_t^Q\left(C_T\right)}\frac{f(z_T)}{f(z_t)}\right)\right\} \\
&= \frac{1}{T-t}\left[\log\left(\frac{P_t^{(C_T)}}{C_t}\right) + \log E_t^Q\left(\frac{C_T}{E_t^Q\left(C_T\right)}\frac{f(z_T)}{f(z_t)}\right)\right].
\end{aligned}$$

Taking limits as $T - t \to \infty$ gives

$$\lim_{T-t\to\infty}(-1) \times \frac{1}{T-t}\log\left(\frac{P_t^{(X_T)}}{X_t}\right) = -\lim_{T-t\to\infty}\frac{1}{T-t}\log\left(\frac{P_t^{(C_T)}}{C_t}\right) - \mathcal{R}, \tag{15}$$

where

$$\mathcal{R} \equiv \lim_{T-t\to\infty} \frac{1}{T-t} \log E_t^Q \left(\frac{C_T}{E_t^Q\left(C_T\right)} \frac{f(z_T)}{f\left(z_t\right)} \right).$$

Equation (15) implies that the yields on the long-term dividend strips to X_T and C_T are not equal except if the term $\mathcal{R}$ is equal to zero. To study this term, note that $\frac{C_T}{E_t^Q(C_T)}$ is positive and has an expectation equal to one; therefore it is a likelihood ratio that can be used to define a new probability measure, $\widetilde{Q}$, equivalent to Q (a Radon-Nikodym derivative). Specifically, we have the following result

Lemma 1 *Define the change of measure* $\frac{d\widetilde{Q}}{dQ} = \frac{C_T}{E_t^Q(C_T)}$. *Under the measure* $\widetilde{Q}$ *the term* $\mathcal{R}$ *can be expressed as*

$$\mathcal{R} = \lim_{T-t\to\infty} \frac{1}{T-t} \log E_t^{\widetilde{Q}} \left(f(z_T) \right), \tag{16}$$

and the dynamics of z_t *under the measure* $\widetilde{Q}$ *are given by*

$$dz_t = \left(\widetilde{A}(w) - \max(\alpha + \beta z_t, \varepsilon) \right) dt + \left\{ w'\sigma - \left[\sigma_C, 0_{1\times(N-1)} \right]' \right\} dB_t^{\widetilde{Q}}, \tag{17}$$

where $dB_t^{\widetilde{Q}}$ *is an* $N-$*dimensional Brownian motion under the probability measure* $\widetilde{Q}$, *and* $\widetilde{A}(w)$ *is given by*

$$\widetilde{A}(w) = r - \left(g - \kappa^{(1)}\sigma_C \right) - \frac{\sigma_C^2}{2} - \frac{1}{2} w'\sigma\sigma' w + w'\sigma \left[\sigma_C, 0_{1\times(N-1)} \right]'. \tag{18}$$

Corollary 1 *The process* z_t *is stationary under the probability measure* $\widetilde{Q}$ *iff* $\widetilde{A}(w) > \varepsilon$ *for some* w.

Lemma 1 implies that if z_t is a stationary process under the probability measure $\widetilde{Q}$, and $\lim_{T-t\to\infty} E_t^{\widetilde{Q}} \left(f(z_T) \right) < \infty$, then $\mathcal{R} = 0$. To provide conditions so that $\mathcal{R} = 0$, note first that the portfolio that maximizes $\widetilde{A}(w)$ is given by

$$w = \left(\sigma\sigma'\right)^{-1} \sigma \left[\sigma_C, 0 \right]', \tag{19}$$

where 0 is an $(N-1)\times 1$ vector. (In the special case where $N=1$, the portfolio w is a scalar given by $\frac{\sigma_C}{\sigma}$.)

The portfolio of equation (19) has an intuitive interpretation as a "replicating" portfolio. The term $(\sigma\sigma')^{-1}$ is the covariance matrix of returns. The term $\sigma\left[\sigma_C, 0_{1\times(N-1)}\right]'$ is the vector of covariances between asset returns and the innovations to cash-flow growth $\sigma_C dB_t^{(C)}$. In other words, the elements of the vector w are the regression coefficients obtained from regressing the cash-flow growth, $\frac{dC_t}{C_t}$, on all asset returns.

Substituting the portfolio (19) into (18) gives

$$\begin{aligned}\max_w \widetilde{A}(w) &= r-\left(g-\kappa^{(1)}\sigma_C\right)-\frac{\sigma_C^2}{2}\left[1,0_{1\times N}\right]\left(I-\sigma'\left(\sigma\sigma'\right)^{-1}\sigma\right)\left[1,0_{1\times N}\right]' \\ &\leq r-\left(g-\kappa^{(1)}\sigma_C\right). \end{aligned} \tag{20}$$

When the market is complete (as I have assumed in this section), the full-rank assumption on the matrix σ implies that the weak inequality (20) is actually an equality. Equation (20) shows that if $r-\left(g-\kappa^{(1)}\sigma_C\right)<0$ then it must be that $\widetilde{A}(w)\leq \max_w \widetilde{A}(w) = r-\left(g-\kappa^{(1)}\sigma_C\right)<0<\varepsilon$, but $\max_w A(w)>\varepsilon>0$. In words, the process z_t is not stationary under the probability measure $\widetilde{Q}$, even though it is stationary under the statistician's ("natural") mature, P.

In summary, even though the cash flows C_t and X_t are co-integrated, the yields on cash-flow streams to C_T or X_T, as T approaches infinity may be entirely different. It is useful to note that $r-\left(g-\kappa^{(1)}\sigma_C\right)<0$ is sufficient to ensure the that $\widetilde{A}(w)<0$, but not necessary. In particular, even if $r-\left(g-\kappa^{(1)}\sigma_C\right)>0$, there will exist portfolios w, so that $\widetilde{A}(w)<0$. As a result of $\widetilde{A}(w)<0$, the process z_t is not stationary under $\widetilde{Q}$ and the long-term yields of cash flows to C_T and X_T will generally differ and converge to different, albeit positive numbers.[13] In other words, long-term yields on the dividend strips to X_T and C_T will converge to different numbers even if the prices of the two dividend strips approach zero as $T\to\infty$.

To complete the analysis, I next show that the divergence of the dividend-yield term

[13]The assumption $r-\left(g-\kappa^{(1)}\sigma_C\right)>0$ implies the positivity of the yields on cash flow streams to C_T.

structures is not only an asymptotic phenomenon. The entire shape of the term structure of dividend-strip-yields differs: Clearly, because $\log C_t$ follows a random walk, the yield curve on dividend strips to C_T is constant and independent of maturity (all yields are simply equal to $r - \left(g - \kappa^{(1)}\sigma_C\right)$.) By contrast, the term structure of dividend-strip-yields to X_T will not be flat. Figure 2 provides an illustration by plotting the yield differences between the dividend strips to X_T and C_T for the same parameters as in Figure 2.

To understand Figure 2, I start by noting that the drift of z_t under the measure $\widetilde{Q}$ is negative and bounded away from zero. Therefore, under the measure $\widetilde{Q}$, z_t becomes arbitrarily negative in the long run with probability one, and $\max(\alpha + \beta z_t, \varepsilon)$ equals ε with probability one in the long run. Accordingly, the dynamics of z_t become $dz_t = \left(\widetilde{A}(w) - \varepsilon\right) dt + \sigma_w dB_t^{\widetilde{Q}}$, where $\sigma_w \equiv w'\sigma - \left[\sigma_C, 0_{1\times(N-1)}\right]'$. Using properties of the log-normal distribution, it follows that the long-term yield-discrepancy, $\mathcal{R}$, is approximately given by

$$\mathcal{R} \approx \lim_{T-t\to\infty} \frac{1}{T-t} \log E_t^{\widetilde{Q}} \left(e^{\log\varepsilon + z_T - z_t}\right) = \widetilde{A}(w) - \varepsilon + \frac{\sigma_w^2}{2}. \tag{21}$$

This asymptotic discrepancy is depicted by the dotted horizontal line in Figure 2.

On the opposite extreme, consider values of T close to t. For such values of T close to t, the negative drift in z_t under the measure $\widetilde{Q}$ will generally exceed ε in absolute value in the short run, especially for large initial values of z_t. Indeed, the higher is the initial value of z_t, the larger the magnitude of the negative drift, because of the payout specification (5). This is reflected in the fact that the lines in Figure 2 decays faster, the larger is the initial value of z_t.

These observations show that two co-integrated claims can exhibit (a) a positive gap in the yields of their long-term dividend strips even at infinity, and (b) the shape of their dividend-strip term structures can be entirely different. An implication of this finding is that even though aggregate dividends and aggregate consumption are co-integrated, the term structure of dividend strips may be uninformative about the term structure of aggregate consumption strips: Typically, economic models make predictions about the latter term structure, whereas

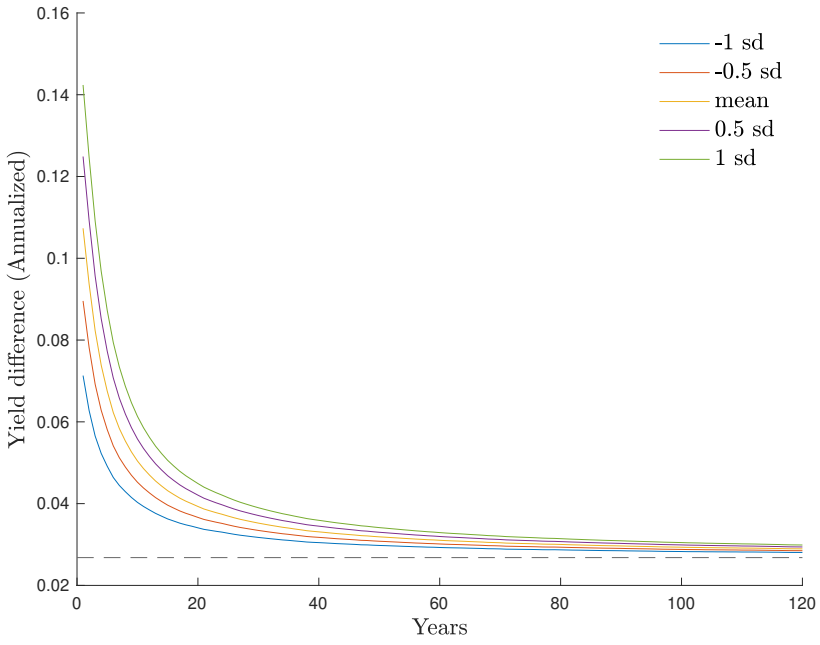


Figure 2: Difference in the yields of dividend strips to X_T and C_T. The different lines correspond to initializing z_t at the stationary mean of z (under the statistical measure, P) plus / minus $0, 0.5, 1$ stationary standard deviations. The dotted line depicts the long-term yield of the cash flow strip in equation (21). The parameters are the same as in Figure 1. The value of $\widetilde{A}(w)$ is -0.02.

in the data one can at most observe the former.[14]

The next section uses the divergence of the term structures at infinity to derive some implications about the relation betwen so-called "dynamic efficiency conditions" and the possibility of bubbles in stochastic economies.

4 Dynamic efficiency conditions, and no-bubble conditions

Proposition 2 helps provide an asset-pricing perspective on why dynamic-efficiency conditions do not imply that an economy is free of bubbles. Specifically, the sufficient condition for an economy to be "dynamically efficient" is that the average logarithmic return to capital

[14] An analysis of dividend strips from S&P options is contained in van Binsbergen et al. (2012), who find that the term structure of dividend strips is downward sloping.

exceeds the average logarithmic growth of the economy. By contrast, a sufficient condition for ensuring the "absence of bubbles" (on positive-supply assets) is that the present value of the aggregate endowment is finite, which implies that the prices of strips to the aggregate endowment at time T converge to zero as T goes to infinity.[15] This section shows that if Proposition 2 were not true, that is, if two cash flows that are co-integrated must either both be finitely-valued or both be infinitely-valued, then dynamic efficiency would actually imply the absence of bubbles even in a stochastic economy. More importantly, Proposition 2 leads to Proposition 3 below, which shows that dynamic efficiency may imply the absence of bubbles even in some stochastic economies.

To provide some background, and keep the paper self-contained, it is useful to provide a brief summary of the notion of dynamic efficiency (Cass (1972)) in a discrete-time, deterministic economy. The criterion of dynamic efficiency is useful for detecting whether the economy is over-accumulating capital or not. Specifically, an economy is over-accumulating capital if it is possible to reduce today's investment (and increase aggregate consumption) by $\epsilon > 0$, without changing aggregate consumption in any future period. Letting $R_{K,t+i}$ denote the gross return on capital and $1+g$ the aggregate economic growth rate, the Cass criterion boils down to testing whether the cumulative product, $\prod_{i=1}^{T} \frac{R_{K,t+i}}{1+g}$, converges to zero or diverges to infinity. Intuitively, if $\prod_{i=1}^{T} \frac{R_{K,t+i}}{1+g}$ goes to infinity, then the consumption increase today will have to lead to some consumption cut in the future.[16] Zilcha (1990), Zilcha (1991) extend this observation to stochastic economies and shows that an economy is dynamically efficient as long as the expected logarithmic return on capital exceeds the expected logarithmic

[15]The condition that the price of these strips converges to zero is also necessary for excluding interim Pareto inefficiency in overlapping generations economies. See Bloise and Reichlin (2023).

[16]Starting in a steady state and reducing next period's capital by $\varepsilon > 0$ implies that output at time $t+1$ will be smaller by $\varepsilon F_{K,t+1}$, where $F_{K,t+1}$ is the marginal product of capital. In turn, since consumption at time $t+1$ is unchanged, capital at time $t+2$ will be smaller by $\varepsilon\left(F_K\left(K_{t+1}\right)+(1-\delta)\right)=\varepsilon R_{K,t+1}$, where δ is depreciation and $R_{K,t+1}$ is the gross return on capital. Iterating forward, capital at time $T+1$ will be smaller by $\varepsilon\prod_{i=1}^{T} R_{K,t+i}$. In the unperturbed economy the capital stock grows at the rate $(1+g)$, so that the ratio of perturbed-to-unperturbed capital stocks at time $T+1$ is $\varepsilon\prod_{i=1}^{T} \frac{R_{K,t+i}}{1+g}$. If the steady state value of the return on capital, $R_{K,t+i}$ exceeds $1+g$, then sooner or later it will be inevitable to cut back aggregate consumption because there will not be enough capital to support any given finite consumption.

growth rate of the economy, $E \log R_{K,t} \geq E \log (1 + g_t)$.

The sufficient condition that guarantees the absence of bubbles on positive supply assets is distinct from the dynamic efficiency condition. Again in the interest of keeping the paper self-contained, it is useful to provide a short derivation of this condition in a fairly simple, continuous-time, overlapping-generations, endowment economy. Suppose that individuals come to life at time s and live for T periods. Denote by $c_{t,s} \geq 0$ an individual's consumption at time t, where s indexes a birth cohort, and t indexes calendar time. Similarly, let $\omega_{t,s} \geq 0$ denote the individual's labor income and $A_{t,s}$ her assets. Assuming that consumers come to life without assets ($A_{s,s} = 0$), and that the financial markets are complete, the individual's life-time budget constraint is

$$A_{t,s} = E_t \left(\int_t^{s+T} \frac{m_u}{m_t} \left(c_{u,s} - \omega_{u,s} \right) du \right), \tag{22}$$

where m_t is the unique stochastic discount factor. To focus on infinite-horizon bubbles and exclude the finite-horizon bubbles that can arise simply because of continuous trading, assume that all Sharpe ratios and the interest rate are bounded. The following result, which is a simplified version of Santos and Woodford (1997b), shows that as long as the present value of the aggregate endowment is finite, the total value of all assets is equal to the present value of dividend income (aggregate consumption minus aggregate wage income.)

Proposition 3 *Define aggregate consumption,* $C_t \equiv \int_{-\infty}^{+\infty} c_{t,s} 1_{\{s \leq t \leq s+T\}} ds$, *and aggregate labor income,* $\Omega_t \equiv \int_{-\infty}^{+\infty} \omega_{t,s} 1_{\{s \leq t \leq s+T\}} ds$. *Also let* $\overline{A}_t \equiv \int_{-\infty}^{t} A_{t,s} 1_{\{s \geq t-T\}} ds$ *denote the aggregate value of all assets in the economy. Assume that* $E_t \int_t^{\infty} \frac{m_u}{m_t} \Omega_u du \leq E_t \int_t^{\infty} \frac{m_u}{m_t} C_u du < \infty$. *Then the aggregate value of all assets is equal to the present value of aggregate dividend income over the infinite horizon,*

$$\overline{A}_t = E_t \int_t^{\infty} \frac{m_u}{m_t} \left(C_u - \Omega_u \right) du. \tag{23}$$

Because the value of all assets obeys the present value relationship (23), there can be no bubbles on positive supply assets. The finiteness of the aggregate endowment, $E_t \int_t^{\infty} \frac{m_u}{m_t} C_u du <$

∞, is the key assumption of Proposition 3. Without this condition, an economy can exhibit bubbles.

In a *deterministic* economy that possesses a steady state, the dynamic efficiency condition boils down to whether the steady state return on capital exceeds the aggregate growth of the economy. Additionally, the absence of risk implies that all rates of return are equal, and accordingly the dynamic efficiency condition is equivalent to the interest rate exceeding the growth rate in the economy, which implies that $\int_t^\infty \frac{m_u}{m_t} C_u du = C_t \int_t^\infty e^{-(r-g)(u-t)} du < \infty$.

In a stochastic economy, the dynamic efficiency condition differs from the condition that the aggregate endowment is finite. Proposition 2 helps explain why the dynamic efficiency condition does not imply the finiteness of the aggregate endowment. Indeed, suppose (counterfactually) that the statement of Proposition 2 is incorrect. Specifically, assume (counterfactually) that if there exists a (finite-valued) fund that finances a cash flow processes X_t that is co-integrated with consumption, C_t, then the present value of C_t must also be finite. If this statement were true, then the dynamic efficiency condition would imply the finiteness of the endowment condition, even in a stochastic economy. The reason is that (a) capital is a tradable asset, and (b) the dynamic efficiency condition (expressed in continuous time) is $r - g + \mu_K^e - \frac{\sigma_K^2}{2} + \frac{\sigma_C^2}{2} > 0$; therefore, the condition $A(w) > 0$ of Proposition 2 is satisfied by the portfolio that invests exclusively in capital. Put simply, the dynamic efficiency condition implies the existence of a portfolio (plain investment in capital) that finances a cash flow that is co-integrated with C_t. Were it not for the (negative) result of Proposition 2, then the dynamic efficiency condition would imply the finiteness of the endowment condition.

The results of the previous section don't only provide an asset-pricing perspective on why dynamic efficiency does not imply finiteness of the aggregate endowment. The results of Proposition 2 also provide a route to derive slightly more general "dynamic efficiency" conditions that do imply the finiteness of the aggregate endowment. The next proposition provides an example

Proposition 4 *Let $dR_K = \mu_K dt + \sigma_K dB_t$ denote the dynamics of the return to capital, and also let $\sigma_{K,1}$ denote the first element of the vector σ_K. Suppose that the excess return on*

capital reflects exclusively compensation for endowment risk

$$\mu_K - r = \kappa^{(1)} \sigma_{K,1}. \tag{24}$$

Let $w = [\frac{\sigma_C}{\sigma_{K,1}}; \overrightarrow{0}_{N-1}]$ *be ther portfolio of risky assets that invests a fraction* $\frac{\sigma_C}{\sigma_K}$ *in capital and zero in everything else. Additionally, assume that* $A(w) > 0$. *Then* $E_t \int_t^\infty \frac{m_u}{m_t} C_u du < \infty$.

In words, Assumption (24) states that the only risk that drives the risk premium of the stock market is the risk associated with the Brownian motion B_t^C. If this assumption holds, and the expected logarithmic return of a portfolio that invests $\frac{\sigma_C}{\sigma_{K,1}}$ in capital and $1 - \frac{\sigma_C}{\sigma_{K,1}}$ in the riskless asset exceeds the logarithmic growth rate of C_t, then the present value of C_t is finite.

The proof of Proposition 4 is short and instructive. The definition of the portfolio w and assumption (24) imply that

$$A(w) = r - g + \frac{\sigma_C}{\sigma_{K,1}} \left(\kappa^{(1)} \sigma_{K,1} \right) + \frac{\sigma_C^2}{2} - \frac{1}{2} \left(\frac{\sigma_C}{\sigma_{K,1}} \right)^2 \sigma_K^2.$$

Accordingly,

$$r - \left(g - \kappa^{(1)} \sigma_C \right) = A(w) + \frac{\sigma_C^2}{2} \left[\frac{\|\sigma_K\|^2}{\sigma_{K,1}^2} - 1 \right] > 0,$$

which follows from the assumption $A(w) > 0$ and $\|\sigma_K\|^2 \geq \sigma_{K,1}^2$. Since $r - \left(g - \kappa^{(1)} \sigma_C \right) > 0$, the present value of C_t is finite.

Proposition 4 says that if the logarithmic return of a specific portfolio exceeds the aggregate logarithmic growth rate, then the present value of the endowment is finite. Unfortunately, condition (24) is a very strong one. In some special situations this condition may be warranted. One such situation is if the return on capital is *perfectly* correlated with the Brownian motion B_t^C. In that case the return on capital (appropriately combined with the risk-free asset) is effectively the replicating portfolio of the cash flow C_t. But then the present value of C_t must be finite because the value of capital is finite. Of course in the

data, the return on capital, which is typically proxied by a delevered version of the return on the stock market, is far from perfectly correlated with either consumption or output. In addition, as I argue in Section 5, only a small fraction of the equity premium appears to reflect compensation for aggregate risk; the fluctuations in stock returns that are orthogonal to GDP growth seem to command a sizeable premium as well.[17]

Another situation where assumption (24) could make sense is if one assumed that the economy is populated by a representative agent, and therefore the only source of risk to her marginal utility emanates from the Brownian shock, B_t^C. But in that case, it would be pointless to check the finiteness of the aggregate endowment. The transversality condition of the representative agent's optimization problem would ensure this finiteness of the aggregate endowment in the first place.

I conclude this section with two remarks.

First, the analysis sofar has focused on situations where the market is complete ($M = N$), so as to remove any ambiguity associated with the pricing of dividend strips. Incompleteness makes the implications of Proposition 2 particularly unfortunate: When the market is complete, the disconnection between dynamic efficiency and the finiteness of the endowment becomes more of a theoretical exercise. From a practical perspective, a complete market allows the determination of the value of a portfolio that exactly replicates the cash flow C_T. Therefore, one can examine directly whether the value of this replicating portfolio tends to grow without bound as T increases. In an incomplete market, one cannot exactly replicate the value of a claim to C_T, and $\kappa^{(1)}$ cannot be tied down uniquely. If Proposition 2 were not true, that is, if it were always the case that two cointegrated claims are either both finitely valued or both infinitely valued, then one could sidestep the challenges of incompleteness: To establish the finiteness of the endowment, it would suffice to find a traded asset (or portfolio) that finances a cash flow that is co-integrated with C_T. In turn, as I discussed above, this would always be possible in a dynamically efficient economy by investing in capital. Unfor-

[17] In incomplete-markets models and OLG models aggregate risk is only part of the equity premium; distributional risks are priced as well. See, e.g., Gârleanu et al. (2012), Gârleanu and Panageas (2021), Gârleanu and Panageas (2023). But even in some representative-agent models, the risk compensation in the stock market may reflect compensation for depreciation shocks, investment specific shocks, etc. See, e.g., Bulow and Summers (1984) or Papanikolaou (2011).

tunately, Proposition 2 stops this line of reasoning in its tracks. To find out if $r-(g+\kappa^{(1)}\sigma_C)$ is positive or not, one needs to elicit $\kappa^{(1)}$ somehow. This is the topic of the next section.

The second comment is parenthetical. In Appendix B I address the following question: Intuitively, how is it possible for the value of the aggregate endowment to be infinite, and yet for the stock market, which is "the present value of aggregate dividends" to be finite? The answer provided in Appendix B is that the stock market is the present value of dividends of the units of the capital stock *currently in existence,* which is finite. Whether the value of the stock market is also the the present value of *aggregate dividends* may or may not be true. Specifically, Appendix B considers a production economy where average q and marginal q are both equal to one (no adjustment costs.) By definition, a unit of capital is equal to the discounted present value of the profits that will accrue to this specific unit of capital from 0 to infinity. In turn, the value of the stock market is the present value of profits that will accrue to the *existing, not future* units of capital, which will depreciate over time. This depreciation implies that the profit share accruing to the existing units of capital will decline as a fraction aggregate profits. As a result the growth rate of profits of a fixed unit of the capital stock is below the growth rate of aggregate profits. Because of this, it is possible for the present value of profits of the existing units of the capital stock to be finite, while the present value of aggregate dividends may not even be a well-defined object. The stock market equals the present value of aggregate profits only under the additional assumption that the present value of a strip that pays a cash flow proportional to the aggregate capital stock , K_T, converges to zero as $T \to \infty$. The intuition for this finding is that when new units of the capital are added to the economy, the net present value of each addition is zero; the investment cost simply equals the present value of profits that will accrue to the new units of the capital stock.

5 The market price of GDP-risk

The previous sections highlighted that it is not possible to argue that the present value of the consumption stream is finite simply because there are claims to cash flows that are

co-integrated with GDP that have finite value. In particular, the dynamic efficiency condition guarantees the existence of a cash flow process that is *co-integrated* with GDP, but unfortunately this does not imply that a cash flow stream that is *proportional* to GDP has finite value. Reaching the conclusion that such a cash-flow stream is finitely-valued requires a stance on the risk neutral dynamics of interest rates and GDP growth. Taking a stance on the risk neutral dynamics of GDP growth (the risk compensation $\kappa^{(1)}$) is particularly difficult, since there are no traded GDP futures. In addition, it appears impossible to create portfolios that exhibit perfect correlation with GDP growth, so as to use the Sharpe ratio of such a portfolio to infer the risk compensation $\kappa^{(1)}$.

In the next two subsections I take a pragmatical approach to the issue of gauging the magnitude of $\kappa^{(1)}$. Section 5.1 examines a historical episode where GDP-linked bonds (more precisely bonds linked to industrial production) were actually traded. Subsection 5.2 takes a more indirect approach by deriving the level of $\kappa^{(1)}$ that is implied by the observation that attempts to introduce GDP futures in the US market resulted in essentially zero uptake.

5.1 The French experience in the mid-1950's

It is hard to find historical episodes where countries issued bonds with coupons linked to GDP growth.[18] Typically, the attempts to introduce GDP-linked bonds occur in developing economies and usually in the aftermath of sovereign default. Not too surprisingly, market participants tend to be unwilling to purchase such bonds, since the possibility of another sovereign default looms in their minds.

An interesting exception, where a large, developed economy introduced contingent bonds is the case of France in the mid 1950s.[19] In addition, there is readily available data on the market prices of these bonds spanning about 15 years.

Specifically, on June 1, 1956, the French Government issued bonds with coupons that

[18]There is a sizeable academic and practitioner literature on the benefits of GDP-linked debt. See Borensztein and Mauro (2004) for a comprehensive survey. See also Benford et al. (2016), Bowman and Naylor (2016), Blanchard et al. (2016), Kamstra and Shiller (2009), Shiller (1998) for indicative policy proposals. Barro (1995), Hatchondo and Martinez (2012), and Caballero and Panageas (2008) are indicative papers that discuss the potential benefits of indexation in macroeconomic models.

[19]See, e.g., Rozental (1959).

were contingent on the path of industrial production ("Bons d'Equipement Industriel et Agricole".) According to the press release in Le Monde (May 24, 1956), these bonds would have denominations of 10,000 , 100,000, and 1,000,000 francs and a 15-year maturity. Redemption would be at 105% of par. The interesting part was the determination of the annual coupon. According the press release, the coupon would be at a 5% minimum; in addition, the coupon increased by 0.05% for each point that the annual index of industrial production exceeds its level in 1955 (base 120 in 1955.) In terms of taxation, "interest and the increase in interest are exempt from tax on securities, the redemption premium benefits from the same exemption and is not subject to the progressive surtax on personal income."

These bonds were traded on the French stock exchange and are available on the "Database for Financial History" (dfih.fr). The database contains 875 price observations (typically on a biweekly frequency from August 1956 to April 1971) across all three denominations.[20] I complement this time series with the following data series: (1) Annual data on French real GDP 1950-1971 from FRED. (2) Monthly data on "Production of Total Industry in France", available from 1960-1971. (3) Monthly data on Long Term Bond Yields (10-year) available from 1960-1971. To obtain interest rate data prior to 1960, I also use Tableu 3 from Chauveau (1975) for the years 1953-1972.

As is frequently the case with historical data, there are some important limitations. The first and most important limitation is that there appears to have only been a single issue of those bonds. This makes it impossible to derive a term structure (since this would require multiple maturities on the same date.)

Second, it is not completely clear which is the reference index that is being used for the determination of coupons. For the analysis that follows I use the index "Production of Total Industry in France." (PTIF.) At an annual frequency, the logarithm of this index is essentially perfectly correlated with the logarithm of real GDP (correlation of approximately 99% from 1960-1971.) The left plot of Figure 3 provides a visual impression. Because of this very tight correlation, I use the log-growth rate of GDP to extend the index of industrial production backwards to 1956, albeit at an annual frequency. (The data on industrial production from

[20]The database indicates that the data corresponds to actual transactions, because the entry "transaction #" takes numbers larger than 1 on several days.

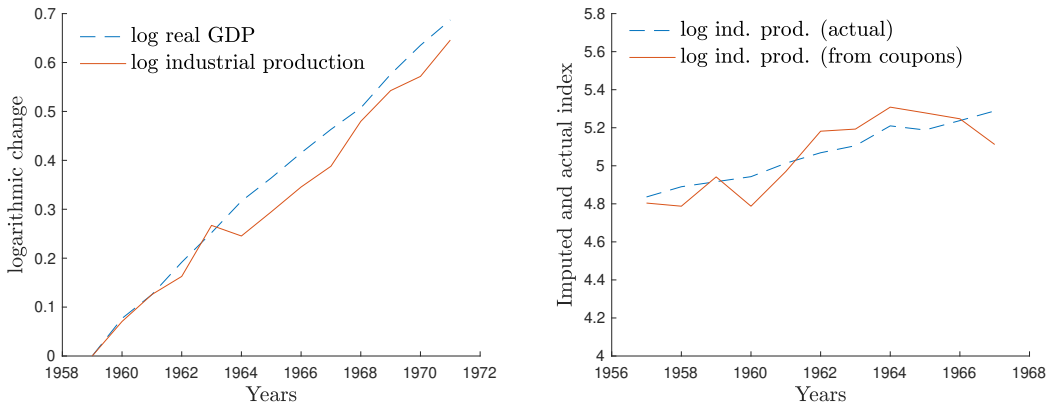


Figure 3: Left subplot: Log real GDP (dotted line) and log industrial production (solid line). Base year=1959. Right subplot: Inferred log-industrial production from coupon payments (solid line), and the log-index of industrial production from FRED (dotted line.)

1960 onwards are monthly.)

A simple way to validate that the PTIF index is the appropriate reference index for these bonds is to use the drops in the price of the bonds each June to infer the annual coupon (payable at the beginning of June). From the inferred coupon, one can use the contractual formula for the determination of the coupon to impute the value of the index that was used for the coupon determination and then compare the result to the PTIF index. The main limitation of this indirect inference approach is that the coupon payment can only be inferred from bi-weekly, not daily data, which introduces some noise in the imputation of the coupon. Nonetheless, Figure 3 shows that the imputed index corresponds reasonably well to the PTIF index. Also, the essentially unit correlation of the PTIF index with GDP implies that effectively these bonds are GDP-contingent bonds.

I next perform the following exercise. I assume that the dynamics of GDP follow the process

$$\frac{dY_t}{Y_t} = g^Q dt + \sigma dB_t^Q \tag{25}$$

under the risk-neutral measure Q. Using (25) and letting $T_i \geq t$ denote the dates of remaining coupon payments after date T, and T the date of maturity, the arbitrage-free price of the

bond $\widehat{P}(Y_t)$ at time t is

$$\begin{aligned}\widehat{P}(Y_t) &= 100 \times \sum_{T_i}^{T} e^{-r_t(T_i-t)} \left[0.05 + 0.0005 \times E_t^Q \left(Y_t - Y_{1955}\right)^+\right] + 105e^{-r_t(T-t)} \qquad (26)\\ &= 100 \times \sum_{T_i}^{T} \left[e^{-r_t(T_i-t)}0.05 + 0.0005 \times C^{BS}\left(Y_t, Y_{1955}, T_i - t, r_t, \sigma, r_t - g^Q\right)\right]\\ &+ 105e^{-r_t(T-t)},\end{aligned}$$

where $C^{BS}(S_t, K, T-t, r, \sigma, q)$ denotes the Black-Sholes Call option formula, and S_t is the price of the underlying, K is the strike price, $T-t$ is the time to maturity, r is the interest rate, σ the volatility of the underlying, and q is the continuous dividend yield on the underlying.[21] Equation (26) allows one to infer the value of Y_t by solving the equation $P_t = P(Y_t)$, where P_t is the observed price of the security at time t. This leads to the implicit value of industrial production $\widetilde{Y}_t(P_t; g^Q, \sigma) = \widehat{P}^{-1}\left(P_t; g^Q, \sigma\right)$. In words, this equation says that conditional on a choice of g^Q, and σ, one can infer a time series of imputed $\widetilde{Y}_t$ from the time-series of P_t. σ is the same under the statistical and the risk-neutral measure, and therefore I use monthly data on the index of industrial production and estimate an annualized volatility of $\widehat{\sigma} = 0.042$ from 1960-1971. I then define $u_{t_i} \equiv \log Y_{t_i} - \log \widetilde{Y}_{t_i}(P_{t_i}; g^Q, \widehat{\sigma}) - \alpha$ and estimate the risk-neutral growth parameter, g^Q, to minimize the expression

$$\widehat{g}^Q = \arg\min_{g^Q,\alpha} \sum_{i=1}^{N} u_{t_i}^2. \qquad (27)$$

If the quantity Y_t was a traded security, the timing of observation for Y_{t_i} and $\widetilde{Y}_{t_i}$ were perfectly synchronized, there is no micro-structure noise, etc. then $u_{t_i} = 0$ for the true value of g^Q. However, because Y_{t_i} is available only once a month, whereas P_{t_i} is an average of bi-weekly observations, and there are liquidity effects that could cause short run disturbances in prices, one would expect non-zero residuals u_{t_i}. Assuming that u_{t_i} is i.i.d. and independent of P_{t_i}, Y_{t_i}, then (27) can be estimated by non-linear least squares. Using the implicit function

[21] Note that Y_t is not a directly traded security, but by assuming a dividend yield equal to $q = r_t - g^Q$, the drift of $\frac{dY_t}{Y_t}$ under the risk neutral measure is given by $r - q = r - (r - g) = g^Q$, which is precisely (25).

	Full Sample	Subsamples		
End Year:	1969	1962	1965	1967
σ =0.043	0.0259	0.0277	0.0282	0.0240
	(0.0026)	(0.0065)	(0.0036)	(0.0031)
σ =0.020	0.0268	0.0300	0.0297	0.0251
	(0.0026)	(0.0062)	(0.0035)	(0.0031)
σ =0.030	0.0265	0.0292	0.0292	0.0248
	(0.0026)	(0.0063)	(0.0036)	(0.0031)
σ =0.040	0.0260	0.0281	0.0285	0.0242
	(0.0026)	(0.0064)	(0.0036)	(0.0031)
σ =0.050	0.0253	0.0266	0.0273	0.0233
	(0.0027)	(0.0068)	(0.0037)	(0.0032)

Table 1: Estimates of g^Q for different subsamples and different volatility assumptions.

theorem to compute $\frac{\partial \widetilde{Y}_t(P_t;g^Q,\sigma)}{\partial g^Q}$ gives[22]

$$Z_t \equiv \frac{\partial \widetilde{Y}_t(P_t;g^Q,\sigma)}{\partial g^Q} = -\frac{\sum_{T_i \geq t}^{T}(T_i - t) \times \frac{\partial C^{BS}\left(Y_t, Y_{1955}, T_i - t, r_t, \sigma, r_t - g^Q\right)}{\partial Y_t}}{\sum_{T_i \geq t}^{T} \frac{\partial C^{BS}(Y_t, Y_{1955}, T_i - t, r_t, \sigma, r_t - g^Q)}{\partial Y_t}}.$$

It follows that $\widehat{g}^Q - g^Q$ is asymptotically normal with mean zero and standard error equal to $\frac{1}{\sqrt{N}} \frac{\widehat{\sigma}\left(u_{t_i}\right)}{\widehat{\sigma}\left(Z_{t_i}\right)}$.

Table 1 shows that the estimated values of g^Q are in a range of $0.0233 - 0.03$ depending on the subsample used and considering many alternative volatility choices. The standard errors of the estimate for g^Q is 0.0026 for the full sample, implying a 95% coverage interval from $0.021 - 0.031$. The estimates don't seem to be particularly sensitive to the subsample used, nor are they particularly sensitive to the assumption about the volatility σ.

To put the numbers in perspective, the average growth rate of real log GDP from 1950

[22]The implicit function theorem implies that

$$\frac{d \log \widetilde{Y}}{dg^Q} = \frac{1}{\widetilde{Y}} \frac{\sum C_q^{BS}\left(\widetilde{Y}, Y_{1955}, T_i - t, r_t, \sigma, r_t - g^Q\right)}{\sum C_{\widetilde{Y}}^{BS}\left(\widetilde{Y}, Y_{1955}, T_i - t, r_t, \sigma, r_t - g^Q\right)}$$

Using standard properties of the Black-Sholes formula $C_{\widetilde{Y}}^{BS} = e^{-q(T-t)} N(d_1)$ and $\frac{C_q^{BS}}{\widetilde{Y}} = -(T-t)\, e^{-q(T-t)} N(d_1) = -(T-t)\, C_{\widetilde{Y}}^{BS}$.

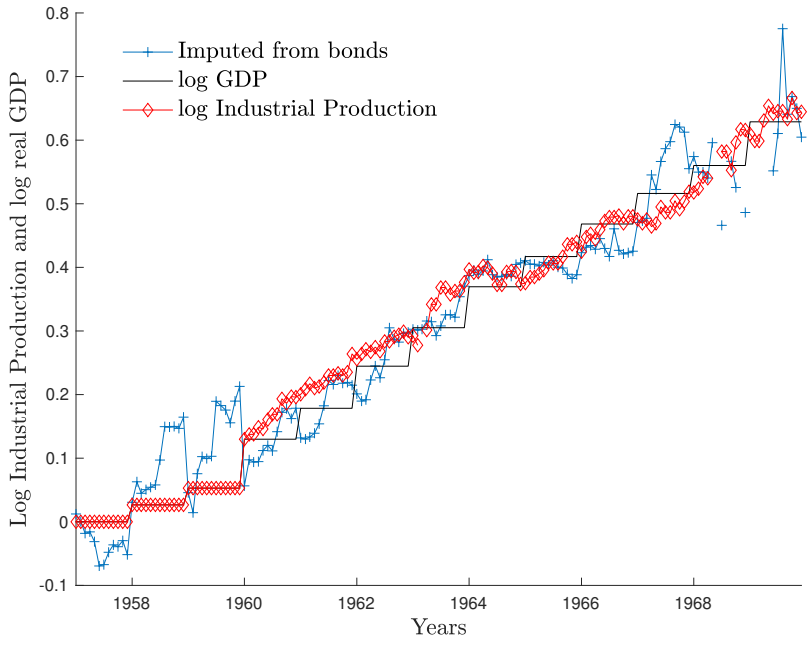


Figure 4: Log-Industrial production index inferred from bond prices (blue line), log real GDP (black line) and the index of log industrial production (red line). The data for the third time series start in 1960.

to 1960 in France was about 0.0425. Even though I don't have readily available data on the industrial production index during this period, the near unit correlation between log GDP and log industrial production suggests that the growth rate of the log industrial production was also 0.0425 for the most recent historical period around the launch of these bonds. Accordingly, the baseline value of 0.026 for g^Q translates into a "Sharpe ratio" $\kappa^{(1)} = \frac{g-g^Q}{\sigma} = \frac{0.0425-0.026}{0.0427} = 0.39$. This number for the Sharpe ratio of aggregate risk is in the same ballpark that one encounters for the Sharpe ratio of equity markets.

Figure 4 gives a visual impression of the goodness of fit. The figure depicts three time series: (1) the inferred index of industrial production (monthly), (2) real log gdp growth (annual), and (3) the index of industrial production (available post 1960-monthly). The figure shows that the index of industrial production that is inferred from bond-price data is somewhat more volatile than the other two indices, but overall the series co-move quite closely. In particular, the discrepancies between the time series (1) and (3) appear quite

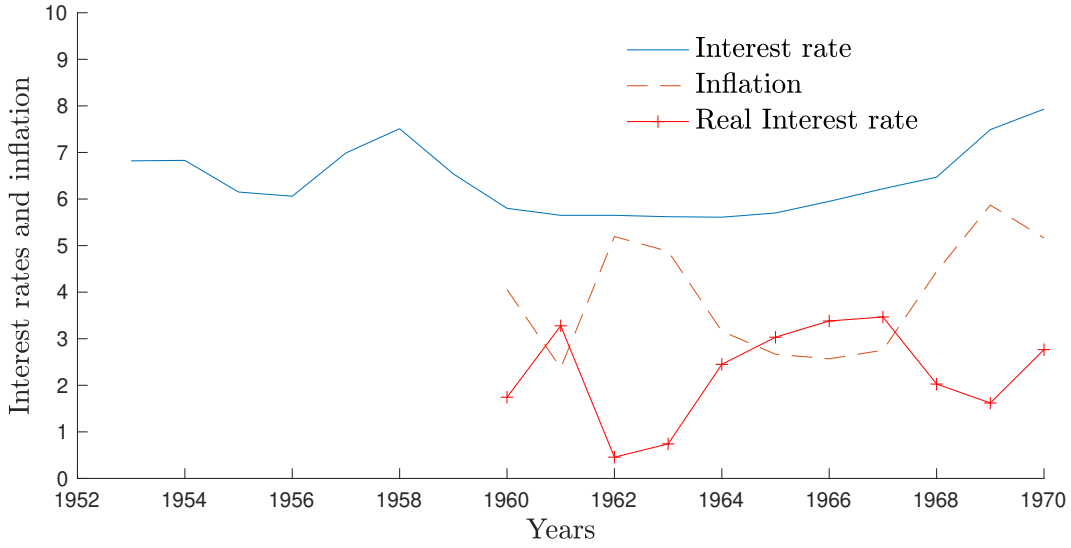


Figure 5: Interest rates (from 10-year bonds), inflation and the real interest rate. (Data on inflation are readily available from 1960 onward.)

transient, consistent with the interpretation that they reflect liquidity effects and other trading frictions.

I conclude this section with a remark. The (realized) real rate of the French economy during this period was 0.021 and quite stable around that value, as Figure 5 shows. This value is lower than the baseline estimate of 0.026 for g^Q implying that — over this sample period – the risk-adjusted value of the growth rate was higher than the real rate.

Of course, this is not conclusive evidence, because one would like to know whether the real rate (at an infinite maturity) as inferred from the term structure is below or above g^Q. The data of that period do not allow one to perform such an exercise.

5.2 The plausibility of the relatively low market price of risk for aggregate output

The previous section inferred the market price of aggregate-output risk from securities whose value is directly dependent on output risk. The main result of the section was that the Sharpe ratio for output risk was not spectacularly high. As a plausibility check, this section takes a different approach to gauging the plausibility of this low estimate.

The analysis starts with the following observation: Historically, there has been a multitude of proposals to introduce gdp-linked bonds in advanced economies. These proposals were primarily discussed in academic circles. On occasion, these proposals reached policy makers and practitioners. Yet, there was no uptake by market participants. The approach taken in this section is to try to infer some information from this reluctance to embrace gdp-linked bonds. Specifically, the section takes the view that market participants must have perceived that the introduction of a zero net supply security (indexed to growth) would not substantially raise the welfare of the representative investor and would not induce her to choose a different portfolio than she is already holding.

To formalize this intuition, suppose that $M = N + 1$ in equation (2) and that in addition there is now an asset (in zero net supply) with a return that is perfectly correlated with dB_t^C and a volatility that is normalized to one. By the absence of arbitrage,

$$\mu^e = \sigma\kappa, \text{ where } \kappa = \begin{bmatrix} \kappa^{(1)} \\ \kappa^{(2..N+1)} \end{bmatrix}, \tag{28}$$

and it follows that the expected (excess) rate of return of this contract is $\kappa^{(1)}$, since its volatility is normalized to one. The optimal portfolio for an investor with risk aversion γ is $w^{**} = \frac{1}{\gamma}\left(\sigma^{new}\sigma^{new'}\right)^{-1}\mu^{e,new}$, where σ^{new} and $\mu^{e,new}$ are the new covariance matrix and the new vector of excess returns. Suppose that investors perceive that the introduction of the new security will not change the return properties of all pre-existing assets, nor will it change investors' optimal portfolios. Mathematically, this idea can be expressed by the pair of equations

$$\sigma^{new} = \begin{bmatrix} [1; 0_{1\times N}] \\ \sigma \end{bmatrix}, \text{ and } \mu^{e,new} = \begin{bmatrix} \kappa^{(1)} \\ \mu^e \end{bmatrix}, \tag{29}$$

and

$$\frac{1}{\gamma}\left(\sigma^{new}\sigma^{new'}\right)^{-1}\begin{bmatrix} \kappa^{(1)} \\ \mu^e \end{bmatrix} = w^{**} = \begin{bmatrix} 0 \\ w^* \end{bmatrix}, \text{ where } w^* = \frac{1}{\gamma}\left(\sigma\sigma'\right)^{-1}\mu^e \ . \tag{30}$$

Solving for $\kappa^{(1)}$ from (30) leads to

$$\kappa^{(1)} = [1; 0_{1\times N}] \left(\sigma^{new}\sigma^{new'}\right) \begin{bmatrix} 0 \\ \sigma^{-1}\mu^e \end{bmatrix} \tag{31}$$

Equation (31) is intuitive. It says that the market price of risk $\kappa^{(1)}$ is given by the covariance between $\frac{1}{\sigma^C}\frac{dC}{C} = dB_t^{(C)}$ and the return of the mean-variance efficient portfolio $\sigma^{-1}\mu^e$. Equations of this form are well understood for the determination of the expected returns of *traded* assets. The equation continues to apply in environments where investors perceive that the introduction of a *new* asset will not alter the return properties of the existing assets. An early paper that introduced this assumption in the context of valuing the investment plans of a firm in incomplete markets is Grossman and Hart (1979) (the "competitive perceptions", or more appropriately "utility-taking" assumption.) While in general this assumption is a strong one, it appears somewhat more plausible in situations where market participants contemplated the introduction of a new security and presumably concluded that it would not affect investors' welfare and marginal utilities; by implication it would leave asset prices of existing securities largely unchanged.[23]

The estimation of $\kappa^{(1)}$ from (31) is straightforward conceptually, but econometrically challenging, because the mean-variance efficient portfolio estimated by using the empirical covariance matrix and the historical returns could differ substantially from its theoretical

[23]The value of $\kappa^{(1)}$ resulting from (31) can also be viewed as the "shadow" market price of risk that one would obtain if an investor were simply constrained to choose a weight of zero in the first asset. Specifically, consider the problem of maximizing the constrained problem

$$V = \max_w \left[w'\mu^{e,new} - \frac{1}{2}w'\left(\sigma^{new}\sigma^{new'}\right)w\right], \text{ subject to } [1; 0_{1\times N}]'\, w = 0.$$

Using the envelope theorem, equation (29), and recalling the constraint $[1; 0_{1\times N}]'\, w = 0$ implies

$$\frac{\partial V}{\partial \kappa^{(1)}} = [1; 0_{1\times N}]'\, w^{opt} = 0. \tag{32}$$

Equation (32) has an interpretation in terms of the theory of "convex duality": The value of $\kappa^{(1)}$ that renders the portfolio (30) optimal is obtained by minimizing V over all possible vectors κ that satisfy the (no-arbitrage) requirment (28). (To see this, note that equation (32) is simply the first order condition associated with the problem of minimizing V over $\kappa^{(1)}$. Equation (32) together with the no-arbitrage restrictions (28) fully determine the vector κ that minimizes V.)

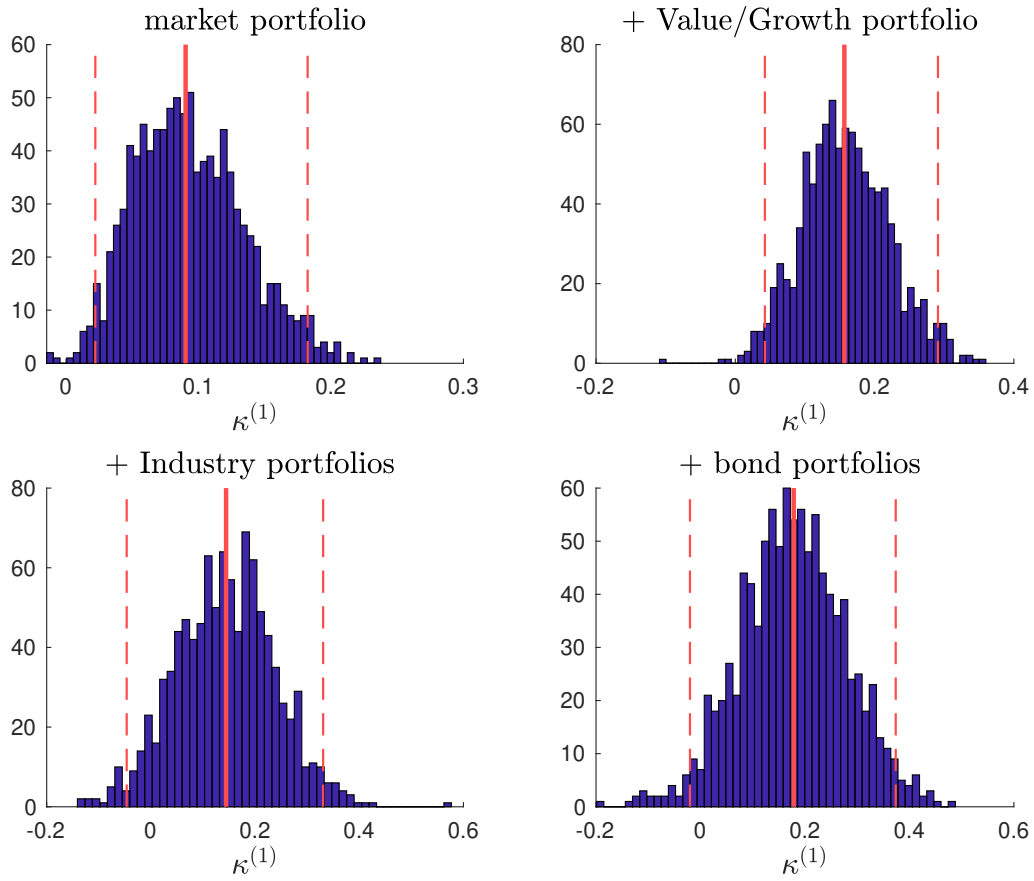


Figure 6: Estimates of the market price of GDP risk, $\kappa^{(1)}$, using a progressively expanding set of portfolios as "test assets." Data are bootstrapped one thousand times and the histogram depicts the resulting values of $\kappa^{(1)}$. The solid red line corresponds to the mean and the dotted red lines to the 95% coverage intervals.

$\kappa^{(1)}$	0.093	0.159	0.143	0.177
	(0.041)	(0.063)	(0.095)	(0.101)

Table 2: Estimates and standard errors of $\kappa^{(1)}$ for different portfolio combinations. Standard errors are computed by bootstrapping 1000 samples obtained with replacement.

counterpart, $\sigma^{-1}\mu^e$. To account for this estimation uncertainty, Figure 6 uses a bootstrap methodology. Specifically, I use different combinations of portfolios to form the vector of returns (the market portfolio, the six portfolios that comprise the Fama-French factors, Fama bliss bond portfolios etc.) For each portfolio combination, the returns are resampled with replacement (1000 times), the variance covariance matrix and the average returns are re-

estimated and the quantity $\kappa^{(1)}$ is re-computed. The figure shows the distribution of the estimates for the different portfolio combinations.

There are two observations worth highlighting. First, the distribution of $\kappa^{(1)}$ does not differ much, irrespective of which combination of portfolios are used as the set of assets in equation (31). Even though the weights of the assets that enter the mean variance portfolio are different in each of the subplots, the covariance of the returns of the resulting portfolio with (real) log gdp growth is essentially unchanged. Second, the magnitude of $\kappa^{(1)}$ is in the same moderate range as the values suggested by the French experience. In all bootstrapped samples one rarely encounters a value of $\kappa^{(1)}$ above 0.4.

6 Conclusion

It is commonplace in macro-finance to impose "transversality" or no-bubble conditions everywhere. To a large extent, these conditions are motivated by the convenience of working within the framework of an infinitely-lived, representative-agent economies. Even though there are plenty of non-pathological models in the OLG and the heterogenous-agents literature that violate some of the transversality conditions imposed by representative-agent models, these models have been encountered with skepticism. Part of that skepticism is rooted in the view that the US economy appears dynamically efficient (by the Abel et al. (1989) test). In finance circles one also encounters the argument that there are some claims whose cash flows appear cointegrated with the aggregate endowment, which are finitely valued. This is taken as suggestive evidence that the aggregate endowment is finitely-valued as well, which excludes the possibility of bubbles.

Neither the dynamic efficiency nor the co-integration argument are correct. Two cash flow processes can be co-integrated, and yet the first cash flow cannot be replicated by any finitely-valued portfolio, while the second can. In this paper, I showed that this paradox of co-integration is actually responsible for why dynamic efficiency does not imply the absence of bubbles. Were it not for this paradox of co-integration, then the dynamic efficiency condition would actually imply the finiteness of the aggregate endowment and the absence

of bubbles. Without taking a stance on the market price of risk for pure GDP risk, it is impossible to conclude anything about the magnitude of risk-adjusted growth, simply by appealing to co-integration. To add to the problem, with the existing set of assets, there is no way to create a portfolio that is perfectly correlated with GDP, so that one can infer the Sharpe ratio of GDP risk.

To gauge the magnitude of this Sharpe ratio, in this paper I used a historical episode where growth-contingent bonds were actually traded. This allows one to infer the implied risk-neutral growth rate that is reflected in the prices of those bonds. The key takeaway is that the Sharpe ratio of GDP risk is 0.39, around the same order of magnitude as the Sharpe ratio of equity investments. This Sharpe ratio for GDP risk leads to two natural conclusions.

First, this magnitude of the Sharpe ratio for GDP risk, coupled with the relatively low correlation between the stock market and GDP growth, implies that a non-trivial part of the equity premium is driven by risks that appear orthogonal to GDP risk. Put more precisely, if GDP futures were traded and had a Sharpe ratio similar to the one that French investors demanded, then a portfolio consisting of an investment in the stock market and short GDP futures to make the portfolio have a beta of zero to GDP shocks, would still command a risk premium of 3.5%. Among other things, this finding implies that it is problematic to view the stock market as simply a levered claim on output. Under this view of the world, the portfolio described in the previous sentence should have an excess return of zero.

Second, the gap between statistical (g) and risk-adjusted (g^Q) differs by about 0.39 times the volatility of output (around 0.023.) This implies a rather small gap of about $0.39 \times 0.023 = 90$ basis points between the statistical and risk-adjusted growth rate of GDP. With this low risk adjustment, it is not clear whether the condition $r - g^Q > 0$ is such an empirically strong assumption that it renders all models that challenge it empirically implausible.

References

Abel, A. B., N. G. Mankiw, L. H. Summers, and R. J. Zeckhauser (1989). Assessing dynamic efficiency: Theory and evidence. *The Review of Economic Studies 56*(1), 1–19.

Abel, A. B. and S. Panageas (2022). Running primary deficits forever in a dynamically efficient economy: Feasibility and optimality. Working Paper 30554, National Bureau of Economic Research.

Aguiar, M. A., M. Amador, and C. Arellano (2021). Micro risks and pareto improving policies with low interest rates. NBER working paper 28996.

Aiyagari, S. and E. McGrattan (1998). The optimum quantity of debt. *Journal of Monetary Economics 42*(3), 447–469.

Amol, A. and E. G. Luttmer (2022). Permanent primary deficits, idiosyncratic long-run risk, and growth. Working Paper 794, Federal Reserve Bank of Minneapolis.

Ball, L., D. W. Elmendorf, and N. G. Mankiw (1998). The deficit gamble. *Journal of Money, Credit and Banking 30*(4), 699–720.

Ball, L. and N. G. Mankiw (2023). Market power in neoclassical growth models. *The Review of Economic Studies 90*(2), 572–596.

Barbie, M., M. Hagedorn, and A. Kaul (2007). On the interaction between risk sharing and capital accumulation in a stochastic olg model with production. *Journal of Economic Theory 137*(1), 568–579.

Barbie, M. and A. Kaul (2009). The Zilcha criteria for dynamic inefficiency reconsidered. *Economic Theory 40*(2), 339–348.

Barro, R. J. (1995, October). Optimal debt management. Working Paper 5327, National Bureau of Economic Research.

Barro, R. J. (2023). r minus g. *Review of Economic Dynamics 48*, 1–17.

Benford, J., M. Joy, and M. Kruger (2016). Sovereign gdp-linked bonds. Technical report. Financial Stability Paper No. 39, Bank of England.

Bertocchi, G. (1991). Bubbles and inefficiencies. *Economics Letters 35*(2), 117–122.

Binswanger, M. (2005). Bubbles in stochastic economies: Can they cure overaccumulation of capital? *Journal of Economics 84*(2), 179–202.

Blanchard, O., P. Mauro, and J. Acalin (2016). The case for growth-indexed bonds in advanced economies today. Technical report. Peterson Institute for International Economics.

Blanchard, O. J. (2019). Public debt and low interest rates. *American Economic Review 109*(4), 1197–1229.

Blanchard, O. J. and P. Weil (2001). Dynamic efficiency, the riskless rate, and debt ponzi games under uncertainty. *Advances in Macroeconomics 1*(2), 1–21.

Bloise, G. and P. Reichlin (2023). Low safe interest rates: A case for dynamic inefficiency? *Review of Economic Dynamics 51*, 633–656.

Bohn, H. (1995). The sustainability of budget deficits in a stochastic economy. *Journal of Money, Credit and Banking 27*(1), 257–271.

Borensztein, E. and P. Mauro (2004). The case for gdp-indexed bonds. *Economic Policy 19*, 165–216.

Bowman, J. and P. Naylor (2016). Gdp-linked bonds. Technical report. Reserve Bank of Australia Bulletin.

Brunnermeier, M. K., S. Merkel, and Y. Sannikov (2021). Debt as safe asset: Mining the bubble. Working Paper.

Bulow, J. I. and L. Summers (1984). The taxation of risky assets. *Journal of Political Economy 92*(1), 20–39.

Caballero, R. and S. Panageas (2008). Hedging sudden stops and precautionary contractions. *Journal of Development Economics 85*(1-2), 28–57.

Cass, D. (1972). On capital overaccumulation in the aggregative, neoclassical model of economic growth: A complete characterization. *Journal of Economic Theory 4*(2), 200–223.

Chauveau, T. (1975). La demande d'obligations émanant des ménages. *Revue Economique 26*(5), 806–833.

Cochrane, J. H. (2021). *The Fiscal Theory of the Price Level*. Princeton University Press, forthcoming.

Diamond, P. A. (1965). National debt in a neoclassical growth model. *The American Economic Review 55*(5), 1126–1150.

Farhi, E. and J. Tirole (2011). Bubbly Liquidity. *The Review of Economic Studies 79*(2), 678–706.

Gârleanu, N., L. Kogan, and S. Panageas (2012). Displacement risk and asset returns. *Journal of Financial Economics 105*(3), 491–510.

Gârleanu, N. and S. Panageas (2021). What to expect when everyone is expecting: Self-fulfilling expectations and asset-pricing puzzles. *Journal of Financial Economics 140*(1), 54–73.

Gârleanu, N. and S. Panageas (2023). Heterogeneity and asset prices: An intergenerational approach. *Journal of Political Economy 131*(4), 839–876.

Grossman, S. J. and O. D. Hart (1979). A theory of competitive equilibrium in stock market economies. *Econometrica 47*(2), 293–329.

Hatchondo, J. C. and L. Martinez (2012). On the benefits of gdp-indexed government debt: Lessons from a model of sovereign defaults. *Economic Quarterly 98*(2), 139–157.

Hellwig, M. (2021). Safe assets, risky assets, and dynamic inefficiency in overlapping-generations economies. Working Paper.

Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2019). The U.S. public debt valuation puzzle. Working Paper 26583, National Bureau of Economic Research.

Kamstra, M. and R. J. Shiller (2009). The case for trills: giving the people and their pension funds a stake in the wealth of the nation. Technical report. Cowles Foundation Working Paper CFDP 1717.

Kocherlakota, N. R. (2023a). Infinite debt rollover in stochastic economies. *Econometrica 91*(5), 1629–1658.

Kocherlakota, N. R. (2023b). Public debt bubbles in heterogeneous agent models with tail risk. *International Economic Review 64*(2), 491–509.

Lucas, D. (2012). Valuation of government policies and projects. *Annual Review of Financial Economics 4*, 39–58.

Martin, A. and J. Ventura (2012). Economic growth with bubbles. *American Economic Review 102*(6), 3033–3058.

Papanikolaou, D. (2011). Investment shocks and asset prices. *Journal of Political Economy 119*(4), 639–685.

Rangazas, P. and S. Russell (2005). The Zilcha criterion for dynamic inefficiency. *Economic Theory 26*(3), 701–716.

Reis, R. (2021). The constraint on public debt when r $<$ g but g $<$ m. Working Paper.

Rozental, A. A. (1959). Variable-return bonds—the french experience. *The Journal of Finance 14*(4), 520–530.

Santos, M. S. and M. Woodford (1997a). Rational asset pricing bubbles. *Econometrica 65*(1), 19–57.

Santos, M. S. and M. Woodford (1997b). Rational asset pricing bubbles. *Econometrica 65*(1), 19–57.

Shiller, R. J. (1998). *Macro markets:creating institutions for managing society's largest economic risks*. Oxford University Press.

Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica 53*(6), 1499–1528.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

Zilcha, I. (1990). Dynamic efficiency in overlapping generations models with stochastic production. *Journal of Economic Theory 52*(2), 364–379.

Zilcha, I. (1991). Characterizing efficiency in stochastic overlapping generations models. *Journal of Economic Theory 55*(1), 1–16.

A Appendix

Proof of Proposition 1. Using the definitions of σ_z and z^*, the dynamics for z_t can be written as

$$dz_t = \begin{cases} (A - \alpha - \beta z_t)\, dt + \sigma_z d\widetilde{B}_t & \text{if } z_t > z^*, \\ (A - \varepsilon)\, dt + \sigma_z d\widetilde{B}_t & \text{if } z_t \leq z^*, \end{cases}$$

where $\widetilde{B}_t$ is a standard Brownian motion. Note that dz_t behaves like an Ornstein-Uhlenbeck process whenever $z_t > z^*$ and dz_t behaves like a Brownian motion with (positive) drift $A(w) - \varepsilon$ whenever $z_t \leq z^*$. The Fokker-Plank equation implies that the stationary density must satisfy the differential equation

$$\frac{\sigma_z^2}{2} \frac{\partial^2 f(z)}{\partial z^2} = \frac{\partial \left[f(z) (A - \alpha - \beta z_t) \right]}{\partial z} \text{ if } z_t > z^*, \tag{33}$$

and

$$\frac{\sigma_z^2}{2} \frac{\partial^2 f(z)}{\partial z^2} = \frac{\partial \left[f(z) (A - \varepsilon) \right]}{\partial z} \text{ if } z_t \leq z^*, \tag{34}$$

subject to the additional conditions $\lim_{z \to +\infty} f(z) = \lim_{z \to -\infty} f(z) = 0$ and $f(z) : R \to R^+$ is continuously differentiable, $f \in C^1$. Integrating both sides of (33) and (34) with respect to z, and imposing the requirements $\lim_{z \to +\infty} f(z) = \lim_{z \to -\infty} f(z) = 0$ leads to the pair of differential equations

$$\frac{\sigma_z^2}{2} \frac{\partial f(z)}{\partial z} = \begin{cases} f(z) (A - \alpha - \beta z_t) & \text{if } z_t > z^*, \\ f(z) (A - \varepsilon) & \text{if } z_t \leq z^*. \end{cases}$$

This differential equation has solution

$$f(z) = \begin{cases} D_1 e^{\frac{2}{\sigma_z^2} \left[(A - \alpha) z - \beta \frac{z^2}{2} \right]} & \text{if } z_t > z^*, \\ D_2 e^{\frac{2}{\sigma_z^2} (A - \varepsilon) z} & \text{if } z_t \leq z^*. \end{cases} \tag{35}$$

for two arbitrary, positive constants D_1 and D_2. The requirement that $f \in C^1$ requires that

$$D_1 e^{\frac{2}{\sigma_z^2}\left[(A-\alpha)z^* - \beta\frac{(z^*)^2}{2}\right]} = D_2 e^{\frac{2}{\sigma_z^2}(A-\varepsilon)z^*}, \tag{36}$$

and

$$D_1 e^{\frac{2}{\sigma_z^2}\left[(A-\alpha)z^* - \beta\frac{(z^*)^2}{2}\right]} \left[(A-\alpha) - \beta z^*\right] = D_2 \left(A-\varepsilon\right) e^{\frac{2}{\sigma_z^2}(A-\varepsilon)z^*}. \tag{37}$$

The definition of z^* implies that if equation (36) is satisfied, then (37) is automatically satisfied. Therefore, D_2 can be arbitrary and D_1 must be set so that

$$D_1 = D_2 \exp\left\{\frac{2}{\sigma_z^2}\left[\beta\frac{(z^*)^2}{2} - (\varepsilon-\alpha)\, z^*\right]\right\}. \tag{38}$$

Substituting (38) into the top equation of (35) and choosing D_1 so that $\int_{-\infty}^{+\infty} f(z)\, dz = 1$ leads to (8). ■

Proof of Proposition 2. The easiest way to prove this result is to assume that $N = M = 1$. In that case $\kappa^{(1)} = \frac{\mu^e}{\sigma}$, where μ^e, σ are scalars corresponding to the expected excess return and volatility of the (single) risky asset. In turn, the condition $r < g - \kappa^{(1)}\sigma_C$, can be expressed as

$$r - g + \frac{\mu^e}{\sigma}\sigma_C < 0. \tag{39}$$

The conditions of Proposition 1 will be satisfied for some portfolio w as long as

$$\max_w A(w) = r - g + \max_w \left\{w\mu^e - \frac{1}{2}w^2\sigma^2\right\} + \frac{1}{2}\sigma_C^2 > 0.$$

Next, use the definition of $A(w)$ from (7), and observe that since $\frac{\sigma_C}{\sigma}$ is a feasible portfolio,

$$\max_w A(w) = A\left(\frac{\mu^e}{\sigma^2}\right) \geq A\left(\frac{\sigma_C}{\sigma}\right) = r - g + \frac{\mu^e}{\sigma}\sigma_C. \tag{40}$$

Assume that $\frac{\mu^e}{\sigma} \neq \sigma_C$, so that the inequality (40) is strict. Then, inequality (40) can be

used to show that there exist values of $r-g$ such that (39) holds while $\max_w A(w) > 0$. Specifically, suppose that $r-g = -\frac{\mu^e}{\sigma}\sigma_C - \delta$ for some $\delta > 0$, so that (39) holds. We have

$$A\left(\frac{\mu^e}{\sigma^2}\right) = r - g + \frac{1}{2}\left(\frac{\mu^e}{\sigma^2}\right)^2 + \frac{1}{2}\sigma_C^2 = -\delta \ - \frac{\mu^e}{\sigma}\sigma_C + \frac{1}{2}\left(\frac{\mu^e}{\sigma}\right)^2 + \frac{1}{2}\sigma_C^2 = \frac{1}{2}\left(\frac{\mu^e}{\sigma} - \sigma_C\right)^2 - \delta \ ,$$

which is positive for $0 < \delta < \frac{1}{2}\left(\frac{\mu^e}{\sigma} - \sigma_C\right)^2$. Accordingly, there exists $\delta > 0$ such that the inequality (39) and $A\left(\frac{\mu^e}{\sigma^2}\right) > 0$ both hold. Choosing the portfolio $w = \frac{\mu^e}{\sigma^2}$ and noting that $\frac{\mu^e}{\sigma} \neq \sigma_C$ implies that the two condition of Proposition 1 hold, namely $\sigma_Z > 0$ and $A\left(\frac{\mu^e}{\sigma^2}\right) > \varepsilon$ for some $\varepsilon > 0$. Accordingly $\log\frac{X_T}{C_T}$ is stationary, and yet $\lim_{T\to\infty} P_t^{(C_T)} = \infty$. The fact that the cash flow X_t can be financed with finite wealth implies that $\lim_{T\to\infty} P_t^{(X_T)} = 0$.

■

Proof of Proposition 3. Write $\overline{A}_t$ as

$$\overline{A}_t \equiv \int_{-\infty}^{t} A_{t,s} 1_{\{s\geq t-T\}} ds = \int_{-\infty}^{t} A_{t,s} 1_{\{s\geq t-T\}} ds + E_t\left(\int_t^{\infty} \frac{m_s}{m_u} A_{s,s} ds\right),$$

since $A_{s,s} = 0$. Using (22) it follows that

$$\overline{A}_t = E_t \int_{-\infty}^{+\infty} \left(\int_t^{\infty} \frac{m_u}{m_t}\left(c_{u,s} - \omega_{u,s}\right) 1_{\{s\leq u\leq s+T\}} du\right) ds.$$

Assuming that the order of integration can be exchanged,

$$\begin{aligned}\overline{A}_t &= E_t \int_t^{\infty} \frac{m_u}{m_t}\left(\int_{-\infty}^{+\infty} \left(c_{u,s} - \omega_{u,s}\right) 1_{\{s\leq u\leq s+T\}} ds\right) du \\ &= E_t \int_t^{\infty} \frac{m_u}{m_t}\left(C_u - \Omega_u\right) du,\end{aligned} \tag{41}$$

where the second equality follows from the definitions of $C_t = \int_{-\infty}^{+\infty} c_{t,s} 1_{\{s\leq t\leq s+T\}} ds$ and $\Omega_t \equiv \int_{-\infty}^{+\infty} \omega_{t,s} 1_{\{s\leq t\leq s+T\}} ds$. This shows that the aggregate value of all assets in the economy, $\overline{A}_t$, is equal to the present value of financial income, $E_t \int_t^{\infty} \frac{m_u}{m_t}\left(C_u - \Omega_u\right) du$. There is no room for bubbles, since bubbles would require that $\overline{A}_t = E_t \int_t^{\infty} \frac{m_u}{m_t}\left(C_u - \Omega_u\right) du + B_t$, where

$B_t > 0$ is the value of the bubble.

In deriving the above equations, the key step is the exchange of the order of integration in equation (41), which requires that

$$E_t \int_{-\infty}^{+\infty} \int_t^{\infty} \frac{m_u}{m_t} \left| c_{u,s} - \omega_{u,s} \right| 1_{\{s \leq u \leq s+T\}} ds du < \infty. \tag{42}$$

A sufficient condition for this inequality to hold is that $E_t \int_t^{\infty} \frac{m_u}{m_t} C_u du < \infty$ and $E_t \int_t^{\infty} \frac{m_u}{m_t} \Omega_u du < \infty$, since then (42) is an implication of the triangle inequality. ■

Proof of Lemma 1. First note that the Radon-Nikodym derivative $\frac{d\widetilde{Q}}{dQ} = \frac{C_T}{E_t^Q(C_T)}$ is given by $\frac{d\widetilde{Q}}{dQ} = \frac{C_T}{E_t^Q(C_T)} = e^{-\frac{1}{2}\sigma_C^2(T-t)+\sigma_C\left(B_T^{C,Q}-B_t^{C,Q}\right)}$. Using Girsanov's theorem, the dynamics of z_t under the probability measure $\widetilde{Q}$ are given by

$$\begin{aligned} dz_t = & \left(r - \left(g - \kappa^{(1)}\sigma_C \right) + \frac{\sigma_C^2}{2} - \frac{1}{2} w' \sigma \sigma' w - \max(\alpha + \beta z_t, \varepsilon) \right) dt \\ & + w'\sigma \left(d\overrightarrow{B}_t^{\widetilde{Q}} + \begin{bmatrix} \sigma_C \\ 0_{N\times 1} \end{bmatrix} dt \right) - \sigma_C \left(dB_t^{C,\widetilde{Q}} + \sigma_C dt \right). \end{aligned}$$

■

B Is the stock market the expected present value of aggregate dividends?

In classical q theory without adjustment costs, the combination of a first order condition for capital accumulation and a transversality condition lead to the familiar "marginal q equals one" condition,

$$\frac{V_t}{K_t} = E_t^Q \int_t^{\infty} e^{-\int_t^u (r_n + \delta_n) dn} F_K(u)\, du = 1, \tag{43}$$

where $F_K(t)$ is the marginal product of capital at time t, r_t is the interest rate and δ_t is depreciation. We therefore have that the value of the stock market is $V_0 = K_0$. Letting $dI_t = dK_t + \delta_t K_t dt$ denote investment at time t, equation (43) also implies that

$$
\begin{aligned}
V_0 &= K_0 + E_0^Q \left\{ \int_0^\infty e^{-\int_0^t r_n dn} dI_t E_t^Q \left(\int_t^\infty F_K(u)\, e^{-\int_t^u (r_n+\delta_n)dn} du - 1 \right) \right\} \\
&= K_0 + E_0^Q \left\{ \left(dK_t + \delta_t K_t dt \right) e^{\int_0^t \delta_n dn} \left(\int_t^\infty F_K(u)\, e^{-\int_0^u (r_n+\delta_n)dn} du \right) \right\} - E_0^Q \left\{ \int_0^\infty e^{-\int_0^t r_n dn} dI_t \right\} \\
&= K_0 + E_0^Q \left\{ d\left(e^{\int_0^t \delta_n dn} K_t \right) \left(\int_t^\infty F_K(u)\, e^{-\int_0^u (r_n+\delta_n)dn} du \right) \right\} - E_0^Q \left\{ \int_0^\infty e^{-\int_0^t r_n dn} dI_t \right\}
\end{aligned} \tag{44}
$$

Using integration by parts gives

$$
E_0^Q \int_0^\infty d\left(e^{\int_0^t \delta_n dn} K_t \right) \left(\int_t^\infty F_K(u)\, e^{-\int_0^u (r_n+\delta_n)dn} du \right) = \tag{45}
$$

$$
\begin{aligned}
&= E_0^Q \left[e^{\int_0^t \delta_n dn} K_t \left(\int_t^\infty F_K(u)\, e^{-\int_0^u (r_n+\delta_n)dn} du \right) \right]_0^\infty + E_0^Q \left\{ \int_0^\infty K_t F_K(t)\, e^{-\int_0^t r_n dn} du \right\} \\
&= \lim_{T\to\infty} E_0^Q \left\{ e^{\int_0^T -r_n dn} K_T \right\} - K_0 + E_0^Q \left\{ \int_0^\infty K_t F_K(t)\, e^{-\int_0^t r_n dn} du \right\}.
\end{aligned} \tag{46}
$$

Combining (44) with (46) leads to

$$
\begin{aligned}
V_0 = E_0^Q &\left\{ \int_0^\infty e^{-\int_0^t r_n dn} \left(\underbrace{K_t F_K(t)}_{\text{Aggregate profits}}\, du - \underbrace{dI_t}_{\text{Aggregate investment}} \right) \right\} \\
&+ \lim_{T\to\infty} E_0^Q \left\{ e^{\int_0^T -r_n dn} K_T \right\}
\end{aligned} \tag{47}
$$

The expression insides square brackets can be interpreted as aggregate dividends at time t, and accordingly, the first line of equation (47) corresponds to the present value of aggregate dividends. However, all expressions in (47) are "formal" in the sense that the integrals may not exist and the limit $\lim_{T\to\infty} E_0^Q \left\{ e^{\int_0^T -r_n dn} K_T \right\}$ may diverge to infinity.