Why are real interest rates so low? Secular stagnation and the relative price of investment goods

Gregory Thwaites
LSE and Bank of England

First version: 23/1/2014
This version 22/10/2014

Abstract

Over the past four decades, real interest rates have risen then fallen across the industrialised world. Over the same period, nominal investment rates are down, while house prices and household debt are up. I explain these four trends with a fifth - the widespread fall in the relative price of investment goods. I present a simple closed-economy OLG model in which households finance retirement in part by selling claims on the corporate sector (capital goods) accumulated over their working lives. As capital goods prices fall, the interest rate must fall to reflect capital losses. And in the long run, a given quantity of saving buys more capital goods. This has ambiguous effects on interest rates in the long run: if the production function is inelastic, in line with most estimates in the literature, interest rates stay low even after relative prices have stopped falling. Lower interest rates reduce the user cost of housing, raising house prices and, given that housing is bought early in life, increasing household debt. I extend the model to allow for a heterogeneous bequest motive, and show that wealth inequality rises but consumption inequality falls.

I test the model on cross-country data and find support for its assumptions and predictions. The analysis in this paper shows recent debates on macroeconomic imbalances and household and government indebtedness in a new light. In particular, low real interest rates may be the new normal. The debt of the young provides an alternative outlet for the retirement savings of the old; preventing the accumulation of debt, for example through macroprudential policy, leads to a bigger fall in interest rates.

[1]Preliminary and incomplete. This paper does not represent the views of the Bank of England. It has benefited from helpful discussions with Gauti Eggertsson, Romain Ranciere, Helene Rey and especially Silvana Tenreyro, and comments from seminar participants at LSE, the Bank of England, the Centre for Macroeconomics, the 2014 Money Macro and Finance conference and the IMF Research Department. All errors are my own.

Non-technical summary

The financial crisis that began in 2007 pushed the interest rates set by central banks in much of the industrialised world close to zero. But 'real' interest rates - interest rates minus expected inflation - had been trending down across the industrialised world for at least twenty years before this, and had already reached historic lows on the eve of the crisis.

There have been many explanations for this fall in industrialised-world interest rates, among which are three leading candidates. The first is demographics - in particular a rise in the share in the population of middle-aged people who tend to save a lot, as baby-boomers have entered late middle age. The second is inequality, whereby a rise in the share of income or wealth going to rich people, who tend to save more, has raised total saving. And the third is emerging markets, which in some cases save a lot and have lent these savings to the rich world.

Each of these explanations has merit. But what they all have in common is a rise in domestic or foreign saving as a cause of the fall in interest rates. Interest rates are the price of savings, so an increase in the supply of them reduces the price.

But all savings must ultimately fund investment. So they all predict, therefore, a rise in investment in the industrialised world. But in contrast, the share of investment in total expenditure has fallen across the industrialised world over the past thirty years, a fall which again long predates the recent financial crisis. Furthermore, alongside the fall in interest rates, much of the industrialised world saw house prices and household debt rise to historic highs before the crisis.

This paper fleshes out a new, complementary explanation for the falls in real interest rates, rises in household debt and falling investment rates across the industrialised world. The story is based on the widespread fall in the price of investment goods - the machines, equipment and buildings that firms buy - relative to the prices of other things the economy produces. This fall has reduced the demand for savings, rather than the supply.

I present a simple model in which households need to save for retirement. As the price of investment goods falls, a given quantity of retirement saving buys more of them. If it is hard for firms to use more machines in place of workers, the increase in the number of machines will mean the extra production each new machine generates will fall a lot. Firms will want to spend less on investment, reducing the competition for households' savings and therefore the price that households receive for them, which is the interest rate.

Lower interest rates make it cheaper to buy houses. If houses are in fixed supply, the price of housing gets bid up and, given that housing is bought on credit early in life, household debt increases too. Housing naturally becomes an alternative destination for retirement savings as machines get cheaper. So does the debt of the young, which prospective retirees invest in (implicitly through banks and pension funds) and then live off when older. This means that preventing a rise in household debt could lower interest rates further.

I test the model by comparing across countries movements in the price of investment goods, the share of investment in the economy, household debt and house prices, and find support for its assumptions and predictions. If the model is right, real interest rates may stay low in the future, even if investment goods have stopped getting cheaper.

1 Introduction

The financial crisis that began in 2007 pushed central banks in much of the industrialised world to the zero lower bound on nominal policy rates. Much ink has been spilled about how this happened, what central banks should have done when they got there, and how to avoid it happening again. But real interest rates had been trending down across the industrialised world for at least twenty years before this, and had already reached historic lows on the eve of the crisis (Summers (2013), King and Low (2014)). Alongside this fall in interest rates, much of the industrialised world saw house prices and household debt rise to historic highs before the crisis. While these series have subsequently fallen back somewhat, they appear at the time of writing to have stabilised at elevated levels in relation to GDP and real incomes in many countries.

There have been many explanations for this fall in industrialised-world interest rates, among which are three leading candidates. The first is demographics - in particular a rise in the weight of high-saving age-groups as baby-boomers enter late middle age. The second is inequality, whereby a rise in the share of income or wealth accruing to the high-saving rich has raised aggregate saving. And the third is emerging markets, whereby an excess of saving in the developing world has pushed down on rich-world interest rates.

Each of these explanations has merit. But what they all have in common is a rise in domestic or foreign saving as a cause of the fall in interest rates. They all predict, therefore, a rise in investment in the industrialised world.[1] But in contrast, nominal investment rates have fallen sharply across the industrialised world over the past thirty years, a fall which again long predates the recent financial crisis.

This paper fleshes out a new explanation for the falls in real interest rates and rises in household debt across the industrialised world, complementary to those which rely on higher saving, but which also explains the fall in investment rates. The story is

[1] With the caveat that some demographic models that featuring slowing population growth may predict falling investment rates

based on the widespread fall in the price of investment goods relative to consumption over the past thirty or so years documented in Karabarbounis and Neiman (2014). I extend their data back in time for some countries, to show that this fall has not been a feature of the very long run, but rather began a few years either side of 1980.

In the model, households need to save to provide for retirement. The corporate sector invests the savings of the household sector in capital goods. If the price of capital goods falls, a given quantity of savings can buy more capital goods, raising the return on investment for a given marginal physical product of capital. But the increase in the volume of capital goods lowers the marginal product, thereby lowering the return on investment. The net impact of these two effects depends on the curvature of the production function.

I parameterise the model with a less-than-unit elasticity of substitution between labour and capital, in line with most estimates but at variance with those in Karabarbounis and Neiman (2014). Consistent with the predictions of the model at these parameter values, I present cross-country evidence showing that nominal investment rates have fallen further in countries where the relative price of investment has fallen further.

Depending on parameterisation and the timing convention in the model, the dynamics of the transition to this new steady state can involve a temporary rise in interest rates, as households attempt to bring forward the extra consumption afforded by the fall in the relative price of capital. This provides a new interpretation of the period of historically high world real interest rates experienced in the 1980s. More generally, the transitional dynamics operate for some decades both before and after the change in relative prices. For example, I find that the fall in the relative price of capital has been particularly good for the baby boomer generation whose housing wealth has been revalued by the shock.

But the new steady state is one of lower interest and investment rates and higher household debt ratios, even after investment goods prices have stopped falling. Lower interest rates reduce the user cost of housing, boosting housing demand. Housing supply is fixed, so house prices (or at least land prices) must rise. Houses are bought early in life and largely on credit, so household debt also increases. Acquiring these debt claims is an alternative form of retirement saving, so the capital investment rate falls in the steady state, as we see in the data. The model's implications for household debt and house prices receive qualified support in cross-country econometric analysis. I extend the model to allow for bequests, and for heterogeneity in the bequest motive. My core findings are robust to this modification. Furthermore I find that the real interest rate moves in the *opposite* direction to wealth inequality, in contrast to Piketty et al. (2014), but moves in the same direction as consumption inequality.

These findings cast recent debates on macroeconomic imbalances and household and government indebtedness in a new light, and have important policy implications. Some prominent policymakers (see, for example, Ingves (2014)) are seeking to prevent what they see as 'excessive' levels of household debt. But if low rates of interest and investment, accompanied by pressure for governments and households to become indebted, represent the transition to a new steady state in which the corporate sector's demand for household savings is weak, then attempts by macroprudential or monetary authorities to prevent this may be futile or counterproductive. I show that preventing the rise in household debt in response to a fall in capital goods prices makes interest rates fall further in response to the initial shock.

The mechanism in this paper builds on a long history of related ideas in the literature. Summers (2013) recently raised the issue of the pre-crisis falls in real interest rates and the possibility that they would stay low for an extended period in the future. But the idea that capitalist economies could be plagued by chronically low returns on capital, and that this could result from an overaccumulation, in some sense, of physical capital goes back at least to Marx (1867) and Hansen (1938). The fall in capital goods prices in the face of a need for retirement savings creates a form of asset shortage reminiscent of Caballero et al. (2008), which is satisfied by the endogenous creation of debt claims on the young. The focus on the fall in the relative price of investment goods builds on the important contribution of Karabarbounis and Neiman (2014), whose data I draw on for this study. Rey (2014) and Summers (2014) have, among others, linked secular stagnation to falls in the relative price of capital goods.

Two papers are particularly close, methodologically speaking, to the present study. Giglio and Severo (2012) examine the effect of a change in production technology in an OLG model and find that the conditions in which asset bubbles may exist are modified. Like the present study, Eggertsson and Mehrotra (2014) address the issue of secular stagnation in an OLG model. They show that a tightening of the debt limits facing young households, reduced population growth and increased income inequality can reduce the equilibrium real interest rate in such a model, and explore the consequences for resource utilisation in a sticky price model. They also show that falling relative capital goods prices can lower the real interest rate. Relative to that study, this paper gives conditions under which interest rates can remain low even after capital goods prices have stopped falling; IMF (2014) finds that relative prices have been stable since 2002. Furthermore, this paper derives the implications of lower capital goods prices for house prices, household debt and wealth inequality. But unlike Eggertsson and Mehrotra (2014), this paper says nothing about resource utilisation or nominal variables.

The remainder of this paper is structured as follows. Section 2 sets out the key facts the model aims to explain. Section 3 describes the core economics of the paper in the simplest possible model. Section 4 describes the baseline model. Section 5 shows the results of model simulations in which I vary the relative price of investment and generate movements in interest rates, investment rates and household debt which are qualitatively similar to those presented in section 2. Section 6 examines the sensitivity of these findings to parameter values, and extends the model to allow for bequests. Section 7 extends the model to allow for bequests, heterogeneous agents, intangible capital and open-economy considerations. Section 8 adduces some econometric evidence in support of the model. Section 9 concludes.

2 Motivating facts

This section sets out the key stylised facts that the model aims to connect. I focus on the widest possible set of industrialised countries for each data series, but also, where possible, show data for a subset consisting of the 11 advanced countries [2] for which the EU-KLEMS database has sufficient data to calculate long time-series of nominal and real capital-GDP ratios.

2.1 Falling real interest rates

Ex-ante real interest rates can now readily be measured in many industrialised countries with reference to the yields on index-linked government liabilities. However, these securities were not issued before the 1980s, complicating the the measurement of ex-ante real interest rates before then. IMF (2014) presents an attempt to solve this problem by constructing a parametric model of inflation expectations and subtracting the result from the yields on nominal government liabilities.

Figure 1 shows the result for the UK and the US. The figure shows that interest rates have been trending generally downwards for the 30 years since their recent peak in the early 1980s. The model-based series in IMF (2014) suggest that US ex ante real rates were close to current levels in the early 1970s, fell below zero in the middle part of that decade, before rising sharply in the late 1970s-early 1980s. King and Low (2014) and Laubach and Williams (2003) (updated to 2014) both also show declining real interest rates from the 1980s to the 2008 crisis.

2.2 Rising household debt ratios

Figure 2 shows the an index of the ratio of household debt to GDP since 1970 for a broad sample of industrialised countries and our restricted sample of 11 countries.

[2] Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US

The figure shows a rise in the average ratio of around 50pp since 1970.

2.3 Price and quantity of capital investment and stock

Figure 3 shows the simple average across OECD countries and across our restricted sample of the ratio of nominal investment to nominal GDP. The nominal investment rate has been trending downwards since at least the mid-1970s. Figure 4 shows that the corresponding stock ratio (the current replacement cost of the capital stock as a proportion of GDP) had also fallen from nearly 4 times annual GDP around 1980 to nearly 3 times by 2007 for the 11 countries in the EU-KLEMS database for which data are available.

Figure 5 shows the real investment - GDP ratio across the same two sets of countries since 1970. The series show no strong trend over the whole sample, although there is weak evidence of an upward trend since the early 1980s. Figure 4 shows that the ratio of the real capital stock to real GDP (both at 1995 prices) has been trending upwards since the 1970s.

These divergent patterns in the nominal and real ratios are of course a manifestation of a trend fall in the price of investment goods relative to consumption or GDP, documented in Karabarbounis and Neiman (2014). Figure 6 shows four series of the ratio of the investment deflator to the consumption deflator. The red and blue lines are taken from the respective countries' national accounts data. The green line is the average change across all the countries in the dataset, and the purple line is the average among our restricted sample in this dataset. All three lines show that the relative price of investment goods has been falling in recent decades, with a fall of perhaps 30% since the mid-1970s. The longer series show that, prior to this fall, there has not been a secular trend in this relative price.

2.4 Factor shares

Karabarbounis and Neiman (2014) document a fall in the corporate and whole-economy labour share since 1975, within a large number of industries and countries. In the baseline model presented below, there are no pure profits in the economy and only two factors of production - capital and labour. In such a world, a falling labour share must imply a rising profit share, which in turn is equal to the product of the average return on capital and the capital-output ratio

$$\frac{\Pi}{Y} = \frac{\Pi}{Kp_K}\frac{Kp_K}{Y} \tag{1}$$

Figures 1 and 4 show that the real interest rate and the nominal capital-output ratio have typically fallen in industrialised countries over past thirty years which, according to the simple equation above, would generate a fall in the labour share. However, there are a number of explanations for the apparent discrepancy between the trends in the labour share, the real interest rate and the capital output ratio:

- The capital-output ratio could be somehow mismeasured, perhaps because of the omission of intangible or nonreproduced factors of production, and has in fact not fallen over time.
- The marginal and average returns on capital correspond to the real interest rate and the average profit rate respectively. There could be an increasing wedge between them, or an increasing wedge between the marginal product of capital and the real interest rate in financial markets, perhaps because of corporate taxes or physical depreciation.
- Corporate profits include a component of 'pure profit' as well as remuneration for capital investment, corresponding, for example, to producer markups over marginal cost. These markups could have risen over time.
- Relatedly, profits could be remunerating highly-skilled or managerial labour, e.g. through the granting of share options, such that the labour share, broadly

conceived, has not fall as much as the wage share would suggest.

The baseline version of the model incorporates none of these features. Equation 1 holds in the model and therefore, in generating a falling real interest rate and a falling capital-output ratio, it also produces a rise in the labour share.

3 Real interest rates, capital goods prices and the curvature of the production function

Other things equal, lower capital goods prices p raise the return on capital when denominated in consumption goods: a foregone consumption good buys more capital goods, so for a given marginal product of capital, the return on investment

$$r = \frac{1}{p}\frac{\partial Y}{\partial K} - \delta \tag{2}$$

is higher.

But other things will not be equal - the fall in capital goods prices will mean that a given volume of savings will finance more of them, pushing down on the marginal product of capital to an extent that depends on the curvature of the production function. Whether the volume effect outweighs the price effect depends on the curvature of the production function. And savings may respond to the resulting change in interest rates in either direction, depending on the properties of the utility function. To crystallise these issues before I present the baseline model, this section of the paper analyses the role of the curvature of the production and utility functions in the simplest possible model with variable capital goods prices.

3.1 Simplest possible model

Consider a world populated by an identical series of overlapping generations, each of which lives for two periods.[3] Each generation of households has a standard isoelastic utility function defined over consumption in each generation of life

$$U(c_1, c_2) = \frac{c_1^{1-\theta}}{1-\theta} + \beta\frac{c_2^{1-\theta}}{1-\theta}$$

Households supply one unit of labour at wage rate W in the first period of life, and can lend money to firms at net interest rate r to provide for their retirement. So their intertemporal budget constraints are as follows

$$c_2 \leq (W - c_1)(1+r)$$

Young households' saving in the first period of life as a fraction of their wage income can be shown to be given by [4]

$$s = \frac{W - c_1}{W} = \frac{\beta^{\frac{1}{\theta}}(1+r)^{\frac{1}{\theta}-1}}{1+\beta^{\frac{1}{\theta}}(1+r)^{\frac{1}{\theta}-1}} \tag{3}$$

This familiar expression shows that the sign of the slope of the savings schedule in $\{s, r\}$ space depends on the intertemporal elasticity of substitution $\frac{1}{\theta}$. When this substitution elasticity is high (i.e. above unity), a fall in interest rates causes a fall in savings, as agents substitute away from relatively expensive retirement consumption. Infinite-horizon households pin the interest rate down at $r = \frac{1}{\beta} - 1$, and are thus equivalent to OLG households with linear period utility functions. When the elasticity is below unity, retirement saving is akin to a Giffen good: lower interest rates *raise* the savings rate out of wages, as the desire to offset the negative effect of lower interest rates on retirement consumption outweighs the higher price of it. When the elasticity is exactly one, these two effects cancel and the savings schedule

[3]Overlapping generations are necessary because the interest rate in an infinite horizon model would be pinned down by the household discount factor

[4]See Appendix for derivation

is vertical.

Turning to the determination of factor prices, firms hire labour and borrow funds from young households, buy capital goods (which depreciatate at rate δ) at relative price p and maximise profits with them.[5] Factor prices $\{W, r\}$ will therefore be set equal to marginal product in the standard fashion

$$W = \frac{\partial Y}{\partial L}$$
$$r = \frac{1}{p}\frac{\partial Y}{\partial K} - \delta$$

In aggregate, the gross savings of the young will equal the replacement cost of the capital stock

$$pK = sW$$

If we assume a CES production function with elasticity of substitution σ and capital share parameter α we can derive an 'investment schedule' that implicitly maps s into $\{r, p\}$[6]

$$s = \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{1 - \alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \tag{4}$$

Which way does the interest rate schedule slope in $\{s, r\}$ space? There are two effects. The effect in the numerator is negative for the standard reasons: for given capital goods prices, more savings reduces the marginal product of capital and hence the interest rate. The effect in the denomnator is of ambigious sign, and comes through the labour share (for a Cobb-Douglas function $\sigma - 1 = 0$ it is absent). For low σ, an increase in r reduces the denominator, raising the quotient. This is because the saving rate is expressed here as a fraction of wages and when $\sigma < 1$, higher interest rates are associated with a lower labour share. To save enough for a given

[5] We can for now think of a class of final goods firms turning intermediate goods into consumption goods one-for-one or into capital goods at rate p^{-1}. This will be made more explicit when describing the full model in section 4

[6] Derived in the Appendix

volume of capital goods, a lower labour share must mean a higher saving rate. For reasonable parameter values, the effect on the numerator will dominate, such that the investment schedule slopes down in $\{s, r\}$ space.

The derivative of the saving rate with respect to the price of capital goods p is the same sign as $1-\sigma$. Consider a fall in the relative price of capital goods of x percent. Holding the marginal product of capital constant, the return on investment increases by x percent as each consumption unit of investment buys x percent more capital goods. But because of the price fall a given volume of savings can finance x per cent more capital goods, and the marginal product of each will fall by $\frac{x}{\sigma}$ per cent, such that the sign of the effect is equal to the sign of $1-\sigma$.

Figure 7 depicts graphically how the effect of a fall in p is governed by the effects of the curvature parameters $\{\sigma, \theta\}$. The top left panel shows the effect of a fall in the relative price of investment goods on the investment schedule for values of σ either side of unity. The top right panel adds an upward-sloping savings schedule, corresponding to a relatively elastic utility function. In this case, the rates of interest and of investment/saving will covary positively, with the sign of the change once again depending on the sign of $dp\,(1-\sigma)$. The cases of a small open economy or of infinite-horizon households correspond to a horizontal saving schedule - no change in interest rates and a change in investment rates of the same sign as $dp\,(1-\sigma)$. The bottom-left panel depicts a highly inelastic utility function. In this case, the changes in the rates of interest and saving are of opposite sign, but the former is still the same sign as $dp\,(1-\sigma)$. Finally, the bottom-right panel shows an extreme case in which the savings schedule slopes downward but is *shallower* than the investment schedule. In this case a fall in the relative price of capital would lead to a fall in the investment rate and a rise in interest rates if $dp\,(1-\sigma) < 1$.

4 The baseline model

In this section we augment the heuristic model above with an intermediate period of working life, and with the requirement for households to buy a house when young. This enables us to analyse the effect of capital goods prices on house prices and household debt, and how the existence of both alters the determination of interest rates.

4.1 Households

The economy is closed and comprises three overlapping generations of constant and equal size. Each generation has a standard separable CES utility function over consumption and 'housing'[7]

$$U(c_1, c_2^{'}, c_3^{''}, h) = \frac{1}{1-\theta}\left(c_1^{1-\theta} + \beta_2 c_2^{'1-\theta} + \beta_3 c_3^{''1-\theta}\right) + \phi \frac{h^{1-\gamma}}{1-\gamma} \tag{5}$$

where I denote leads one and two periods hence with primes and double-primes respectively. In period 1 (young adulthood), the household supplies η units of labour inelastically (remunerated at wage W), consumes goods and buys a house. She can borrow or save a net amount S_1 at rate r. In period 2 (middle age), the household remains in said house, supplies $(1-\eta)$ units of labour, and can again borrow or save $S_2^{'}$. In period 3 (retirement), she sells her house and consumes the proceeds plus her accumulated savings. So each of the three periods is associated with a budget constraint as follows

$$c_1 + hp_h + S_1 = \eta W \tag{6}$$

$$c_2^{'} + S_2^{'} = (1-\eta)W + (1+r)S_1 \tag{7}$$

$$c_3^{''} = (1+r^{''})S_2^{'} + hp_h \tag{8}$$

[7] Housing is in fixed supply so might be more usefully thought of as land, or more generally any non-produced asset that yields utility

Forming and solving the Lagrangean yields standard consumption Euler equations thus

$$\frac{c_1^{-\theta}}{(1+r')(1+r'')} = \frac{\beta_2 c_2^{'-\theta}}{(1+r'')} = \beta_3 c_3^{''-\theta} \tag{9}$$

We also get a housing demand equation that depends on future house prices and consumption as you would expect

$$\phi h^{-\gamma} + \beta_3 c_3^{''-\theta} p_h^{''} = c_1^{-\theta} p_h \tag{10}$$

This is intuitive. The LHS is the marginal utility of housing plus the discounted marginal utility of the retirement consumption paid for by the sale of the house. The RHS is the consumption utility cost of buying a unit of housing.

4.2 Firms

A measure of perfectly competitive firms produce intermediate goods, combining capital and labour with a CES production technology

$$Y = A[(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \tag{11}$$

These intermediates can then either be consumed directly, or transformed into capital goods at rate p units of intermediate for every one unit of capital. The relative price of investment goods - the key exogenous parameter in our model - is therefore p. This means of introducing investment-specific technological change is isomorphic to that in Greenwood et al. (1997).

Wages are set equal to the marginal product of labour

$$W = \frac{\partial Y}{\partial L} \tag{12}$$

Firms equate the user cost of capital to its marginal product, both denominated in consumption goods

$$1 + r^{'} = \frac{1}{p_K}\frac{\partial Y^{'}}{\partial K^{'}} + \frac{p_K^{'}}{p_K}(1-\delta) \tag{13}$$

4.3 Market clearing

At the end of each period, the net savings of households of young and middle age are transformed into next period's capital stock (at this period's relative prices), such that the following capital-market clearing condition holds in stock terms

$$S_1 + S_2 = K^{'} p_K \tag{14}$$

There is a fixed measure $\bar{H}$ of housing or land for each of the first two generations to live in, so that in equilibrium

$$h = \bar{H} \tag{15}$$

5 Results

5.1 Parameterisation

Each of the three periods of adult life lasts twenty years. The discount factors $\{\beta_1, \beta_2\}$ and capital share parameter α are set to hit an annualised steady-state interest rate of 3% and a capital share of one-third respectively. The depreciation rate δ is set at set at the standard value of .05 in annualised terms. The elasticities $\{\theta, \gamma\}$ in the utility function are set to unity (log utility), while the production elasticity σ is set to 0.7, consistent with most estimates and with the econometric evidence presented in section 8. The importance of these curvature parameters are explored in sensitivity analysis below. I setϕ to hit a sensible value for the ratios of housing wealth to GDP respectively.

A fall in p_K amounts to an improvement in the overall level of technology, in the sense that the lower is p_K, the larger is the total volume of consumption and investment goods a given factor endowment can produce. However, the overall growth rate of TFP has not notably accelerated over the past several decades. So when considering changes in p, I change A so as to keep potential GDP unchanged given existing factor endowments. The fall in capital goods prices in the simulations prompts an accumulation of capital goods, so potential GDP rises nonetheless.

5.2 Solution method

We first solve for the steady state of the model for a given value of capital goods prices p. An initial assumption is made about house prices and the savings of each generation $X^0 = \{S_1^0, S_2^0, p_h^0\}$, which implies a certain constellation of factor prices $\{W^0, r^0\}$. Household behaviour is then optimised taking these prices as given, the resulting optimal values of $X^* = \{S_1^*, S_2^*, p_h^*\}$ are computed, and the initial guess is updated toward them - i.e. $X^1 = \lambda X^* + (1-\lambda) X^0$, where $\lambda \in (0,1)$ is a gain parameter. This process is repeated until the solution converges to a fixed point, i.e. until $X^n \approx X^{n-1}$.

To assess the dynamic effects of a change in p, I consider a simulation path of sufficient length T that the economy will be at the steady state at the beginning and end of the simulation, with the exogenous changes to p occurring in the middle. I first calculate the steady state in each period t of the simulation $\{X_t^{ss}\}_{t=1}^{T}$, given the extant values of the exogenous parameters. I then optimise the behaviour of each generation t, taking the behaviour of the other generations as given, obtaining $X_t^* \left(\{X_s^{ss}\}_{s \neq t}\right) \forall t$. As above, the initial guess is updated towards this solution until it converges. I verify that the model converges to the steady state well inside the endpoints of the simulation.

5.3 Comparative statics

The blue lines in figure 8 show the the effect of varying the relative price of capital goods p on the steady state of the model at the baseline parameter values. In the baseline model, the annualised real interest rate falls by 20 basis points. The nominal investment rate falls about 1 percentage point in response to the lower relative price of capital (bottom left panel). This implies a somewhat upward-sloping savings schedule in the model, notwithstanding the assumption of log utility which, in the simple two-period model of section 3. This is because the fall in interest rates lowers the user cost of housing for a given house price. House prices must rise to choke off the resulting increase in demand - by about 10 per cent in the baseline, relative to GDP. To fund the purchase of more expensive houses, the ratio of the net debt of young households to GDP increases by about 20 percentage points (top right panel). Housing is a store of value as well as a consumption good, so the purchase of a house is an alternative to the purchase of capital goods as a means to fund retirement consumption; in general equilibrium, part of the money that would have gone to fund the purchase of capital goods is instead lent to the young to fund their house purchase, who live off the sale proceeds in old age.

5.4 Dynamic results

What are the dynamic consequences of the experiment considered above? We analyse the dynamic impact of a 30% fall in the relative price of capital goods over one model period (20 years). The exercise of mapping into the data 20-year model time periods, each of which contains a series of supposedly discrete events, is somewhat nuanced. According to the model's timing conventions, savings accumulated at the end of period $t-1$ become productive in period t. An important question is which period's capital goods prices are used to convert savings into capital goods, and back into consumption goods. In the baseline simulation shown here, consumption foregone in period $t-1$ becomes productive capital goods in period t, with the conversion happening at period t prices. For this reason, the first period of low capital goods prices - period 10 in the charts - corresponds roughly to the 1980s and 1990s in the data. Interest rates are measured ex ante - so the period 10 interest rate is the return on savings made in the 1990s, paying a return in the 2010s. The interest rates observed at the time of writing (the mid-2010s) correspond to period 11 of the model. Alternative timing conventions are possible - for example turning period $t-1$ savings into period t productive capital goods at period $t-1$ prices - and are explored in the sensitivity analysis below. Timing conventions would matter less in a model with shorter time periods or in periods with more stable capital goods prices, and of course do not matter at all when analysing the steady state.

Figure 9 shows the results of the baseline dynamic simulation. In each panel, the blue line shows the relative price of capital goods produced in the period in question. The top left panel shows the path of the ex-ante real interest rate. The ex-ante interest rate earned on savings made at the end of period 9 (before the fall in capital goods prices) rises. This corresponds to the late 1970s, a period of rising world interest rates. The middle left panel shows that the saving (or investment) rate falls before the shock hits, recovers partially, and then resumes its fall. A fall in the saving rate combined with a rise in the interest rate is indicative of a shift inwards in the saving schedule. Consistent with this, the top right panel shows the path of household debt, which begins to rise in advance of the fall in capital goods prices. Younger generations can look forward to funding their retirement in part by selling more expensive houses, and thus begin consuming and dissaving more. The middle-right panel shows that the rise in housing wealth in relation to GDP takes several generations to be completed.

The bottom right panel shows that the profit share initially rises and then falls when the shock hits, as the fall in the interest rate outweighs the rise in the capital-output ratio at the assumed parameter values. How this feature of the model relates to the evidence is discussed in section 3.

Finally, the bottom-right panel shows the response of output and the consumption of each age group. Output initially falls very slightly as capital is decumulated, but eventually rises as the fall in capital goods prices affords a larger real capital stock.

GDP rises despite the assumed fall in Hicks-neutral productivity in the intermediate goods sector, which is calibrated to be sufficient to offset the improvement in technology that the fall in p represents without any increase in factor endowments. The consumption of the young generations rises sooner, and by more, than that of older generations, such that in the steady state the age-consumption profile is flatter. It rises more because the steady-state interest rate is lower, encouraging households to consume earlier, and sooner because households who are young on the eve of the shock anticipate capital gains on their house purchases. This pattern of capital gains can also be observed in the consumption of the old - the generation that is old in period 11 (i.e. the baby boomers) consumes more in retirement than any other retired generation, because it enjoyed the biggest capital gains on housing, buying them relatively cheaply in the 1960-1970s and then trading down in the early 21st century.

Overall, the simulation results generate a qualitatively similar pattern in the real interest rate, housing wealth, the real and nominal capital-output ratios and the household debt-GDP ratio to those which we have observed over the past four decades. The shock is particularly beneficial for the baby-boomer generation. Furthermore, the simulations provide forecasts of what may happen in years to come. In particular, even if the relative price of capital has stopped falling, the interest rate may continue to fall somewhat, as the capital deepening process brought on by the fall in the relative price of capital runs its course. And future generations of retirees will consume less than the current one.

6 Sensitivity analysis

6.1 Housing and debt

The availability of debt and housing as alternative savings vehicles attenuates the fall in interest rates in the model. This is illustrated in figure 10, which considers two alternative regimes for household debt alongside the baseline model. The green line represents a regime in which household debt is forbidden (the net savings of the young must be nonnegative, i.e. $S_1 > 0$). The red represents a simulation in which there is an upper bound on debt that binds at an intermediate level of p. These constraints attenuate the rise in house prices (bottom right panel), as young consumers cannot spread the extra cost of housing over their lives. The fall in the aggregate savings rate (bottom left panel) is also attenuated, as the debt of the young and more expensive houses are less readily available as savings vehicles. Higher savings means more capital and thus lower real interest rates - the top left panel shows that, without household debt, a fall in capital goods prices gives rise to a fall in real interest rates of about 60 basis points, i.e. about three times larger than in the baseline simulation.

Figure 11 shows the dynamic effects of a shock to p when household debt is prohibited. The key difference is that the path of interest rates is now monotone. Interest rates do not rise ahead of the shock because young households are not able to borrow to bring forward consumption. The investment rate follows the same falling-rising pattern but now settles at a higher rate than before the shock, as household debt is no longer available as an alternative destination for retirement savings. The bottom-right panel of the figure shows that it is now the middle-aged rather than the young whose consumption rises the most. As before, lower interest rates dissuade retirement saving, but the young cannot respond by dissaving more; only the middle-aged can respond, by reducing retirement consumption at the expense of higher consumption in middle age.

6.2 Parameterisation

6.2.1 Curvature of the production function

The key parameter in this model is the elasticity of substitution in the production function σ between capital and labour. Figure 12 shows how the impact of p

on the steady state of the model depends on the elasticity of substitution between capital and labour σ. When the production function is Cobb-Douglas, the relative price of capital has no effect on the interest rate, house prices, household debt or the investment rate in the steady state. The heuristic model presented in section 3 explains why. The volume of capital goods bought with a given quantity of consumption goods is inversely proportional to the relative price. With a Cobb-Douglas production function, the marginal product of capital is inversely proportional to the real capital-output ratio, so these two effects exactly offset. There are nonetheless some dynamic effects during the transition to lower relative capital prices (figure 13). In particular, the interest rate rises and then falls, as consumers attempt to bring forward some of the higher consumption afforded by the lower capital goods prices.

Figures 12 also the effect of the fall in relative capital goods prices on the steady state of the model when the $\sigma = 1.3$, in line with the estimates in Karabarbounis and Neiman (2014), and symmetric around the Cobb-Douglas case with the baseline value of $\sigma = 0.7$. Everything now goes in the opposite direction to the baseline: interest rates and the investment rate rise, while house prices and the household debt ratio fall. Figure 14 shows the dynamics of the transition. The interest rate and investment rate overshoot their long run value during the transition. The profit share rises and the household debt ratio falls monotonically.

For all the values of σ considered here, the interest rate falls in the period after the shock hits, and in this sense the model can account qualitatively for interest rates being lower now than during the early 1980s. However, the amount further that the interest rate is expected to fall, and where it will settle relative to its previous value, depend crucially on σ, and in particular whether it is bigger or smaller than one. Furthermore, the model can only account for rising house prices and debt when σ is below one. This result can be viewed in one of two ways. Either the model is 'wrong', or at least insufficiently general to account for the facts it is trying to explain without particular values for key parameters. Or it helps to identify a value for σ that is in line with the range of estimates reported in Chirinko (2008) and with the econometric evidence presented in section 8, but not with the findings in

Karabarbounis and Neiman (2014).

6.2.2 Parameters of the utility function

The curvature parameters in the utility function are important because, as discussed in section 3, there could be no effect on interest rates in a model with overlapping generations and infinitely elastic utility (and hence savings) functions. Figure 15 shows how the effect of p on the steady state of the model depends on the the curvature parameter $\{\theta\}$ in the utility function. In these experiments we recalibrate β to hit the same initial interest rate but leave fixed the other parameters, in particular the utility function parameters $\{\gamma, \phi\}$. The figure shows that, if the utility function is more inelastic (setting $\theta = 1.5$), the interest rate falls by less, while house prices and household debt fall by somewhat more. This is in sharp contrast to the simplified model presented in section 3, which predicts that the interest rate varies by less when the utility function is more elastic. The reason for the discrepancy is the addition of housing to the model. If we omit housing from the model (see figure 16), the effect of the curvature of the utility function is in line with section 3: more elastic utility means that the interest varies by less.

6.3 Timing conventions in the model

Time periods in this model are 20 years long. Within-period timing conventions may accordingly have an important effect on the dynamics of the model. In particular, in the baseline simulations above, savings accumulated in period t are assumed to be turned into capital goods at the end of period t, at period t prices, and yield a physical return in period $t+1$ before implicitly being transformed back into intermediate goods at period $t+1$ prices. A reasonable alternative would be to assume that savings made in period t are transformed into capital goods at the beginning of the following period $t+1$, at the prices extant in that period, and then transformed back into intermediate goods in the same period at the same prices. There would then be no price depreciation component in interest rates.

Figure 17 shows the dynamic effects of a change in capital goods prices when the timing convention is altered in this manner. In a steady state with constant capital goods prices, these timing conventions clearly would not affect the steady state, so only the dynamic solution is shown. The most striking difference is in the path of interest rates, which fall in the period before the shock, rise above their pre-shock value and then fall again to their new, lower steady state value. Depending on when one dates the fall in capital goods prices that took place over this period, this pattern may rationalise the relatively low world interest rates observed in the early to mid-1970s. Figure 6 shows that the fall in capital goods prices was relatively steady over the period between the mid-1970s to mid-1990s, such that either timing assumption is reasonable in a model with 20-year periods.

7 Extensions

7.1 Bequests

In the baseline model, households spend all their wealth by the end of their lives, including their housing wealth. In practice, bequests form a large part of households' total resources and a large fraction of GDP is bequeathed in any one year (Piketty (2011)). Retirees often live in owner-occupied housing until the end of life (Yang (2009). These features can be introduced into our framework by adding bequests $\{b', b\}$ respectively given and received to the utility function and budget constraints as follows[8]

$$U(c_1, c_2^{'}, c_3^{''}, h, b) = \frac{1}{1-\theta}\left(c_1^{1-\theta} + \beta_2 c_2^{'1-\theta} + \beta_3 c_3^{''1-\theta}\right) + \phi\frac{h^{1-\gamma}}{1-\gamma} + \xi\frac{b^{'1-\zeta}}{1-\zeta} \tag{16}$$

[8] These are 'warm glow' preferences over bequests. Households still care about the consumption value of their assets in retirement, because they evaluate their bequests in consumption rather than utility terms. Adding later generations' utility directly to the utility function would collapse the model into an infinite horizon setup.

$$c_1 + hp_h + S_1 = \eta W \tag{17}$$

$$c_2^{'} + S_2^{'} = (1-\eta)W + (1+r)S_1 + b \tag{18}$$

$$c_3^{''} + b^{'} = (1+r^{''})S_2^{'} + hp_h \tag{19}$$

Again forming and solving the Lagrangean we get Euler equations thus

$$\frac{c_1^{-\theta}}{(1+r')(1+r'')} = \frac{\beta_2 c_2^{'-\theta}}{(1+r'')} = \beta_3 c_3^{''-\theta} = \xi b^{'-\zeta} \tag{20}$$

Piketty (2011) finds that bequests in France are around 15% of GDP. With log utility bequests will be a constant fraction of old-age consumption. Consistent with this, we set $\xi = 0.75\beta_3$. Figures 18 and 19 show that, qualitatively speaking, the steady state and dynamic solutions of the model are unaffected by this change, although the change in interest rates is attenuated somewhat. The levels of debt and house prices are nonetheless higher at any given level of p. Households save for bequests much like they save for retirement consumption. By the time the middle-aged receive their requests, they themselves are on the cusp of retirement and are therefore no long debtors. Anticipating bequests in middle age, the young accumulate more debt and push up house prices.

7.2 Heterogeneous agents

In the versions of the model presented above, the only dimension along which agents are heterogeneous is age. The change in p affects different generations differently, but there can be no intra-generational inequality. In reality, inherited wealth is distributed highly unequally across the members of any given generation. The asset-price consequences of a change in p are accordingly likely to have consequences for intra-generational inequality. To study this, we simulate a version of the model in which the population is divided into two kinds of dynasty - life-cycle households without a bequest motive, as in the baseline model, and households with a bequest motive. For illustrative purposes, we set the proportions of each kind to one-half, and apportion to them equal labour endowments (and thus labour income).

Figure 20 sets out the dynamic consequences of a change in p on the ratio of the consumption of households with a bequest motive to life-cycle households. In the steady state before the fall in p, households with the bequest motive accumulate more wealth (about one-and-a-half times as much) and consume about 13% more than those without it, even though they have the same labour income. In the long run, as in the baseline and dynastic models, the fall in p lowers the real interest rate and raises house prices. Consumption inequality falls because lower interest rates reduce the returns to inherited wealth, and thus the extra consumption that can sustainably be financed from bequests. But the fall in interest rates revalues non-produced assets (land) such that the wealth-income ratio rises. Households that do not receive bequests have only life-cycle saving, which falls somewhat, as a source of lifetime wealth, so wealth inequality increases. In this model, therefore, and in contrast to Piketty et al. (2014), r moves in the *opposite* direction to wealth inequality, but moves in the same direction as consumption inequality.

The dynamic consequences of the shock to p are non-monotone and vary a great deal according to date of birth. Households in bequest-giving dynasties who are young on the eve of the shock do especially well, because they receive a disproportionate slice of the one-off capital gains that accrue to asset holders.

7.3 Open economy

The baseline model in this paper treats the industrialised world as a large, closed economy, with a view to explaining a global trend. The world is of course composed of many economies which are open to trade in goods and financial assets with each other as well as with emerging markets. For any one of these countries, foreign assets are an important store of value, such that we might expect the external position of any given economy to depend on the domestic relative price of its capital goods. The real interest rate in any given country may accordingly not depend to a great extent on the relative price of capital goods in that country, even if the interest rate and the relative price of capital goods are linked at a global level. Furthermore, testing the implications of the model presented above is hampered by the fact that, at a global level, we only have one very short time series when time is denominated in model units, whereas an open-economy version will lend itself to testing along the cross-country dimension. Last but not least, the low-frequency behaviour of the current account dynamics is of independent interest.

To study the open-economy implications of capital goods prices, and to take our model more readily to the data, we therefore consider a simple open-economy version of our baseline model. We assume that intermediate goods, and financial claims denominated in them, are perfectly tradable across borders. They are transformed into consumption and investment goods at home using the domestic technology. This technology can vary across countries, and hence so can the relative price of capital goods. Each country takes the world real interest rate as given. Relative to the closed-economy baseline model set out in section 4, all prices and quantities except the interest rate r acquire country subscripts i, and the only equation to materially change is the asset-market clearing condition (14)

$$S_{1i} + S_{2i} - K_i^{'} p_{Ki} = NFA_i \tag{21}$$

where NFA_i denotes the net foreign assets of country i.

Figure 21 shows the steady-state effect of changing the relative price p_{Ki} of capital goods in country i holding the world interest rate fixed and starting from a position in which, at $p_{Ki} = 1$, the net foreign asset position is zero. The experiment can be thought of as describing the behaviour of a small open economy in which the relative price of capital goods falls by more than the world average. The blue lines show the impact in a closed economy (i.e. the baseline simulation), and the green lines show the small open economy. The top left panel reminds us that, by construction, interest rates do not change. The fall in capital goods prices leads the corporate sector to demand less in the way of investable funds, as we see by the fall in the nominal investment rate (top right panel). Middle-aged households' savings go overseas rather than to young households when the economy is open, such that net foreign assets rise from zero to about 60% of GDP (bottom right), while the household debt ratio is essentially unchanged. House prices relative to incomes (bottom left panel) rise nonetheless, albeit by about 10% rather than the 30% we see in the closed-economy case. This is because the fall in investment and the rise in net foreign assets afford a rise in the consumption-GDP ratio, and consumption and housing demand are positively related.

Figure 22 shows the dynamic behaviour of a small open economy when faced with a 30% fall in the relative price of capital over one period as above, but where world (and hence domestic) real interest rates are fixed. The blue line (read against the left-hand axes) show the relative price of capital, and the green lines show the behaviour of the closed-economy baseline model for ease of comparison. The top left panel shows that, by assumption, the interest rate in the small open economy is unchanged. The two most striking results are in the top and bottom right panels. House prices rise ahead of the fall in capital goods prices (middle panel), in anticipation of higher housing demand after the shock. This raises household debt as young households seek to smooth consumption in the face of higher house prices. However, once capital goods prices fall, output and wages rise sharply, such that the ratio of house prices to wages falls and young households need less debt. Net foreign assets rise sharply around the shock as the interest rate that would prevail in a closed economy falls below the world interest rate: the fall in domestic capital goods prices makes available savings that can fund foreign investments.

8 Econometric evidence

8.1 Evidence on the elasticity of substitution between capital and labour σ

Karabarbounis and Neiman (2014) present a model in which, like the model above, the relationship between the labour share and relative price of investment goods depends on the elasticity of substitution between capital and labour

$$\frac{s_{Lj}}{1-s_{Lj}}\hat{s_{Lj}} = \gamma + (\sigma - 1)\,\hat{\xi}_j + u_j$$

where s_{Lj} is the labour share in country j, ξ_j is the relative price of investment, and hats denote low-frequency, country-specific time trends. They regress the time trend in the labour share on the time trend in the relative price of investment goods, obtain a coefficient averaging 0.28 across datasets and and infer that the elasticity of substitution between capital and labour is 1.28.

Their model also implies an analogous relationship between the nominal investment rate $\frac{I_N}{Y}$ and the relative price of capital [9]

$$\widehat{\frac{I_N}{Y}}_j = \tilde{\gamma} + (1-\sigma)\,\hat{\xi}_j + \tilde{u_j}$$

Similarly, the baseline model presented in this paper predicts that, when the supply of funds is perfectly interest-elastic, as in the small open economy case in section 7.3 or a model with infinite-horizon consumers, the investment rate and the price of investment are related as follows in the steady state

$$\widehat{\frac{pI}{Y}} = c + (1-\sigma)\,\hat{p}$$

These equations motivate regressing the nominal investment rate on the relative price of investment goods across countries as a way of quantifying the crucial parameter σ: if the elasticity is greater than one, we would expect a negative relationship between the relative price of capital and the nominal investment rate as the quantity falls (rises) proportionally more than the price falls (rises). We employ the dataset in Karabarbounis and Neiman (2014) to this end.

Table 1 sets out the results, showing the central estimate for the coefficient on the

[9] See Appendix A.3 for derivation

relative price of capital, along with its standard deviation and implied confidence intervals for σ. Results are shown for three different estimators - robust regression (the estimator used by Karabarbounis and Neiman) and OLS, both on country-specific time trends, and panel fixed effects, on country-year observations - and for two different sources for the relative price of investment (Penn World Tables and the World Bank). In all cases, following Karabarbounis and Neiman, the sample is restricted to contain only countries with 15 or more years of observations, and contains the corporate investment rate where it is available, and its whole-economy analogue where it isn't. The results are clearly sensitive to the choice of estimator and sample. Using Karabarbounis and Neiman's preferred robust regression methodology, the central estimates of σ are 0.2 and 0.3, depending on the source data for relative prices. None of the confidence intervals for σ contain 1; on the face of it, these results do not corroborate the greater-than-unit elasticity presented in Karabarbounis and Neiman (2014), instead leaning toward a less-than-unit elasticity in line with most estimates in the literature (Chirinko (2008)). Nominal investment rates are typically increasing in the relative price of capital, rather than decreasing, as a model with a greater-than-unit elasticity would predict.[10]

8.2 Testing the model's predictions on house prices, household debt and net foreign assets

The baseline closed-economy model predicts that countries with relatively low capital-goods prices will have low steady-state interest rates. In the open-economy version, these lower shadow real rates translate into positive net foreign asset positions. And in a state state with growth in nominal GDP, these more positive external positions would necessitate more positive current account balances. Furthermore, around the time of the transition, the bigger the fall in capital goods prices, the more positive the

[10]The results also constitute a puzzle - why the investment share suggests an elasticity below unity, whereas the capital share suggests an estimate the other side of it. The implication is that markups are negatively related to the relative price of capital, perhaps because the model is missing a third factor.

current account balance. The model therefore predicts a negative relation between capital goods prices and the current account, both in the steady state and during the transition.

The closed-economy model also predicts a negative steady-state relationship between capital goods prices and household debt. There is no such long-run relationship between debt and relative prices - taking world interest rates as given - in the open economy model. Finally, both models predict a negative relationship between relative prices and real house prices.

Given the rising, but on average intermediate, degree of de facto capital-account openness in the world economy over the past forty years, it is not clear a priori whether the open- or closed-economy versions of the model will turn out to be better approximations to the real world.

With these caveats in mind, we can take the model to the data in a manner analogous to the previous subsection, regressing the household debt-GDP ratio, real house prices and the current account-GDP ratio on the level of the relative price of capital. We use the same three estimators used in the previous subsection: panel fixed effects on annual country-year observations; OLS on country-specific time trends in relative capital goods prices and the current account, the household debt-GDP ratio, and real house prices; and robust regression on the same.[11]

Tables 2, 3 and 4 set out the results. Table 2 shows that, for panel fixed effects and robust regression on time trends, we find a large and significant negative relationship between the relative price of capital goods and household debt. This is in line with the prediction of the closed-economy version of the model: lower capital goods prices reduce interest rates, stimulating household borrowing. For OLS on time trends, we find a positive but insignificant effect. Table 3 displays a similar pattern - large negative and significant coefficients when using panel fixed effects; negative coefficients when using robust regression, but which are significant for only one of

[11] At any point in time, the behaviour of the current account in country i will depend on the path of capital goods prices in country i relative to other countries. So in the panel regression of current account, we first condition p on time and country fixed effects

the measures of relative prices; and inconclusive results from OLS. Finally, table 4 shows the results for the current account. Here the results for cross-country trends are inconclusive, but panel fixed-effects deliver a significant negative coefficient, in line with the predictions of the open-economy model.

Overall, cross-country econometric analysis provides qualified support for the assumptions in and predictions of the model. There is strong evidence that nominal investment rates are increasing in the relative price of capital p, and thus that the key elasticity parameter σ is below 1. There is some evidence that household debt and house prices are both negatively related to p, consistent with the predictions of the closed-economy model. And there is weak evidence that the current account is negatively related to p, consistent with the open economy model. But taken together, the economies in our sample appear to have behaved more like financially closed economies on average over the period in question.

9 Conclusion

This paper presents a model of 'secular stagnation' - persistently low real interest rates - driven by the interaction of life-cycle savings motives and an improvement in the technology for producing investment goods. The model is complementary to other explanations for low real interest rates that rely on demographics, emerging markets and inequality. Using standard parameter values and the observed path for capital goods prices over the past few decades, it is able to reproduce part of the rising-falling pattern in real interest rates, the falling ratios of nominal investment and capital to GDP, and the rise in household debt observed across the industrialised world.

The dynamic simulations predict that the real interest rate will stay low, even if the relative price of capital goods has stopped falling. The model suggests that limiting the accumulation of household debt would have made the fall in interest rates larger. And it suggests that the rise in the wealth-income ratio the shock has produced may have increased inherited wealth inequality, even though interest rates have fallen.

References

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008). An equilibrium model of "global imbalances" and low interest rates. *American Economic Review 98*(1), 358–93.

Chirinko, R. S. (2008, June). σ: The long and short of it. *Journal of Macroeconomics 30*(2), 671–686.

Eggertsson, G. and N. Mehrotra (2014, 4). A model of secular stagnation. mimeo.

Giglio, S. and T. Severo (2012). Intangible capital, relative asset shortages and bubbles. *Journal of Monetary Economics 59*(3), 303–317.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-run implications of investment-specific technological change. *The American Economic Review 87*(3), pp. 342–362.

Hansen, A. (1938). *Full Recovery Or Stagnation?* W. W. Norton, Incorporated.

IMF (2014, April). World economic outlook. Technical report, IMF.

Ingves, S. (2014, 11). Interview with dagens nyheter. Interview with Dagens Nyheter.

Karabarbounis, L. and B. Neiman (2014). The global decline of the labor share. *The Quarterly Journal of Economics 129*(1), 61–103.

King, M. and D. Low (2014, February). Measuring the "world" real interest rate. Working Paper 19887, National Bureau of Economic Research.

Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Marx, K. (1867). *Capital: A Critique of Political Economy.* Number v. 1 in Capital: A Critique of Political Economy. Penguin Books Limited.

Piketty, T. (2011). On the long-run evolution of inheritance: France 18202050. *The Quarterly Journal of Economics 126*(3), 1071–1131.

Piketty, T., G. Postel-Vinay, and J.-L. Rosenthal (2014). Inherited vs self-made wealth: Theory & evidence from a rentier society (Paris 18721927). *Explorations in Economic History 51*(C), 21–40.

Rey, H. (2014, 1). Panel on secular stagnation. Remarks by Helene Rey at Panel on Secular Stagnation.

Summers, L. (2013, 11). Crises yesterday and today. Remarks by Larry Summers at 14th Annual IMF Research Conference, Nov 8 2013.

Summers, L. H. (2014, Apr). U.s. economic prospects. *Bus Econ 49*(2), 65–73.

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics 12*(3), 423 – 443.

A Derivations

A.1 Saving

$$
\begin{aligned}
c_1^{-\theta} =& \beta\left(1+r\right)c_2^{-\theta} \\
c_2 =& \left[\beta\left(1+r\right)\right]^{\frac{1}{\theta}} c_1 \\
S = W - c_1 =& \frac{c_2}{\left(1+r\right)} \\
=& \left(1+r\right)^{-1} c_1 \left[\left(\beta\left(1+r\right)\right)\right]^{\frac{1}{\theta}} \\
=& c_1 \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1} \\
\frac{W}{c_1} =& 1 + \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1} \\
\frac{c_1}{W} =& \frac{1}{1 + \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}} \\
\frac{W - c_1}{W} =& 1 - \frac{1}{1 + \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}} \\
=& \frac{\beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}}{1 + \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}} \\
S =& W \frac{\beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}}{1 + \beta^{\frac{1}{\theta}} \left(1+r\right)^{\frac{1}{\theta}-1}}
\end{aligned}
$$

A.2 Investment schedule in $\{s, r\}$ space

From the CES production function we have

$$\begin{aligned} r+\delta =& \frac{1}{p}\frac{\partial Y}{\partial K} \\ =& \frac{1}{p}\alpha\left(\frac{Y}{K}\right)^{\frac{1}{\sigma}} \\ \frac{(r+\delta)\,p}{\alpha} =& \left(\frac{K}{Y}\right)^{-\frac{1}{\sigma}} \\ \frac{K}{Y} =& \left[\frac{\alpha}{p\,(r+\delta)}\right]^{\sigma} \end{aligned}$$

From CRS and Euler's theorem we have

$$\begin{aligned} W+K\frac{\partial Y}{\partial K} =& Y \\ \frac{W}{Y} =& 1-\frac{K}{Y}\frac{\partial Y}{\partial K} \\ =& 1-\alpha\left(\frac{K}{Y}\right)^{\frac{\sigma-1}{\sigma}} \end{aligned}$$

Now we can rewrite the saving rate as

$$
\begin{aligned}
s =& \frac{S}{W} \\
=& \frac{\left(\frac{pK}{Y}\right)}{\left(\frac{W}{Y}\right)} \\
=& \frac{p\left[\frac{\alpha}{p(r+\delta)}\right]^{\sigma}}{1-\alpha\left[\frac{\alpha}{p(r+\delta)}\right]^{\sigma-1}} \\
=& \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{1-\alpha p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma-1}} \\
=& \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}}
\end{aligned}
$$

So the derivative of the saving rate with respect to the price of capital goods p is the same sign as $1-\sigma$.

Which way does the interest rate schedule slope? There are two effects. The effect in the numerator is negative for the standard reasons: for given capital goods prices, more savings reduces the marginal product of capital and hence the interest rate. The effect in the denomnator is of ambigious sign, and comes through the labour share (for a Cobb-Douglas function $\sigma - 1 = 0$ it is absent). For low σ, an increase in r reduces the denominator, raising the quotient. This is because we are expressing the saving rate as a fraction of wages and when $\sigma < 1$, higher interest rates are associated with a lower labour share. To save enough for a given volume of capital goods, a lower labour share must mean a higher saving rate.

For reasonable parameter values, the effect on the numerator will dominate, such that the investment schedule slopes down in $\{s, r\}$ space. To see this, differentiate the schedule with respect to r

$$\begin{aligned}
\frac{ds}{dr} =& \frac{-\sigma p^{1-\sigma}\alpha^{\sigma}\left(r+\delta\right)^{-\sigma-1}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \\
&- \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{\left(1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}\right)^{2}}\left(-\alpha p^{1-\sigma}\alpha^{\sigma-1}\left(1-\sigma\right)\left(r+\delta\right)^{-\sigma}\right) \\
=& \frac{-\sigma}{r+\delta}s + s\frac{\alpha p^{1-\sigma}\alpha^{\sigma-1}\left(1-\sigma\right)\left(r+\delta\right)^{-\sigma}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \\
=& \frac{-\sigma}{r+\delta}s + \left(1-\sigma\right)s^{2} \\
=& s\left(\left(1-\sigma\right)s - \frac{\sigma}{r+\delta}\right)
\end{aligned}$$

A.3 Estimating σ from the nominal investment share in the Karabarbounis and Neiman (2014) model

The Karabarbounis and Neiman (2014) model decomposes income into the capital share s_K, the labour share s_L and markups μ

$$\mu\left(s_K + s_L\right) = 1$$

Taking logs and then the derivative with respect to time gives

$$\begin{aligned}
0 = \frac{d}{dt}\log\left(\mu\left(s_K+s_L\right)\right) =& \hat{\mu} + \frac{1}{s_K+s_L}\left[\frac{ds_K}{dt}+\frac{ds_L}{dt}\right] \\
=& \hat{\mu} + \mu s_K \hat{s_K} + \left(1-\mu s_K\right)\hat{s_L}
\end{aligned}$$

Following Karabarbounis and Neiman (2014), we set $\mu = 1, \hat{\mu} = 0$ and get

$$s_K \hat{s_K} + s_L \hat{s_L} = 0$$

and therefore we can rewrite the left-hand side of their equation (19) as follows

$$\frac{s_L}{1 - s_L}\hat{s_L} = \frac{s_L}{1 - s_L}\frac{-s_K \hat{s_K}}{s_L} = -\hat{s_K}\frac{s_K}{1 - s_L} = -\hat{s_K}$$

Karabarbounis and Neiman (2014) write the capital share as

$$\begin{aligned} s_K =& \frac{RK}{Y} \\ =& \frac{1}{\mu} F_K \frac{K}{Y} \\ =& \frac{\alpha_K A_K^{\frac{\sigma-1}{\sigma}}}{\mu} \frac{Y}{K}^{\frac{1}{\sigma}} \frac{K}{Y} \\ =& \frac{\alpha_K A_K^{\frac{\sigma-1}{\sigma}}}{\mu} \frac{K}{Y}^{\frac{\sigma-1}{\sigma}} \end{aligned}$$

If we assume away changes in technology, capital shares and markups we have

$$\hat{s_K} = \frac{\sigma - 1}{\sigma} \widehat{\left(\frac{K}{Y}\right)}$$

In the steady state, the nominal investment rate is proportional to the nominal capital-output ratio

$$\frac{I_N}{Y} = \frac{p\delta K}{Y}$$

so that

$$\frac{K}{Y} = \frac{I_N}{Y}\frac{1}{p\delta}$$

If we assume away changes in depreciation rates we have

$$\hat{s_K} = \frac{\sigma - 1}{\sigma}\widehat{\frac{K}{Y}} = \frac{\sigma - 1}{\sigma}\left(\widehat{\frac{I_N}{Y}} - \hat{p}\right)$$

Combining these results with the estimating equation (19)

$$\frac{s_{Lj}}{1 - s_{Lj}}\hat{s_{Lj}} = \gamma + (\sigma - 1)\,\hat{p} + u_j$$

we have

$$\begin{aligned}
\frac{s_{Lj}}{1 - s_{Lj}}\hat{s_{Lj}} &= -\,\hat{s_K} = \frac{1 - \sigma}{\sigma}\left(\widehat{\frac{I_N}{Y}} - \hat{p}\right) \\
\widehat{\frac{I_N}{Y}} - \hat{p} &= \tilde{\gamma} + \frac{\sigma}{1 - \sigma}\,(\sigma - 1)\,\hat{p} + \tilde{u_j} \\
\widehat{\frac{I_N}{Y}} &= \tilde{\gamma} + (1 - \sigma)\,\hat{p} + \tilde{u_j}
\end{aligned}$$

In other words, if the elasticity of substitution between capital and labour σ is greater than unity, then a fall in the relative price of capital should lead to a rise in the ratio of nominal investment to GDP, as the volume of investment rises by a greater proportion than the fall in its price.

B Tables and charts

Table 1: Estimates of the elasticity of substitution σ

Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
Log(p)	0.491***	1.121***	0.776***	0.290***	0.999***	0.695***
	[0.04]	[0.21]	[0.17]	[0.04]	[0.25]	[0.16]
$\hat{\sigma}$	0.509	-0.121	0.224	0.71	0.001	0.305
$\hat{\sigma_H}$	0.589	0.299	0.564	0.79	0.501	0.625
$\hat{\sigma_L}$	0.429	-0.541	-0.116	0.63	-0.499	-0.015
N	1632	54	54	1643	52	52
no. of countries	99			100		

Table 2: Regression of household debt on relative price of capital

Left-hand side variable	Household debt/GDP					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-0.993***	0.702	-0.779***	-1.179***	0.571	-0.888***
	[0.05]	[0.65]	[0.25]	[0.07]	[0.72]	[0.30]
N	535	18	18	551	18	18
no. of countries	21			21		

Table 3: Regression of real house prices on relative price of capital

Left-hand side variable	Real house prices					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-1.082***	0.121	-0.672	-0.976***	-0.277	-1.520**
	[0.10]	[0.89]	[0.79]	[0.12]	[0.91]	[0.65]
N	535	18	18	551	18	18
no. of countries	21			21		

Table 4: Regression of current account on relative price of capital

Left-hand side variable	Current account/GDP					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-0.055***	0.006	0.020	-0.025**	0.025	0.028
	[0.01]	[0.05]	[0.05]	[0.01]	[0.05]	[0.05]
N	1004	35	35	992	34	34
no. of countries	50			51		

Figure 1: 10-year real interest rates

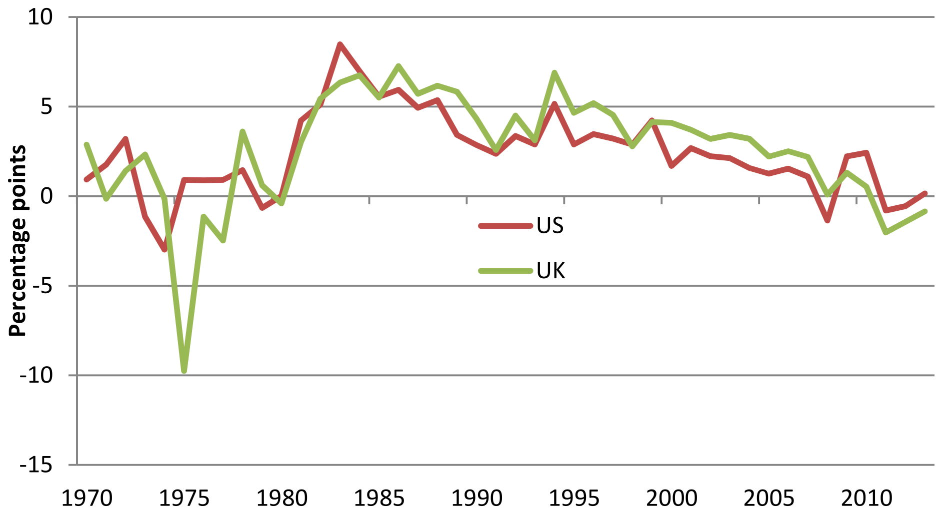

This figure shows estimates of ex-ante 10-year real interest rates for the US and UK, calculated as the difference between nominal government bond yields and model-based estimates of inflation expectations taken from Figure 3.2 of IMF (2014).

Figure 2: HH debt-GDP ratio, % of GDP

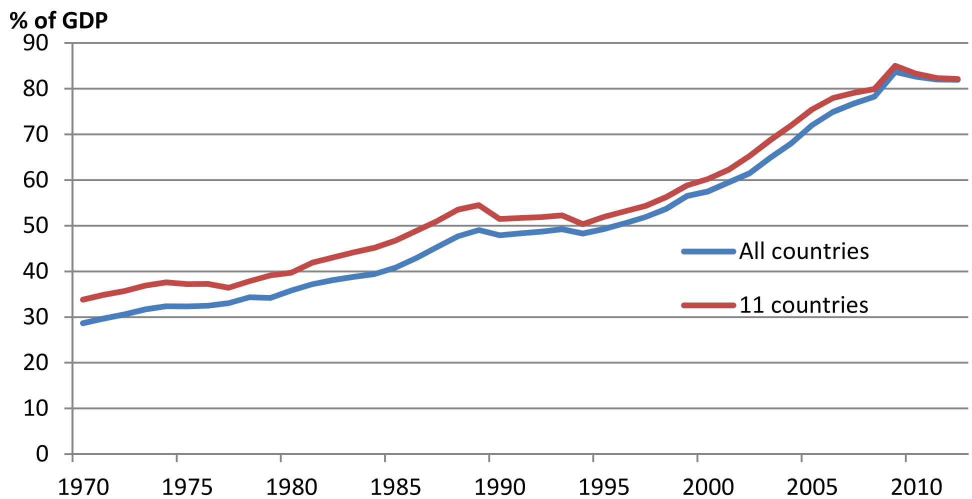

This figure shows the change in the ratio of household debt to GDP since 1970 for a broad sample of industrialised countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US). The source for household debt is the BIS, and the source for GDP is OECD StatBase. The chart is constructed from an unbalanced panel of data by running a fixed-effects panel regression of the household debt ratio on year dummies, then adding the dummy for each year to the intercept of the equation. This allows other countries to affect the change in the ratio in years after they have been added to the sample.

Figure 3: Nominal investment-GDP ratios

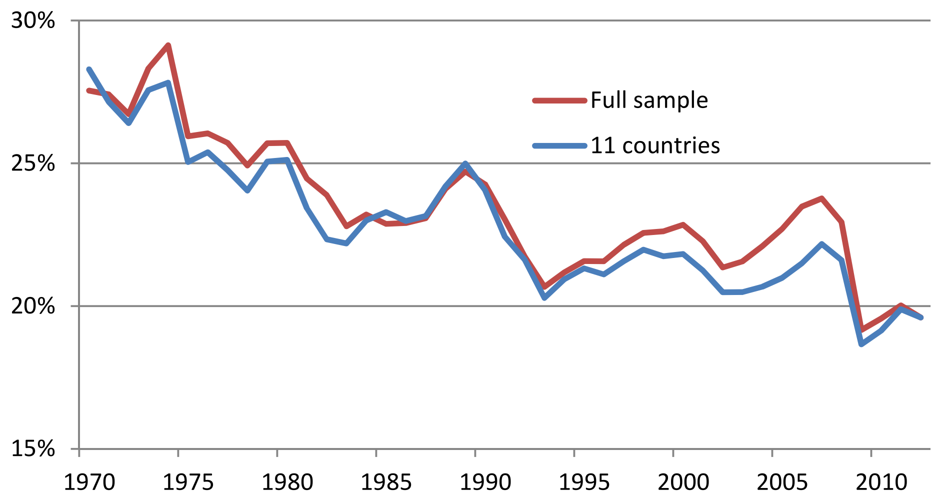

This figure shows a simple average across 24 OECD countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US) of the ratio of nominal gross capital formation to nominal GDP. The source is OECD Statbase.

Figure 4: Nominal and real capital stock-GDP ratio

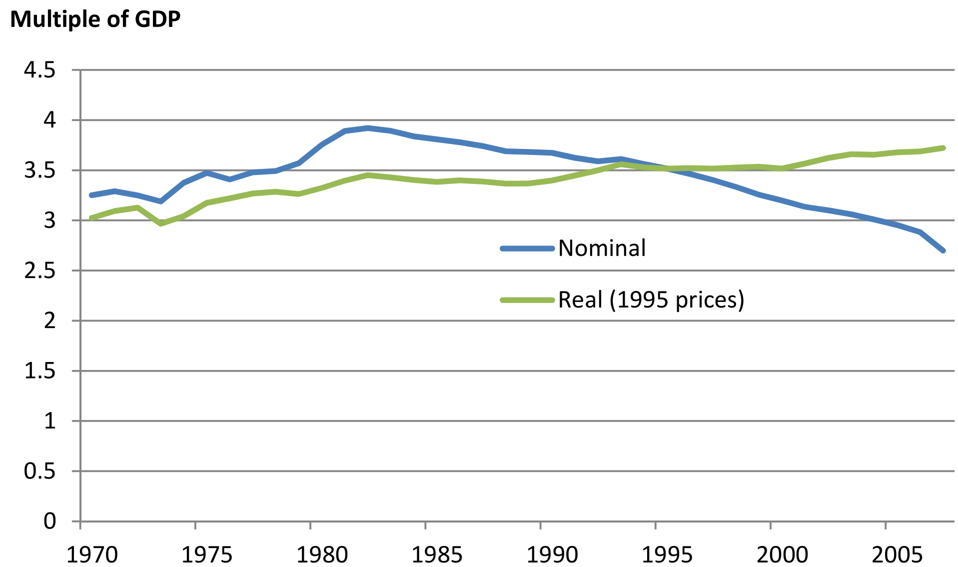

This figure shows the average nominal and real capital-output ratios for Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US. The source is the 'All capital input files' file from the November 2009 release of the EU-KLEMS database. The nominal capital stock was computed for each country as the product of the real capital stock at 1995 prices and the gross fixed capital formation price index rebased to 1995, and then divided by nominal GVA taken from EU-KLEMS to give the nominal capital output ratio. Country-year observations were regressed on country and year dummies, and average index in year t was obtained as the sum of the intercept and the dummy for year t. The real capital-output ratio was constructed analogously, then by rebasing the average ratio of the real capital stock and real GDP to the 2005 nominal ratio.

Figure 5: Real investment-GVA ratio, 11 industrialised countries, 2007=1

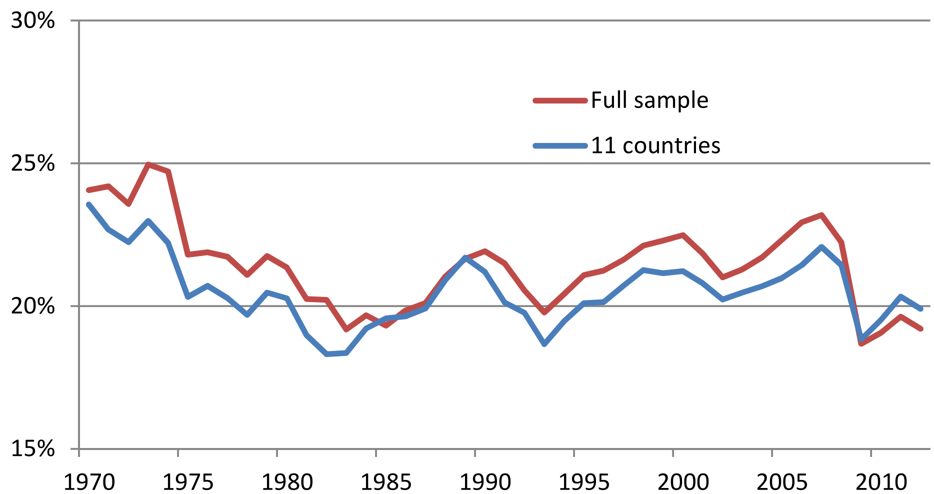

This figure shows a simple average across 24 OECD countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US) of the ratio of gross capital formation to GDP, both at constant prices. The source is OECD Statbase.

Figure 6: Price of investment relative to consumption

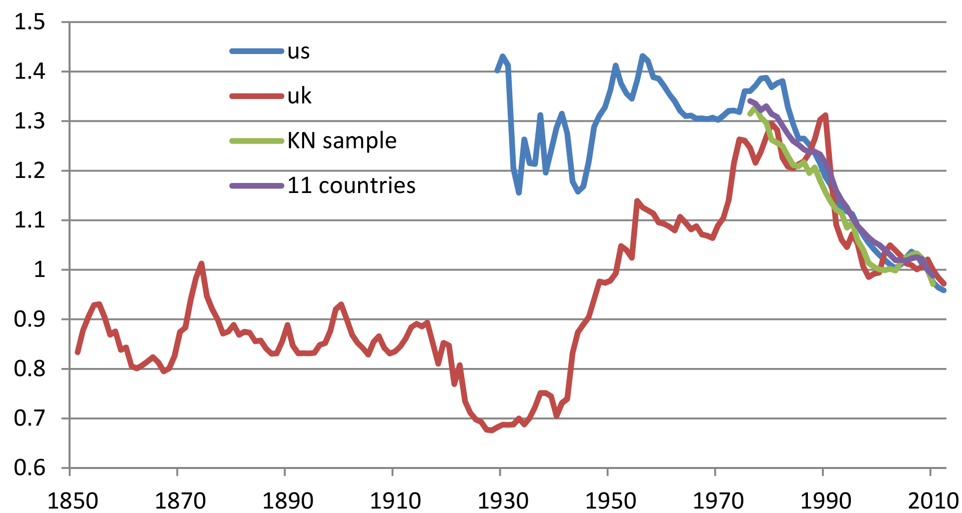

This figure shows four series of the relative price of investment to consumption. The red line is the ratio of the deflators of gross capital formation and private consumption in the UK, taken from the Bank of England's internal long-run database. The blue line is the analogous ratio for the US, taken from the FRED database. The green line is constructed from an unbalanced panel of data by running a fixed-effects panel regression on the data for all countries in Karabarbounis and Neiman (2014), then adding the dummy for each year to the intercept of the equation. The purple line is constructed in the same way for Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US.

Figure 7: Simple savings-investment diagram

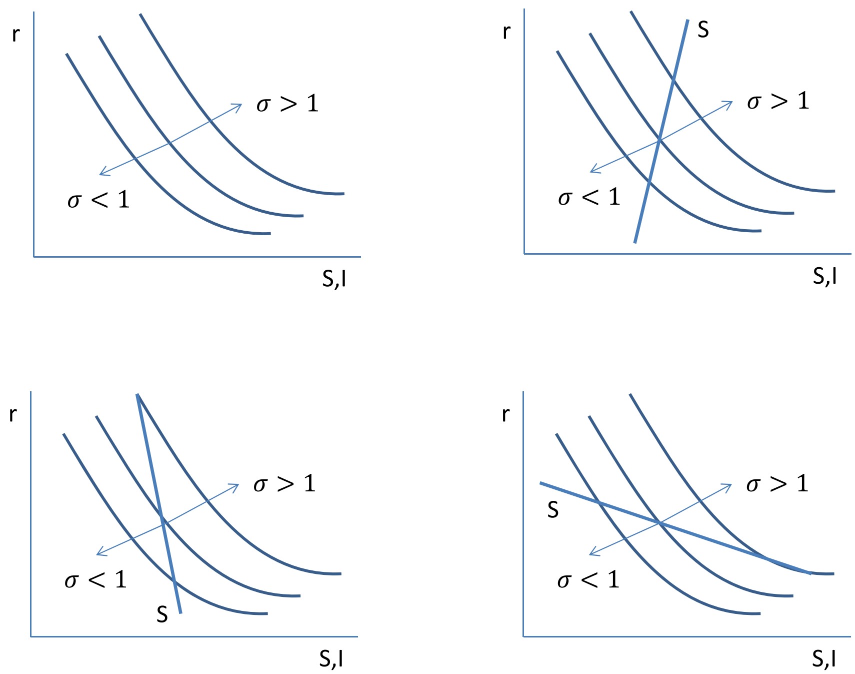

Figure 8: Steady state, baseline setup

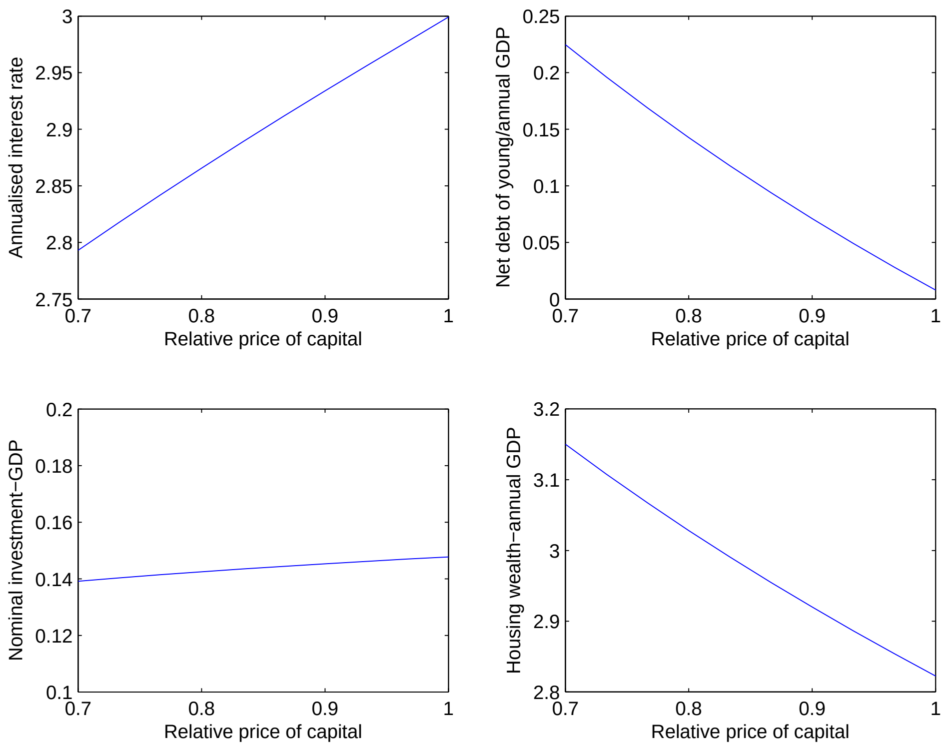

Figure 9: Dynamic solution, baseline setup

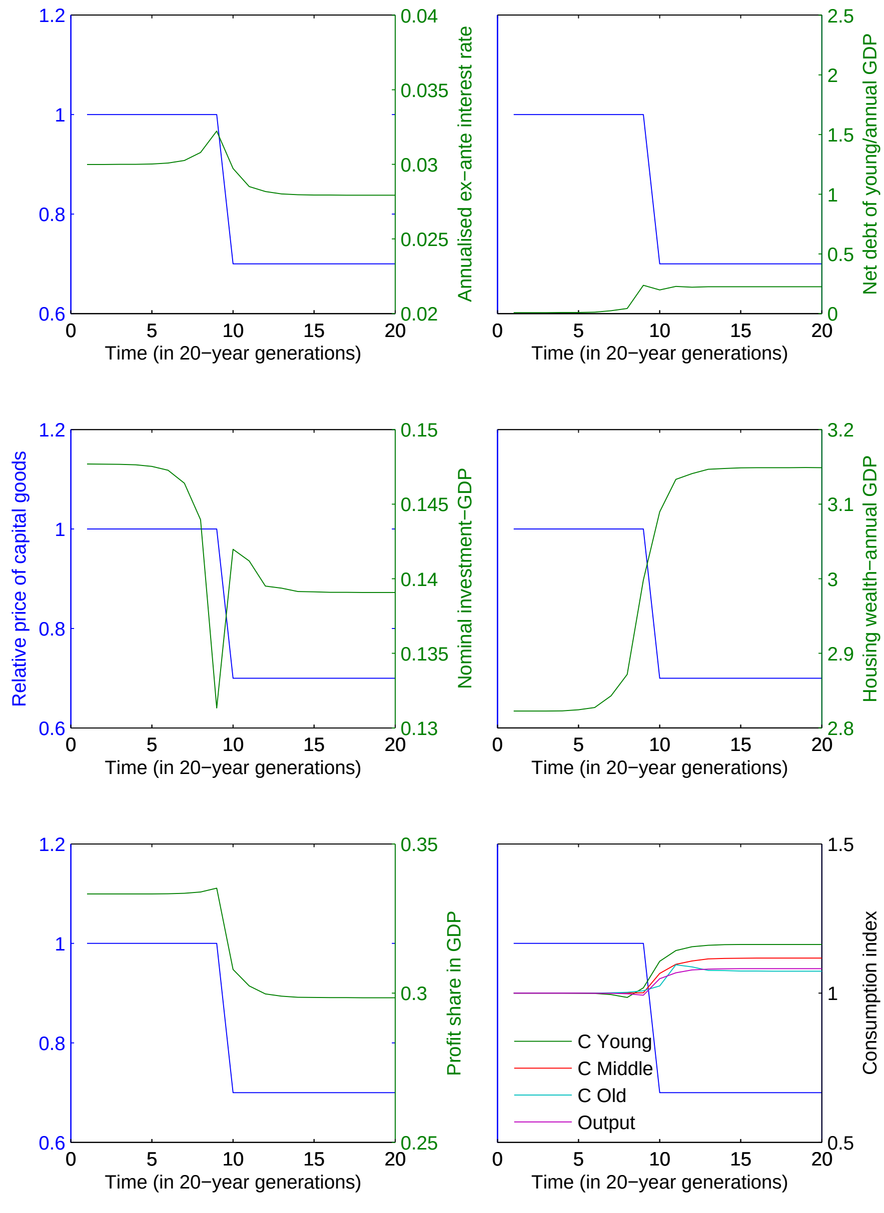

Figure 10: Steady state, varying availability of household debt

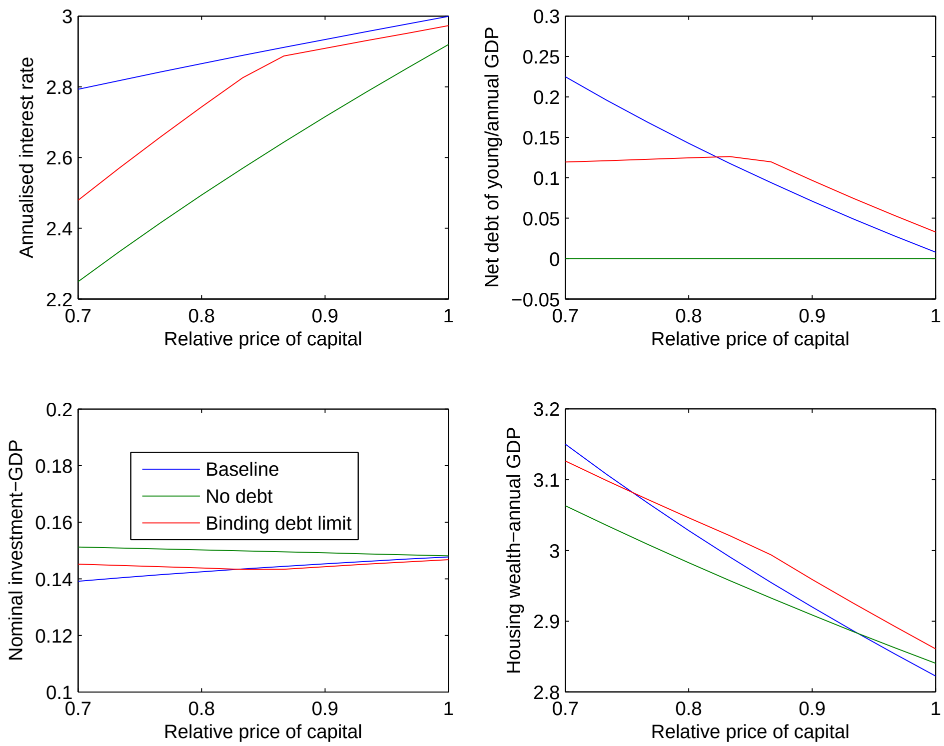

Figure 11: Dynamic solution, no household debt

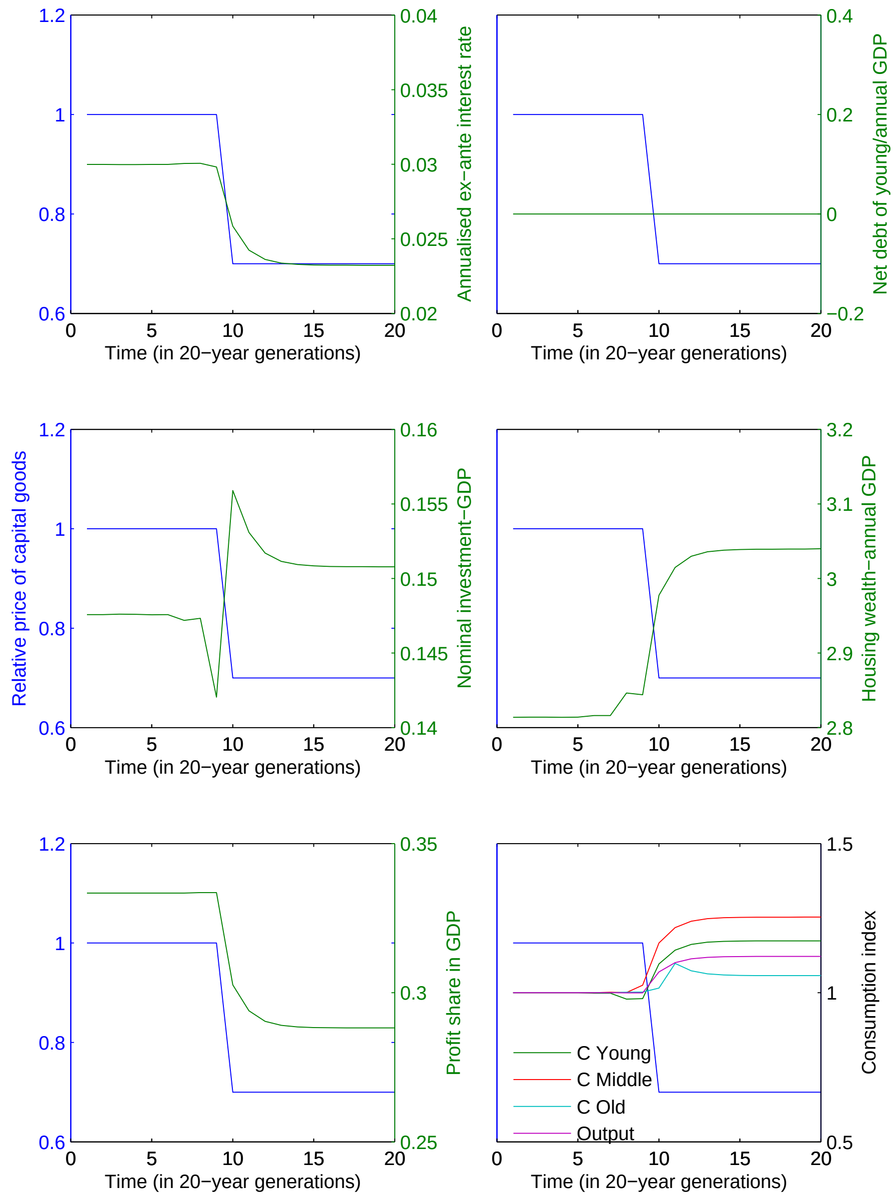

Figure 12: Steady state, varying curvature of the production function

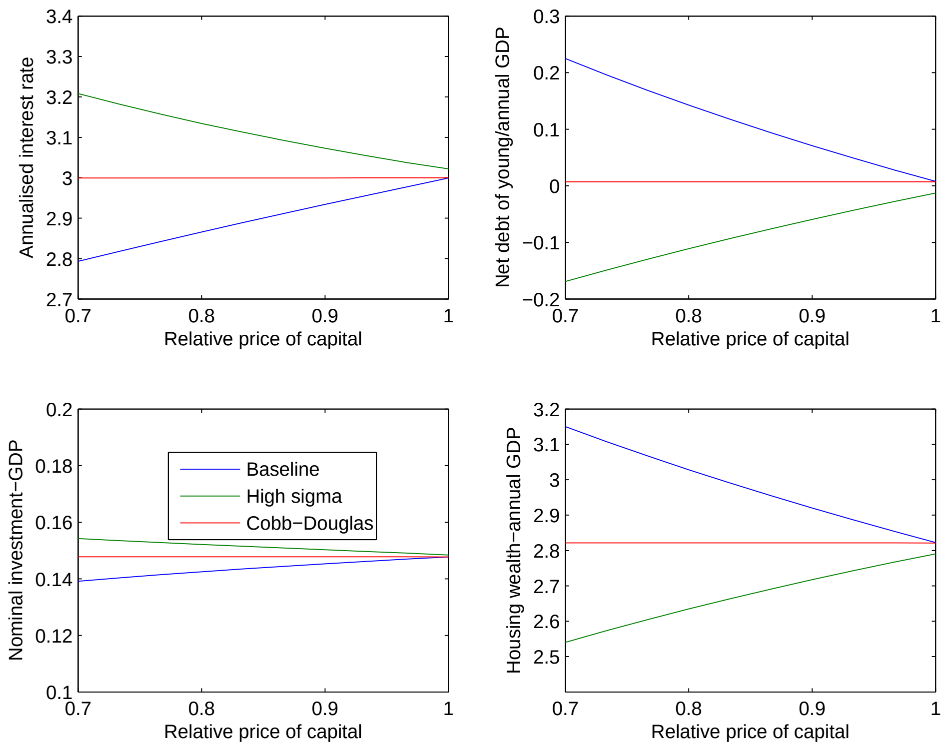

Figure 13: Dynamic solution, Cobb-Douglas production function

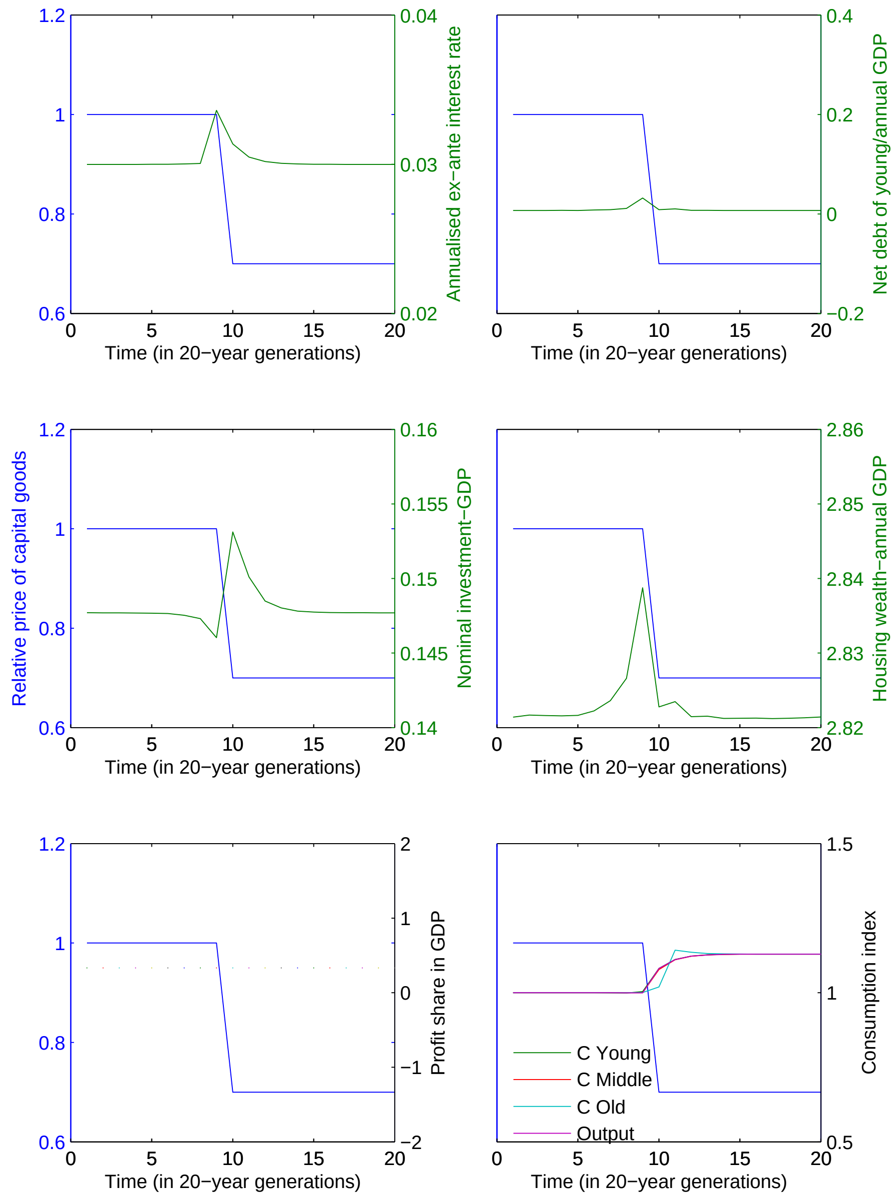

Figure 14: Dynamic solution, highly elastic production function

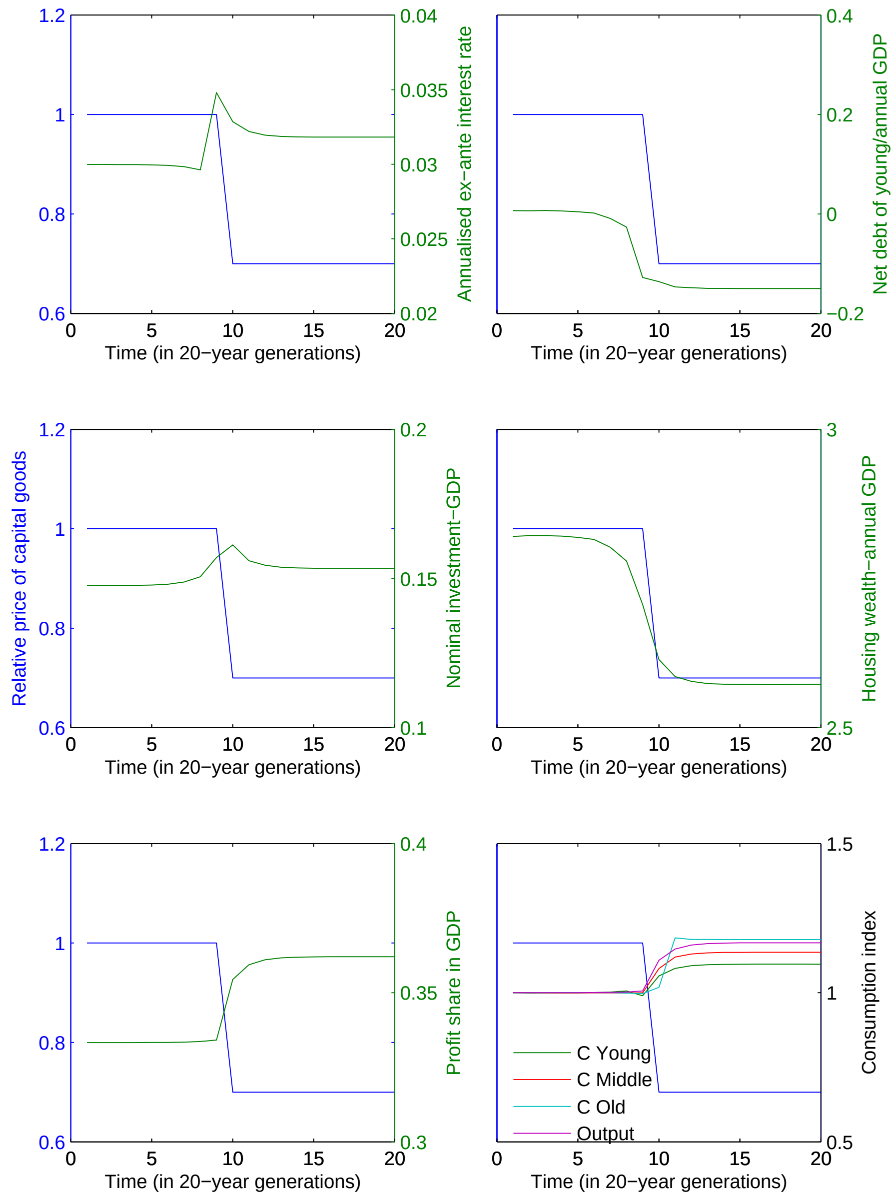

Figure 15: Steady state, varying curvature of the utility function

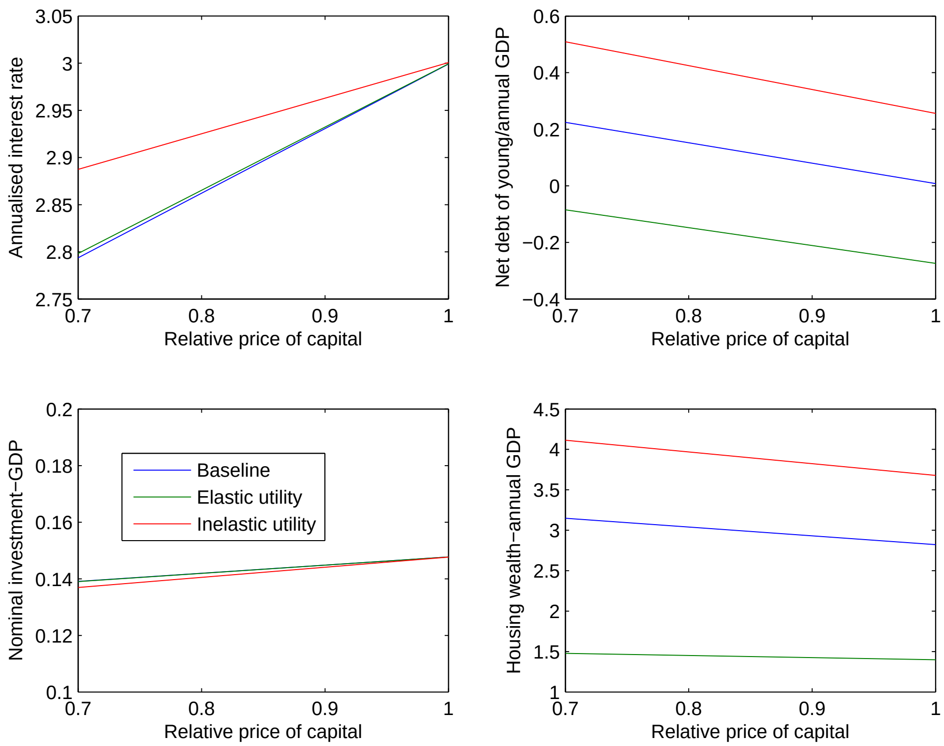

Figure 16: Steady state, no housing, curvature of the utility function

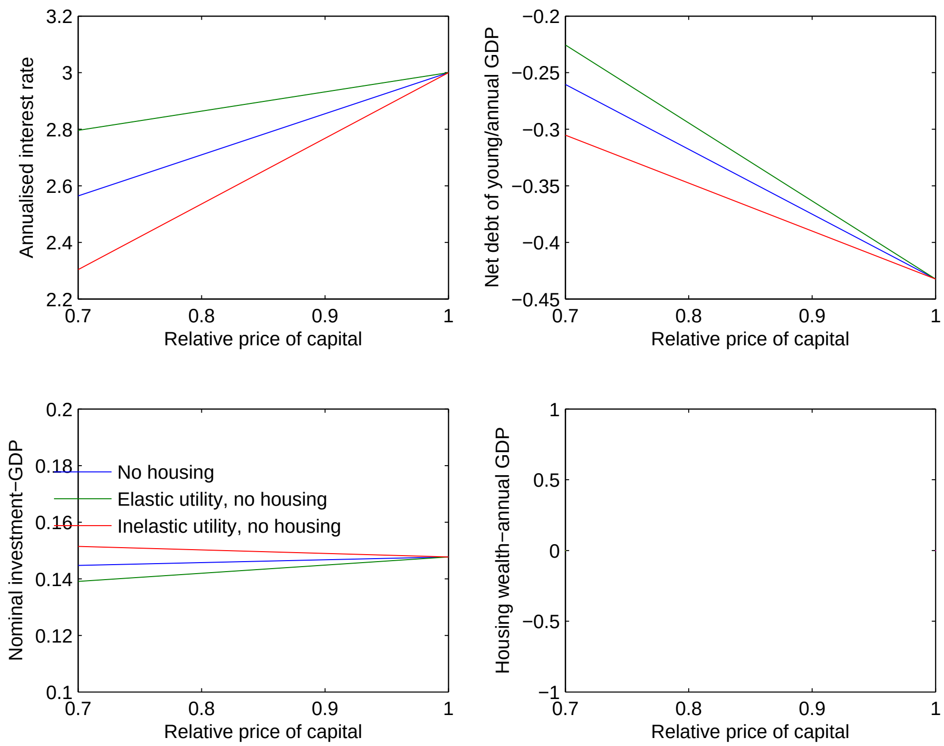

Figure 17: Dynamic solution, alternative timing convention

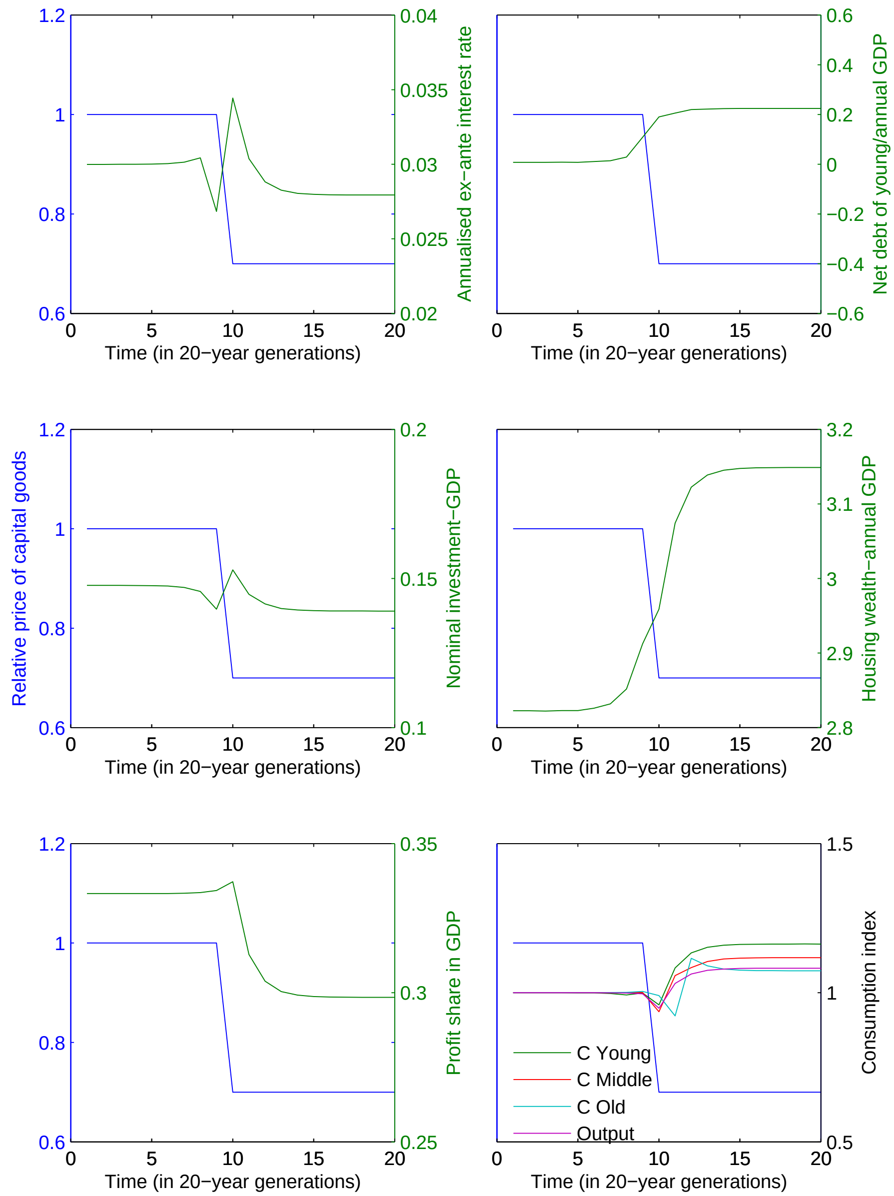

Figure 18: Steady state, bequests and heterogeneous agents

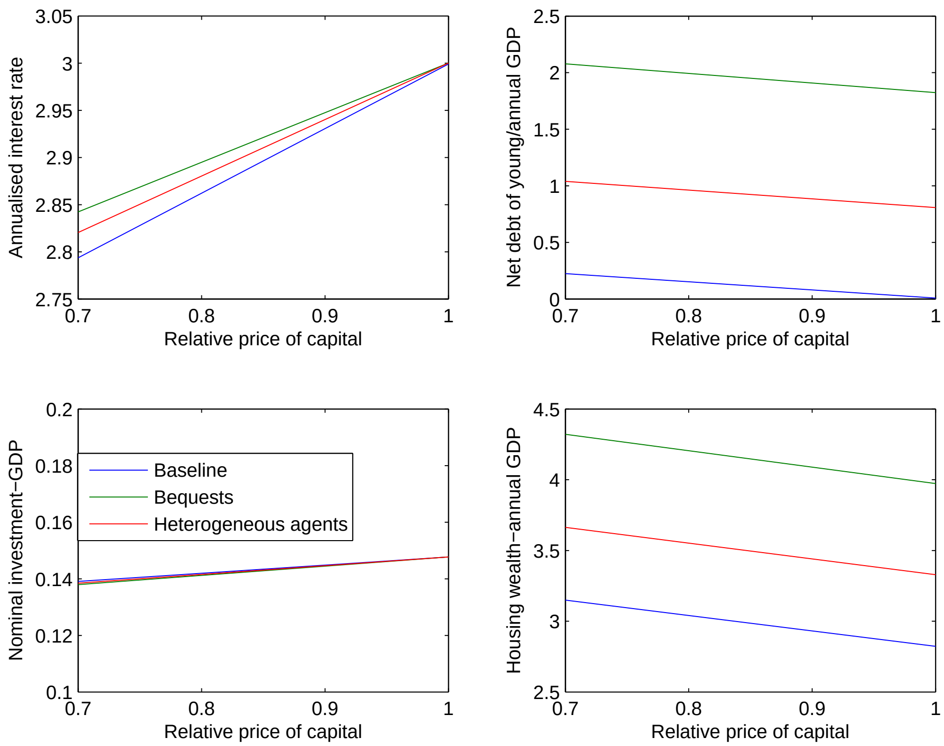

Figure 19: Dynamic solution with bequests

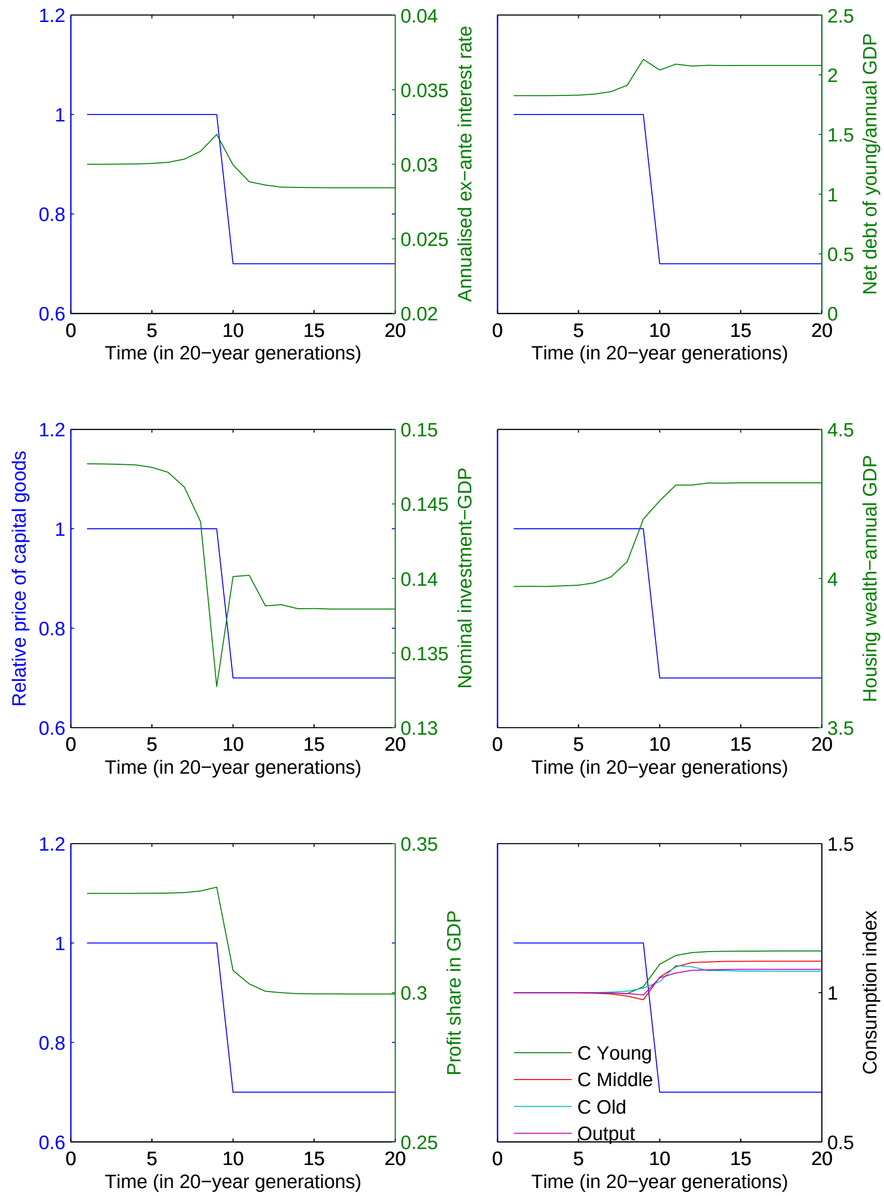

Figure 20: Inequality within generations

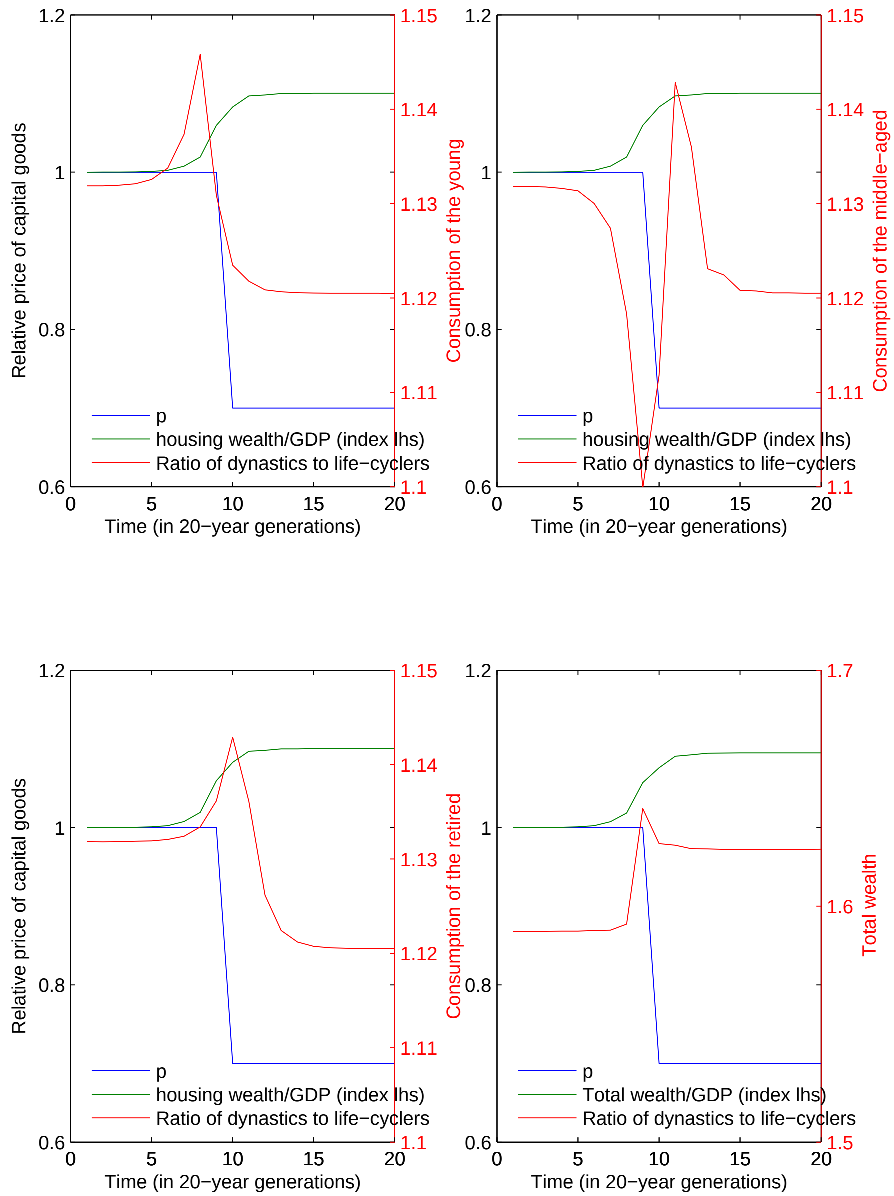

Figure 21: Steady state, small open economy

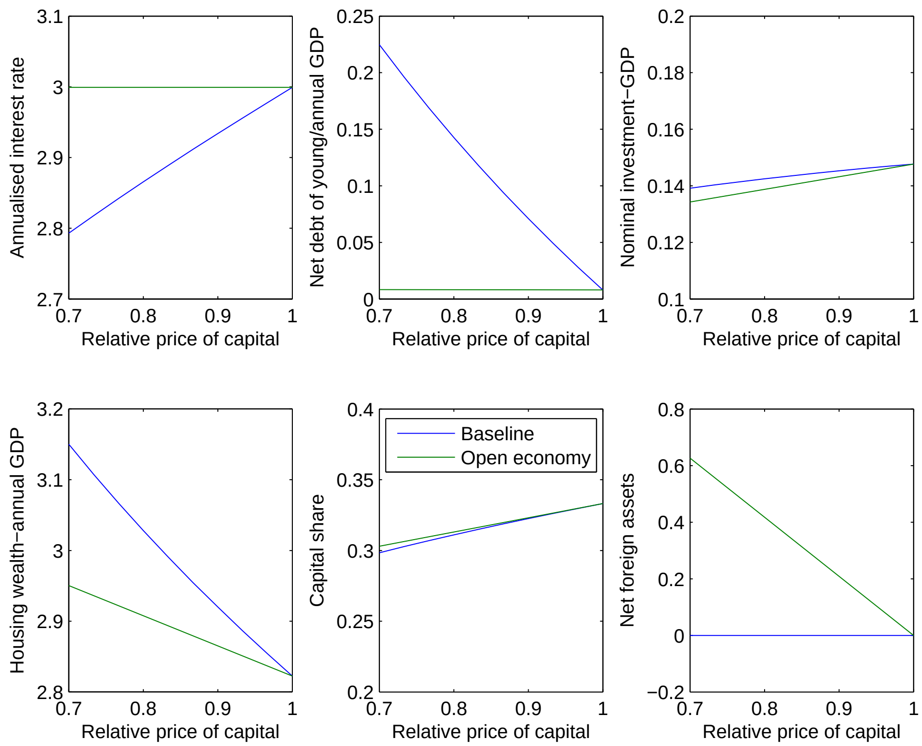

Figure 22: Dynamic solution, small open economy