

[Why are real interest rates so low? Secular stagnation and the relative price of investment goods](#)

CFM-DP2014-28

Gregory Thwaites^{1,2,3}

¹Centre for Macroeconomics; ²London School of Economics; and ³Bank of England

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 these theories all predict 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 movements across countries 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.

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

First version: 23 Jan 2014

This version 7 July 2015

Abstract

Across the industrialised world, real interest rates and nominal investment rates have fallen, while house prices and household debt ratios have risen. I present an OLG model which explains these four trends with a fifth - the widespread fall in the relative price of investment goods. When capital goods are cheaper, a given quantity of saving buys more of them, but the resulting capital deepening lowers the marginal product of each one. Interest rates fall, reducing the user cost of housing, raise house prices and household debt. Preventing the accumulation of debt leads to a bigger fall in interest rates.

JEL codes: E22, E43, E60, E13

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 in relation to incomes before the crisis. While these ratios 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 the supply of domestic or foreign saving to the industrialised world as the cause of the fall in interest rates. They all predict, therefore, a rise in investment in the industrialised world. 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

based on the widespread fall in the price of investment goods relative to consumption over the past thirty or so years documented in Karabarounis and Neiman (2014). I extend their data back further back in time for some countries, to corroborate their finding that this fall has not been a feature of the very long run, but rather began in the late 1970s.

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 in the literature (see e.g. Acemoglu and Robinson (2014), Chirinko (2008)). Consistent with the predictions of the model at these parameter values, I present cross-country evidence showing a positive association across countries between nominal investment rates and the relative price of capital goods. An extension of the model with three factors of production in the spirit of Autor et al. (2003) can reproduce both a falling labour share (as wage labour is highly substitutable with capital) and a falling investment rate (as the product of the two are complementary to the third factor, such as entrepreneurial labour or intangible capital).

The new steady state after investment goods prices have stopped falling is one of lower interest and investment rates and higher household debt ratios. Lower interest rates reduce the user cost of housing, boosting housing demand. Housing supply is fixed, so house prices must rise.¹ 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

¹In the model, housing is effectively residential land, in that its supply is fixed. Knoll et al. (2014) find that the bulk of the increase in house prices is attributable to the increase in the supply of residential land.

see in the data. The model's implications for household debt and house prices are consistent with global trends over the period in question and receive qualified support in cross-country econometric analysis. But they are not central to the mechanism that delivers low interest rates - removing housing from the model makes the fall in interests even more stark.

The dynamics of the transition to the new steady state of low interest rates 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 can operate for some decades both before and after the change in relative prices.

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. In parallel to this work, Rey (2014), Summers (2014) and Eichengreen (2015), among others, have 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 a move away from tangible capital in production can make bubbles more likely when there are financial frictions and intangible capital makes for poor collateral. Like the present study, Eggertsson and Mehrotra (2014) address the issue of secular stagnation in an OLG model. They show that an exogenous 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. It does not rely on a binding exogenous debt limit for young households, but instead generates endogenously the implications of lower capital goods prices for house prices, household debt and wealth inequality. But unlike Eggertsson and Mehrotra (2014), the neoclassical model in 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. Section 7 extends the model to allow for 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: falling real interest rates (subsection 2.1); rising household debt (subsection 2.2); rising house prices (subsection 2.3); and falling capital goods prices and nominal investment rates (subsection 2.4). 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 for which the EU-KLEMS database has sufficient data to calculate long time-series of nominal and real capital-GDP ratios.²

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 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.

It is possible that government bond yields have fallen because of a rise in the convenience yield of holding government liabilities, or relatedly because a ‘shortage’ of safe assets has introduced a wedge between government and private yields Caballero and

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

Farhi (2013). Figure 2 shows that spreads between corporate and government bond yields have indeed risen. But the magnitude of the rise - less than 100 basis points since 1980 - is too small to offset the more than 200 basis points fall in government bond yields. Figure 2 also shows that the dividend and earnings yields on the broad US stock market index - alternative measures of the cost of funds for firms - have both fallen over the period in question. So broader measures of the cost of funds for the corporate sector in the US, which do not include any convenience yield or safety premium on government assets, corroborate the point that real returns have fallen since the 1980s among major economies.

2.2 Rising household debt ratios

Figure 3 shows an index of the ratio of household debt to GDP since 1970 for a broad sample of industrialised countries and the restricted sample of 11 countries. The figure shows a rise in the average ratio of around 50 percentage points since 1970.

2.3 Rising house prices

Knoll et al. (2014) study house prices in 14 advanced economies since 1870 and find that real house prices have tripled since the early 20th century, with ‘virtually all the increase occurring in the second half of the 20th century’. They decompose this into changes in the price of residential structures and residential land and find that around 80 per cent of the rise in house prices since 1950 is attributable to land prices.

2.4 Price and quantity of capital investment and capital stock

Figure 4 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 6 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 6 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 IMF (2014) and Karabarbounis and Neiman (2014). Figure 7 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.

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

$$r = \frac{1}{p} \frac{\partial Y}{\partial K} - \delta \quad (1)$$

is higher.

But other things will not be equal - lower 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.

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

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

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

$$s = \frac{W - c_1}{W} = \frac{\beta^{\frac{1}{\theta}} (1 + r)^{\frac{1}{\theta} - 1}}{1 + \beta^{\frac{1}{\theta}} (1 + r)^{\frac{1}{\theta} - 1}} \quad (2)$$

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 is vertical.

Turning to the determination of factor prices, firms hire labour and borrow funds from young households, buy capital goods (which depreciate at rate δ) at relative price p and maximise profits with them.⁴ Factor prices $\{W, r\}$ will therefore be set equal to marginal product in the standard fashion

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

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

⁴We 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

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 $\{r, p\}$ ⁵

$$s = \frac{p^{1-\sigma} \left[\frac{\alpha}{(r+\delta)} \right]^\sigma}{1 - \alpha p^{1-\sigma} \left[\frac{(r+\delta)}{\alpha} \right]^{1-\sigma}} \quad (3)$$

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 denominator is of ambiguous 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 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 $1 - \sigma$. The intuition is familiar. 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 $1 - \sigma$.

⁵Derived in Appendix A.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’⁶

$$U(c_1, c'_2, c''_3, h) = \frac{1}{1-\theta} \left(c_1^{1-\theta} + \beta_2 c'^{1-\theta}_2 + \beta_3 c''^{1-\theta}_3 \right) + \phi \frac{h^{1-\gamma}}{1-\gamma} \quad (4)$$

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. It can borrow or save a net amount a_1 at rate r . In period 2 (middle age), the household remains in the same house, supplies $(1-\eta)$ units of labour, and can again borrow or save a'_2 . In period 3 (retirement), it sells its house and consumes the proceeds plus its accumulated savings. So each of the three periods is associated with a budget constraint as follows

$$c_1 + hp_h + a_1 = \eta W \quad (5)$$

$$c'_2 + S'_2 = (1-\eta)W + (1+r)a_1 \quad (6)$$

$$c''_3 = (1+r'')a'_2 + hp_h \quad (7)$$

⁶Housing is in fixed supply so might be more usefully thought of as land, or more generally any non-produced asset that yields utility. The model is robust to including ‘warm glow’ preferences over bequests but results are not shown here.

Forming and solving the Lagrangean yields standard consumption Euler equations thus

$$\frac{c_1^{-\theta}}{(1+r')(1+r'')} = \frac{\beta_2 c_2'^{-\theta}}{(1+r'')} = \beta_3 c_3''^{-\theta} \quad (8)$$

This also yields a housing demand equation that depends on future house prices and consumption

$$\phi h^{-\gamma} + \beta_3 c_3''^{-\theta} p_h'' = c_1^{-\theta} p_h \quad (9)$$

This is intuitive. The left-hand side is the marginal utility of housing plus the discounted marginal utility of the retirement consumption paid for by the sale of the house. The right-hand side 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

$$Y = A[(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \quad (10)$$

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

$$W = \frac{\partial Y}{\partial L} \quad (11)$$

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

consumption goods

$$1 + r' = \frac{1}{p} \frac{\partial Y'}{\partial K'} + \frac{p'}{p} (1 - \delta) \quad (12)$$

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

$$a_1 + a_2 = K' p \quad (13)$$

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

$$h = \bar{H} \quad (14)$$

In practice, to varying degrees across space and over time, housing supply responds to price incentives, and in this sense the assumption of fixed housing supply is not realistic. Housing in this model can instead be reinterpreted as residential land, the price of which forms a large share of the total cost of a house, and accounts for the majority of the increase in house prices over recent years (Knoll et al. (2014)). Consistent with interpreting the housing variable in the model as land, the residential structures component of housing investment forms part of total investment and GDP in the evidence presented in sections 2 and 8.

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. I set ϕ to hit a value of housing wealth to GDP of around three, in line with the data. The elasticity in the utility function θ is set to unity (log utility).

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 a given percentage changes in p , I change A by the product of this change and the nominal investment rate so as to keep GDP unchanged given existing factor endowments

$$\begin{aligned}
 Y &= C + p_k I \\
 \frac{\partial A}{\partial p} \frac{p}{a} \Big|_{\frac{\partial Y}{\partial p} = 0} &= \frac{p}{a} \frac{I}{p \frac{\partial I}{\partial A}} \\
 &= \frac{I_N}{Y}
 \end{aligned}$$

The elasticity of substitution between capital and labour σ is a crucial parameter and deserves special attention. Chirinko (2008) discusses a number of approaches for estimating this parameter and the resulting range of estimates. These are typically based either explicitly on a firm’s optimisation problem, choosing capital subject to adjustment costs, or on estimating a relation between investment or capital intensity and the level or change in the user cost of capital. Chirinko (2008) cites over thirty estimates, typically based on firm-level panel or aggregate time-series data. The median of these estimates is 0.6 and the mean 0.5, with about 15% of the estimates above the critical value of unity. The author quotes a preferred range of $\sigma = 0.4-0.6$, i.e. much further below the critical value of unity than the baseline assumption used in this paper. Acemoglu and Robinson (2014) survey this and other estimates and reach the same conclusion. Section 8 adds to this evidence of a below-unit elasticity using cross-country panel data on nominal investment rates and prices.

A noteworthy recent addition to this set of estimates is provided by Karabarbounis and Neiman (2014), who compile a large cross-country panel dataset on the relative price of capital goods and the labour share in the corporate sector. They carefully document a fall in the labour share in the corporate sector across most countries in their database. They write down a standard closed-economy model in which the production side of the economy is very similar to that in the present study, but in which consumers have infinite horizons and therefore the interest rate is pinned down by the discount factor. They map country-specific time trends in the data into transitions between steady states in their model, and find using a ‘robust regression’ algorithm that the labour share has fallen faster on average in countries in which the relative price of capital has fallen faster.⁷ In their model, this points to a value of σ exceeding unity; the authors central estimate significantly greater than one.

A number of authors have raised questions about the conclusion that the elasticity of substitution is greater than unity. Rognlie (2015) notes that the value of the capital stock has not increased enough to drive a fall in the labour share as in their model (I expand on this point in Section 8 below, finding evidence for a less than unit value of σ using the investment rate rather than the capital stock). Lawrence (2015) finds that a combination of relatively rapid labour-augmenting technical progress omitted from combined with inelastic production technology is responsible for the fall in the labour share.

In the body of this paper the elasticity of substitution between capital and labour σ is set to 0.7. This is a conservatively high assumption, above the central tendency in the literature. But the true value is ultimately an empirical question. As demonstrated below in the sensitivity analysis of the full model and above in the exposition of simple one, this parameter is crucial for the behaviour of the model.

⁷The authors use the *rreg* command in STATA as a means of downweighting outliers

5.2 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.

5.3 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.

In the simulation above, the change in the relative price of capital is perfectly foreseen. This is a strong assumption. At the other end of the spectrum is possibility that the shock was completely unforeseen. Chart 10 shows the effect of the same shock where agents at each point expect the relative price of capital goods will remain at its present value indefinitely, but otherwise compute the dynamics of the system correctly. By construction, all variables remain at their pre-shock steady-state values until the shock hits, at which point interest rates decline monotonically towards their new, lower steady-state values

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.

5.4 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

$$\frac{\Pi}{Y} = \frac{\Pi}{Kp_K} \frac{Kp_K}{Y} \quad (15)$$

The baseline model generates falls in the real interest rate and the nominal capital-output rate - matching the observations in figures 1 and 6 - and therefore generates a fall in the capital share and a *rise* in the labour share.

There are a number of possible reasons for the apparent discrepancy between the trends in factor shares, the real interest rate and the capital-output ratio in the data. First, the capital-output ratio could be mismeasured, inter alia because of the omission of intangible or nonreproduced factors of production, and has in fact not fallen over time. For example, the value of land has increased over time (Knoll et al. (2014)), and may be attracting an increased share of national income in profits. 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. Second, corporate profits include a component of ‘pure profit’ as well as remuneration for capital investment, corresponding, for example, to producer markups over marginal cost. (Rognlie (2015)) imputes these pure profits in US data and finds them to have risen over time. Lastly, there could be an increasing wedge between government bond yields and the average gross return on capital in the corporate sector, attributable to depreciation, taxes or a safety/convenience yield.

The baseline version of the model incorporates none of these features. Equation 15

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. Appendix C outlines a variant of the baseline model with three factors of production which can generate both a rising profit share and a falling investment rate. In this model, similar in spirit to Autor et al. (2003), capital and wage labour are substitutes in an intermediate production function, the output of which is complementary to a nonreproducible factor (such as entrepreneurial labour or land) which earns profits. A rise in the supply of capital goods reduces both interest rates and the labour share, by increasing the share of profits going to the non-reproducible factor.

6 Sensitivity analysis

6.1 Housing and debt

The availability of debt and housing as alternative savings vehicles is not essential to the story, and actually attenuates the fall in interest rates in the model. This is illustrated in figure 11, which considers the case in which household debt is prohibited (the net savings of the young must be nonnegative, i.e. $a_1 > 0$). 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 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 - 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. 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

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 . 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 12 also 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), Acemoglu and Robinson (2014) and with the econometric evidence presented in section 8, but not with the findings in Karabarbounis and Neiman (2014).

7 Small open economy

This section of the paper extends the model to study the behaviour of asset prices and the external balance in a small 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 (13)

$$a_{1i} + a_{2i} - K'_i p_i = NFA_i \tag{16}$$

where NFA_i denotes the net foreign assets of country i .

The red lines in figure 13 show 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 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 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. Middle-aged households' savings go overseas rather than to young households when the economy is open.

8 Econometric evidence

This section of the paper confronts the predictions of the model with econometric evidence. Subsection 8.1 details evidence on the elasticity of substitution between capital and labour σ . Subsection 8.2 looks at the model's predictions for house prices, debt and the external balance.

8.1 Evidence on the elasticity of substitution between capital and labour

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 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 ⁸

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

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 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

⁸See Appendix B for derivation

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 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 (2014)) 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 (2014), 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 the 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. So on the face of it, these results do not corroborate the greater-than-unit elasticity, 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, suggesting as a model with a less-than-unit elasticity of substitution.⁹

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

⁹Appendix C sets out a model to recognise the seemingly conflicting results obtained from estimating σ from the labour share and the nominal investment rate.

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 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.¹⁰

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

¹⁰At 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

- large negative and significant coefficients when using panel fixed effects; negative coefficients when using robust regression, but which are significant for only one of 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.

The model has important normative and positive implications. First, it and the accompanying econometric results provide additional evidence on a below-unit elasticity of substitution between capital and labour. Secondly, it stresses the point that capital goods are simultaneously a social savings technology, the means of production, and a produced asset in themselves. So changes in the way that capital goods are produced will have implications for other stores of value, such as housing, land, public debt, and any ‘bubbly’ asset Giglio and Severo (2012). And thirdly, it suggests that real interest rates may remain low, or have further to fall, even if the relative price of capital goods has stopped falling. Fiscal and monetary rules that are calibrated implicitly on the real interest rate may need to be revised as a result.

References

- Acemoglu, D. and J. Robinson (2014). The rise and fall of general laws of capitalism. Conference draft for the spring 2015 brookings panel on economic activity, MIT mimeo.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118(4), 1279–1333.
- Caballero, R. J. and E. Farhi (2013). A model of the safe asset mechanism (sam): Safety traps and economic policy. Technical report, National Bureau of Economic Research.
- 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.
- Eichengreen, B. (2015, January). Secular stagnation: The long view. Working Paper 20836, National Bureau of Economic Research.
- 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.
- Knoll, K., M. Schularick, and T. Steger (2014). No price like home: Global house prices, 1870-2012. CEPR Discussion Papers 10166, C.E.P.R. Discussion Papers.
- Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *Review of Economics and Statistics* 85(4), 1063–1070.
- Lawrence, R. Z. (2015, June). Recent declines in labor's share in us income: A preliminary neoclassical account. Working Paper 21296, National Bureau of Economic Research.
- Marx, K. (1867). *Capital: A Critique of Political Economy*. Number v. 1 in *Capital: A Critique of Political Economy*. Penguin Books Limited.

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

Rognlie, M. (2015). Deciphering the fall and rise in the net capital share. Conference draft for the spring 2015 brookings panel on economic activity, Brookings Institution.

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.

A Derivations

A.1 Investment schedule in $\{s, r\}$ space

Here we derive the investment schedule for the simple two-period model in Section 3. From the CES production function we have

$$r + \delta = \frac{1}{p} \frac{\partial Y}{\partial K}$$
$$\frac{K}{Y} = \left[\frac{\alpha}{p(r + \delta)} \right]^\sigma$$

From CRS and Euler's theorem we have

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

Now we can rewrite the saving rate as

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

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

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

$$\frac{ds}{dr} = s \left((1 - \sigma) s - \frac{\sigma}{r + \delta} \right)$$

B Estimating σ from the nominal investment share in the Karabarounis and Neiman (2014) model

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

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

Taking logs and then the derivative with respect to time gives

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

Following Karabarounis 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

$$s_K = \frac{RK}{Y} = \frac{\alpha_K A_K^{\frac{\sigma-1}{\sigma}} K^{\frac{\sigma-1}{\sigma}}}{\mu} \frac{K^{\frac{\sigma-1}{\sigma}}}{Y}$$

For constant technology, capital shares and markups we have

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

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

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

For constant depreciation rates we have

$$\hat{s}_K = \frac{\sigma-1}{\sigma} \frac{\widehat{K}}{Y} = \frac{\sigma-1}{\sigma} \left(\frac{\widehat{I}_N}{Y} - \hat{p} \right)$$

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

$$\frac{s_{Lj}}{1 - s_{Lj}} \hat{s}_{Lj} = -\hat{s}_K = \frac{1 - \sigma}{\sigma} \left(\frac{\widehat{I}_N}{Y} - \hat{p} \right)$$

$$\frac{\widehat{I}_N}{Y} = \tilde{\gamma} + (1 - \sigma) \hat{p} + \tilde{u}_j$$

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.

C Three-factor model

This appendix sketches out an alternative production function for intermediate goods which breaks the negative relation between the labour share and the investment rate and can, at certain parameter values, deliver a labour share and a nominal investment rate which are both increasing in the relative price of capital. It then illustrates the region of the parameter space at which this is so.

C.1 Modified production function

Consider a production function defined over three factors - capital K , unskilled labour L and entrepreneurial labour M , which is paid in profits. K and L are aggregated first into an intermediate X , then combined with M as follows

$$X = [(1 - \alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}}$$

$$Y = [\mu M^{\frac{\theta-1}{\theta}} + (1 - \mu)X^{\frac{\theta-1}{\theta}}]^{\frac{\theta}{\theta-1}}$$

$$Y = \left[\mu M^{\frac{\theta-1}{\theta}} + (1 - \mu) \left[[(1 - \alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \right]^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}$$

C.2 Labour share and investment rate

The nominal investment rate is proportional to the capital share at fixed interest rates, because the marginal product of capital is inversely proportional to the price of capital. So a shift inward in the investment schedule is equivalent to a fall in the capital share

$$\frac{I_N}{Y} = \frac{\delta PK}{Y} = \frac{\delta}{r + \delta} MPK \frac{K}{Y}$$

The factor share for factor Z is

$$\Omega_Z = \frac{\partial Y}{\partial Z} \frac{Z}{Y}$$

Denoting logs with lower case we have

$$\omega_K = A + \left(\frac{1 - \theta}{\theta}\right) y + \left(\frac{\sigma - 1}{\sigma}\right) k + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right) x$$

and (because $L = 1$)

$$\omega_L = A + \left(\frac{1 - \theta}{\theta}\right) y + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right) x$$

C.3 Parameter space

To calculate the derivative of each factor share with respect to the log relative price of capital p note that

$$\frac{d\omega_Z}{dp} = \frac{d\omega_Z}{dk} \frac{dk}{d(mp_k)} \frac{d(mp_k)}{dp}$$

For a fixed interest rate, the user cost condition implies that $\frac{d(mp_k)}{dp} = 1$. The middle term in the chain is negative and depends on the curvature of the production function. So for both the labour share and capital share (and hence investment rate) to be increasing in the relative price of capital, as suggested by the econometric evidence

in Karabarbounis and Neiman (2014) and this paper respectively, we require

$$\frac{d\omega_L}{dp} < 0$$

$$\frac{d\omega_K}{dp} < 0$$

The elasticities of output and the subfunction X with respect to capital are equal to the factor shares of capital in each, so the parameter condition for the labour share is

$$\left(\frac{1-\theta}{\theta}\right)\omega_K + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right)\frac{\omega_K}{\omega_K + \omega_L} < 0$$

and for capital is

$$\left(\frac{1-\theta}{\theta}\right)\frac{dy}{dk} + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right)\frac{dx}{dk} + \frac{\sigma-1}{\sigma} < 0$$

Setting the labour and capital shares $\{\omega_L, \omega_K\}$ to reasonable values of 0.6 and 0.25 respectively (such that the share going to entrepreneurial labour M is 0.15), figure 14 displays the region of the parameter space in which both conditions are met. The elasticity of substitution between capital and unskilled labour σ must lie in a region around unity that is increasing in size the further below unity is the elasticity between entrepreneurial labour M and the other factors.

The intuition behind this result is as follows. A fall in the relative price of capital increases the quantity of it employed in production, for a given interest rate. For fixed supply of unskilled labour L and entrepreneurial labour M , an increase in K means an increase in the intermediate aggregate X and, for $\theta < 1$, a fall in the income share going to X . If the substitutability between K and L in X (i.e. σ) is low, then the rise in K pushes up the share of L in X quickly enough to offset the fall of the share of X in overall output, and the labour share rises. Conversely, if σ is high enough, then in response to the increase in K the capital share in X will rise quickly enough to allow it (and therefore the nominal investment rate) to rise as a share of Y . The further θ is below unity, the more that the overall share of X falls,

and so the bigger is the region of σ in which these offsetting effects do not dominate.

D 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: US corporate bond spreads and equity yields



This figure shows the yield spread of seasoned US corporate bonds over 10-year US Treasury bonds calculated by Moody's and dividend and earnings yields on the SP500 stock index. Data is taken from Robert Shiller's webpage and the FRED database.

Figure 3: 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 4: Nominal investment-GDP ratios






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




These figures show simple averages 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 ratios of gross capital formation to GDP, at current and constant prices respectively. The source is OECD Statbase.

Figure 6: 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 7: 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 Eichengreen (2015). 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 8: Steady state, baseline setup

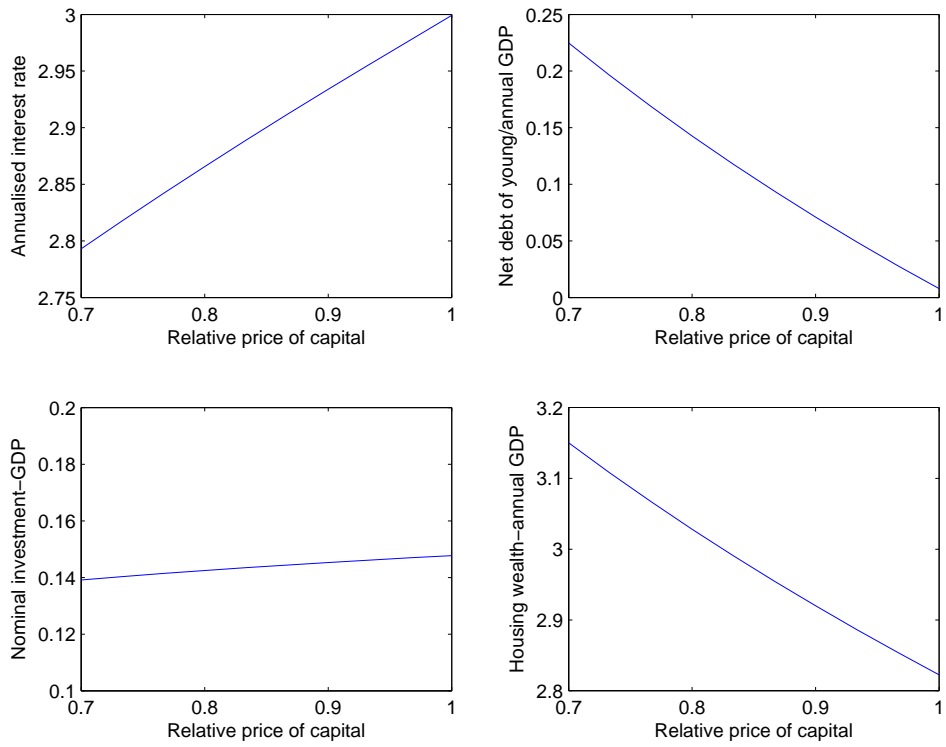


Figure 9: Dynamic solution, baseline setup




Figure 10: Dynamic solution, unanticipated shock

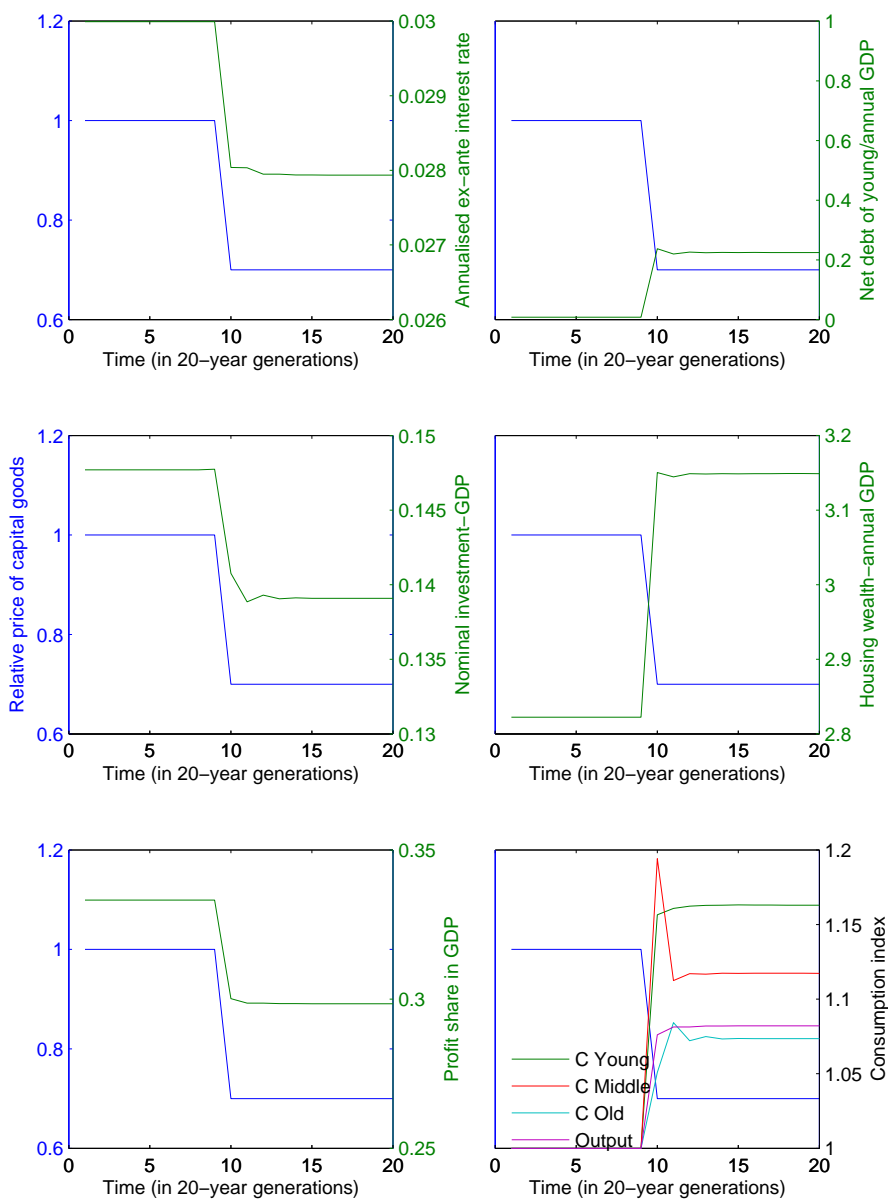


Figure 11: Dynamic solution, no household debt




Figure 12: Dynamic solution, highly elastic production function

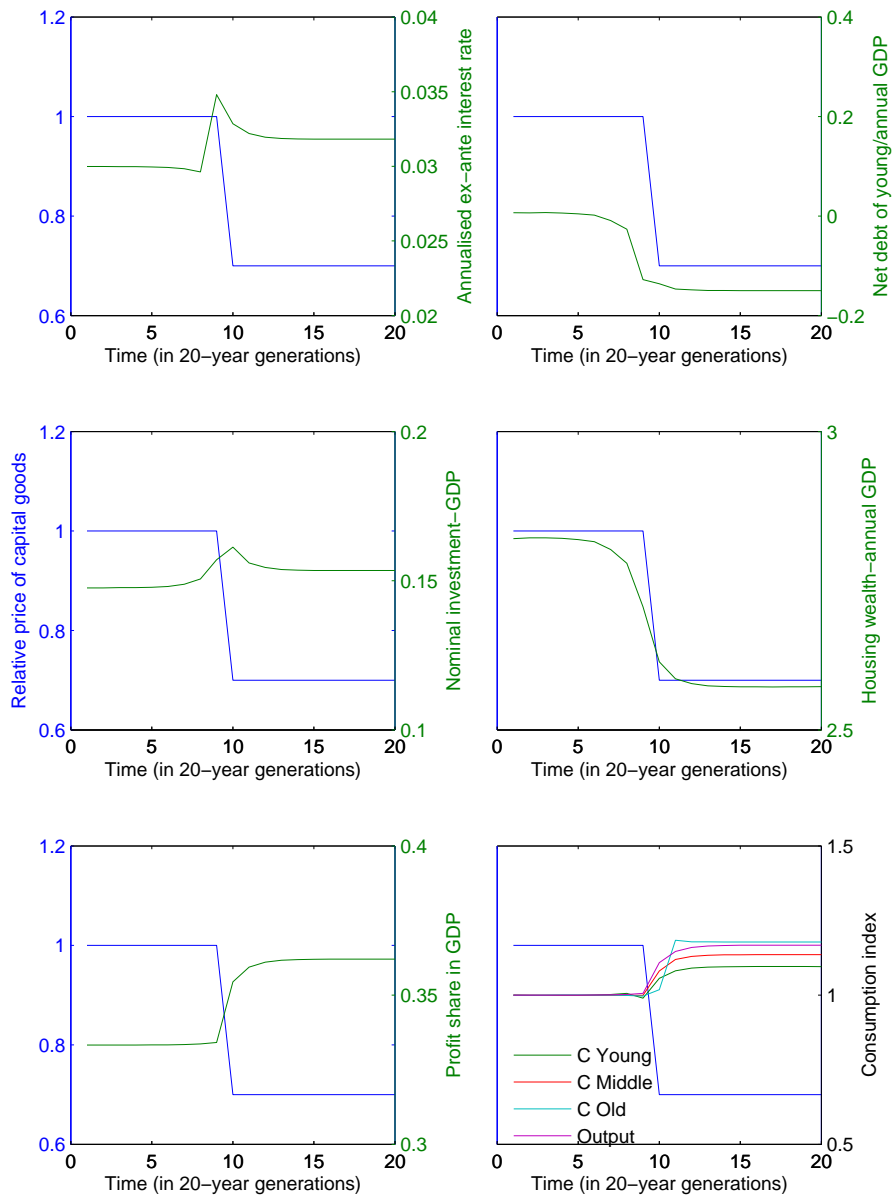


Figure 13: Dynamic solution, small open economy

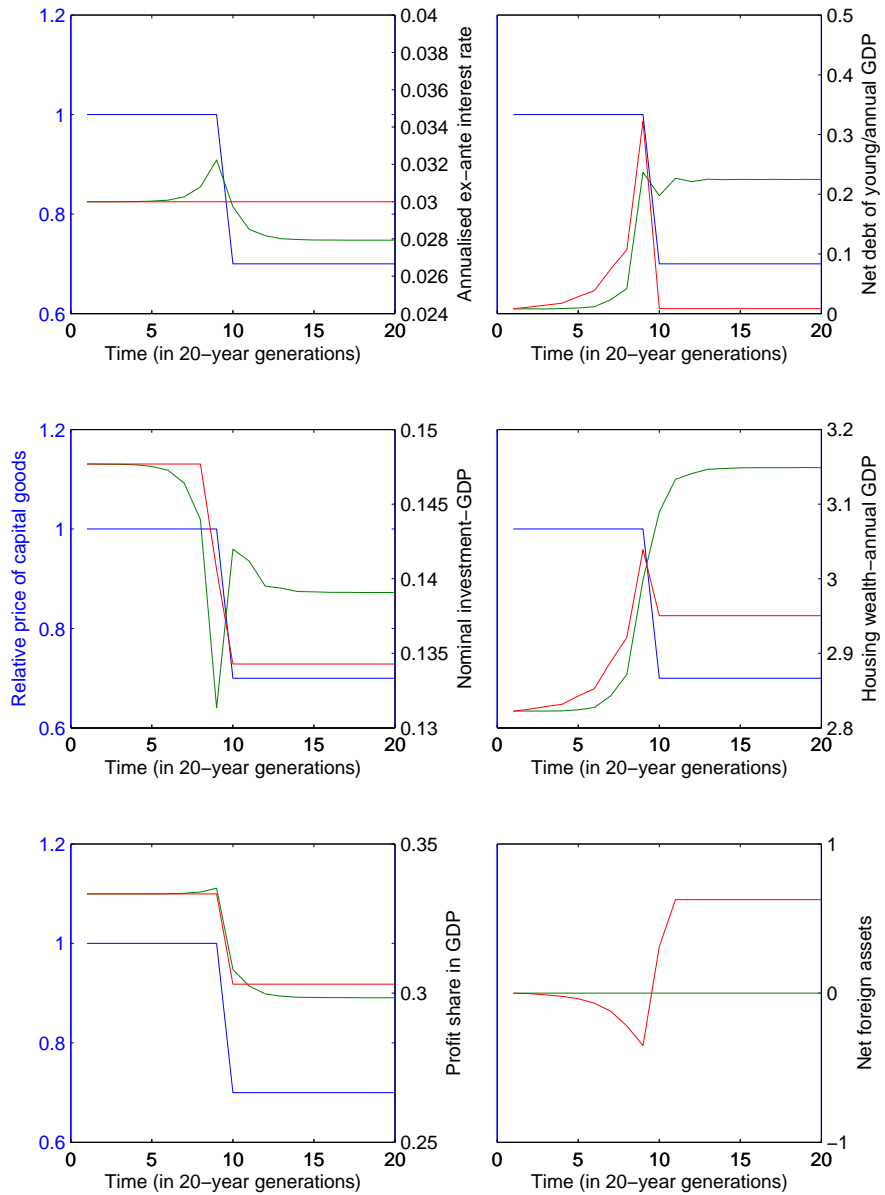



Figure 14: Region of parameter space in which labour share and investment are both increasing in p



The shaded area shows the region of the parameter space of the three-factor production function set out in Appendix C in which the labour share and the nominal investment rate are both increasing in the relative price of capital