Are Low Real Interest Rates Here to Stay?*

Lukasz Rachel and Thomas D. Smith
Bank of England

Long-term real interest rates across the world are low, having fallen by about 450 basis points (bps) over the past thirty years. To understand whether low real rates are here to stay, we need to understand what has caused the decline. The co-movement in rates across both advanced and emerging economies suggests a common driver: the global neutral real rate may have fallen. In this paper we attempt to identify which secular trends could have driven such a fall. Although there is huge uncertainty, under plausible assumptions we think we can account for around 400 bps of the 450 bps fall. Our quantitative analysis highlights slowing global growth expectations as one force that may have pushed down on real rates recently, but shifts in saving and investment preferences appear more important in explaining the long-term decline. We think the global saving schedule has shifted out in recent decades due to demographic forces, higher inequality, and, to a lesser extent, the glut of precautionary saving by emerging markets. Meanwhile, desired levels of investment have fallen as a result of the falling relative price of capital, lower public investment, and an increase in the spread between risk-free and actual interest rates. Looking ahead, in the absence of sustained changes in policy, most of these forces look set to persist, and some may even build further. This suggests that the global neutral rate may remain low and perhaps settle at around 1 percent in the medium to

*This paper is based on material from Bank of England Staff Working Paper No. 571 (Rachel and Smith 2015), which contains additional detail, charts, and analysis. The views expressed in this paper are those of the authors and not necessarily those of the Bank of England or its committees. We would like to thank Scott Simmons for his research assistance during the project and colleagues in the Bank of England's International Directorate for helpful comments. All remaining errors and omissions are our own. Author e-mails: lukasz.rachel@bankofengland.co.uk, thomas.smith@bankofengland.co.uk.

long run. If true, this will have widespread implications for policymakers—not least in how to manage the business cycle if monetary policy is frequently constrained by the zero lower bound.

JEL Codes: E02, E10, E20, E40, E50, E60, F00, F41, F42, F47, J11, O30, O40.

1. Introduction

The downward trend in long-term risk-free interest rates is not a new phenomenon. Alan Greenspan famously highlighted the decline in long-term U.S. bond yields in 2005, which occurred despite the Federal Reserve tightening policy. As Greenspan noted, while the downward trend in yields was clear, the explanations for the fall were not—it was a "conundrum." This conundrum is not unique to the United States either. Since the 1980s, market measures of long-term risk-free real interest rates have declined by around 450 bps across both emerging and developed economies (figure 1). Although there is a lot of variation across countries, the presence of a discernible common trend suggests that global factors are at work.

The decline in global real interest rates has largely occurred against a backdrop of low and stable inflation with little sign of demand overheating. This suggests that the sustained fall in long-term market rates is symptomatic of a fall in the global neutral rate. The global neutral rate is an important policy variable, as it acts as an anchor for a country's equilibrium real rate in the long run.[1]

The purpose of this paper is to contribute to the wider debate about why the global neutral rate may have fallen and what may happen going forward (Summers 2014a, 2014b). Our main contributions are twofold: first, we assemble a rich collection of global data to analyze the main secular trends that could be driving the global neutral real rate. Second, we develop a simple accounting framework to quantify the relative importance of these trends in a coherent way. We then use these insights to explain the fall in the global neutral

[1] For a great overview of the importance of the equilibrium rate for policy, see Fischer (2016).

Figure 1. Comparison of Real Interest Rates

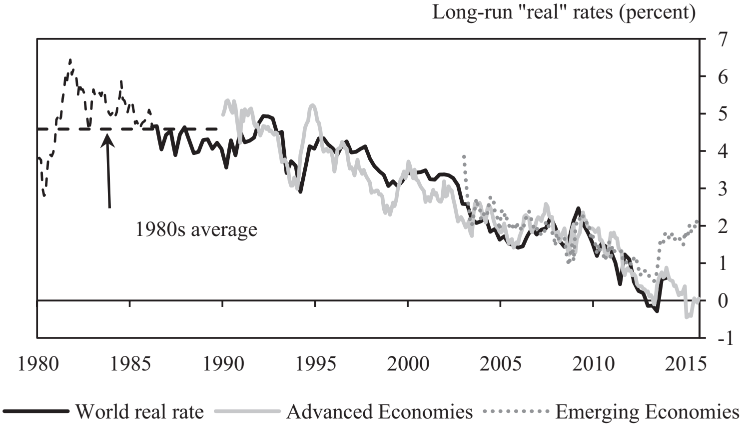

Notes: The "world" real rate (solid black line) is taken from King and Low (2014) and shows the average ten-year yield of inflation-linked bonds in the G7 countries (excluding Italy) over the period 1985–2013. This line has been extended back to the start of the 1980s (dashed black line) using a simple regression linking it to movements in UK ten-year nominal yields and RPI inflation. The solid and dotted grey lines show simpler measures of real rates for different country groups, calculated as the nominal yield on ten-year sovereign bonds minus one-year-ahead inflation expectations from Consensus Economics. Figures have been GDP-weighted together for twenty advanced economies (solid grey) and seventeen emerging markets (dotted grey).

rate in the past, and offer a prediction of how the neutral rate could evolve in the future.

The global neutral rate is largely determined by expectations of global trend growth and other factors shaping preferences for desired saving and investment. We analyze each in turn. First we use a modified growth-accounting framework to analyze the various secular trends that could be affecting global growth. Then we use a simple saving-investment framework to analyze global shifts in desired saving and investment to analyze how changes in preferences could have affected the neutral rate.

Although changes in global trend growth are probably the most commonly cited driver of changes in real interest rates, we find it difficult to account for much (if any) of the pre-crisis fall in global real rates by just appealing to past changes in growth—global growth was fairly steady in the pre-crisis decades.[2] However, the financial crisis does appear to have triggered a wider reassessment of growth prospects, and lower expectations of future growth seem to be playing a role in driving the more recent decline in real rates. Our analysis suggests that slower global labor supply growth (due to demographic forces) and headwinds at the technological frontier (such as a plateau in educational attainment) may cause global growth to slow by up to 1 percentage point (pp) over the next decade. We think expectations of this decline could account for about 100 bps of the fall in real rates seen recently.

Shifts in the balance of desired saving and investment appear quantitatively even more important than changes in growth expectations. Our analysis suggests that the desired saving schedule has shifted out materially due to demographic forces (90 bps of the fall in real rates), higher inequality within countries (45 bps), and a preference shift towards higher saving by emerging market governments following the Asian crisis (25 bps). If this had been the whole story, we would have expected to see a steady rise in actual saving rates globally. But global saving and investment ratios have been remarkably stable over the past thirty years, suggesting that desired investment levels must have also fallen. We pin this decline in desired investment on a fall in the relative price of capital goods (accounting for 50 bps of the fall in real rates) and a preference shift away from public investment projects (20 bps). Also, we note that the rate of return on capital has not fallen by as much as risk-free rates. The rising spread between these two rates has further reduced desired investment and risk-free rates (by 70 bps). Together these effects can account for 300 bps of the fall in global real rates.

When combined, lower expectations for trend growth and shifts in desired saving and investment can account for about 400 bps of the 450 bps decline in the global long-term neutral rate since the late 1980s. Even more difficult than accounting for the past is predicting what might happen from here. Absent major policy changes, our analysis suggests that many of these secular trends look likely to persist and some may even build further. If so, the global neutral

[2]This is consistent with the finding that historically there is only a weak relationship between realized GDP growth and the real interest rate—see Hamilton et al. (2015).

rate may stay low, settling at around 1 percent over the medium to long run.

The rest of this paper is structured as follows. Section 2 starts by describing the concept of the global neutral rate as a long-run anchor for country-specific equilibrium policy rates and the data we use to measure movements in the neutral rate over time. Section 3 discusses the role of economic growth in driving real rates over the past and future. Section 4 analyzes the role of shifts in desired saving and investment as a driver of real rates. Finally, section 5 concludes.

2. Concept and Data

In order to analyze changes in the (unobservable) global neutral rate over time, we need to define what the global neutral rate is and establish an empirical basis for how it has moved.

We define the global neutral rate as the rate to which country-specific equilibrium rates will converge in the long term, absent distortions and shocks. Put differently, the global neutral rate acts as an anchor for equilibrium real rates in open economies.[3] In reality, plenty of distortions and shocks will drive a wedge between country-specific real rates and this long-term anchor. These can be divided into global factors and country-specific factors.[4] Among the global factors, it is useful to distinguish between persistent headwinds that can take several years or even decades to subside (such as the global deleveraging process under way since the crisis) and short-run global cyclical factors (such as global credit conditions or levels of confidence).[5] Among the country-specific factors, a country's cyclical position could drive a country's real rates temporarily higher or lower than the global level. Additionally, a country's structural characteristics—such as its demographic structure, its trend

[3]See Mendes (2014) for an intuitive exposition of the model where a small open economy is a price taker in the global market.

[4]See Mishkin (1984) for evidence of how country-specific risk factors have prevented equalization of real returns across countries.

[5]The evolution of the global neutral real rate since the global financial crisis is particularly uncertain because it is too early to tell how much of the most recent decline is cyclical, and hence will prove temporary, or persistent, but not reflective of the long-term neutral rate.

rate of productivity growth, or the quality of its institutions—may drive individual country rates persistently above or below global R*. An individual country's monetary policy stance can also temporarily drive a wedge between a country's equilibrium real rate and the actual (real) policy rate.[6]

In this paper we use ten-year government bond yields, adjusted for inflation, as an indicator for how the global neutral rate has moved over time. These market measures are subject to the same shocks and distortions as those for country-specific equilibrium rates (highlighted above) as well as others such as the impact monetary policy regime shifts can have on term premia and distortions specific to government bond markets. To account for such distortions, we make a number of adjustments in our analysis.

To smooth out the effect from country-specific factors, we focus our analysis on global measures of real rates aggregated across countries. To sidestep cyclical issues, we focus our analysis on very low frequency movements in global data. For example, in section 3 we consider five-year averages of the global growth rate, and in section 4 we focus primarily on decade averages of the real interest rate. To avoid the impact that monetary policy regime shifts may have had on term premia, we focus our analysis on explaining changes in real rates since the average of the 1980s, rather than the 1980s peak. The average of the 1980s also corresponds to the point where real rates hovered in the early 1990s, which is after many of the largest monetary policy regime shifts had occurred (see International Monetary Fund 2014b).[7]

[6]For a model-based estimate of the equilibrium real interest rate which is similar in spirit to what we focus on in the present paper, see Bomfim (1997). For an example of how the equilibrium rate features in policy briefings and deliberations, see the Federal Open Market Committee briefing documents available at https://www.federalreserve.gov/monetarypolicy/files/fomc20010515bluebook20010510.pdf.

[7]In principle, long-term forward rates would be a better market measure to use in this analysis than long-term spot rates, as forward rates exclude some of the distortions associated with term premia. A lack of historical data on forward rates across countries hinders such analysis at the global level, but King and Low (2014) show that for countries where data are available (such as the United Kingdom), forward rates have tended to co-move with spot rates over the past thirty years.

Changes to the structure of government bond markets could also have affected long-term yields over time. Over the past thirty years there have been significant shifts in both the demand and supply of government bonds. In the decades before the crisis, the supply of high-quality government debt increased steadily. This rise in supply was partly met by increased demand for safe assets by emerging market governments, but in net terms the supply of safe assets increased in the decades before the crisis, potentially putting upward pressure on government bond yields and counteracting some of the fall in long rates observed in the data. However, since the crisis that picture has changed markedly. Deteriorating fiscal positions and ratings downgrades have seen the supply of high-quality debt fall sharply, while demand for safe assets has risen as a result of permanent regulatory changes and cyclical central bank action. This suggest that since the crisis it may well be true that the net supply of safe assets has deteriorated (Caballero and Farhi 2013), and this may have contributed to the fall in actual real interest rates seen recently. However, we still need to appeal to other factors to explain the steady decline in real rates before the crisis and arguably some of the fall since—as highlighted in sections 3 and 4 below.

Another potential issue with market measures of global real rates based on long-term bond yields is that they are affected by low short-term interest rates. Some have even argued that "global rates are low because monetary policy is loose." If this were the case, low rates would not pose a policy dilemma. The solution would be trivial and the downward trend in long rates would simply reverse when central banks tightened policy. But as Bernanke (2015a) and Broadbent (2014) have pointed out, this view of the world is unlikely to be correct, as the decline in actual real interest rates has occurred against a backdrop of contained inflation with little sign of exuberant demand growth. Indeed, global growth and inflation have, if anything, disappointed in the most recent recovery, despite policy rates being historically low. This suggests that observable interest rates have merely followed their unobservable "equilibrium" counterparts—if policy had been tighter, inflation would have been lower and demand would have been too weak to deliver full employment. Ever looser monetary policy is not the cause but the consequence of the fall in long-term rates.

This intuition has been formalized in econometric models, which aim to extract measures of equilibrium interest rate from observed data. Laubach and Williams (2003) perform this sort of exercise for the United States and find that U.S. R* has declined by around 450 bps since the 1960s, and by around 300 bps since the 1980s. The authors suggest that secular trends related to changes in trend growth and shifts in saving and investment preferences are responsible for this decline—not monetary policy. We use a similar taxonomy in our analysis at the global level: section 3 focuses on growth and section 4 on preferences.

3. Global Growth and Real Interest Rates

One of the most frequently cited drivers of changes in real interest rates is changes in trend growth. Before analyzing how global trend growth might have changed, it is worth dwelling on how changes in growth affect real rates. The Euler equation in the neoclassical model pins down the real rate by time preferences, the pace of technological progress, and, in some formulations, population growth (equation (1)).

$$r^* = q/\sigma + \theta + (\alpha \cdot n), \tag{1}$$

where $r*$ is the real interest rate consistent with inflation at target and zero output gap in the long run; σ is household's intertemporal elasticity of substitution in consumption (preference for smoothed consumption); q is the rate of labor-augmenting technological change; θ is household's rate of time preference (patience); n is the rate of population growth; and α is the coefficient on the rate of population growth.[8]

The preference parameters are particularly noteworthy, as they affect the link between growth and real rates. Estimates of these parameters are difficult to obtain, but one meta-study by Havranek et al. (2015) suggests that the global average for σ (household preferences for smoothed consumption) could be around 0.5. This implies

[8]The infinite-horizon representative-agent Ramsey model does not include population growth in the steady-state real rate formulation. But there may be good reasons to include it (e.g., see Baker, Delong, and Krugman 2005).

that the mapping between productivity growth and real rates may not be one-for-one, but potentially one-for-two, i.e., a 1 pp fall in productivity growth could lead to a 2 pp fall in real rates.

The link between population growth (n) and real interest rates is even less certain. The standard neoclassical model does not include population growth as a driver of real interest rates at all (so $\alpha = 0$). Yet excluding population growth entirely from the analysis seems an omission. Alvin Hansen's (1934, 1938) original secular stagnation hypothesis emphasized the role of population growth in driving down the rate of return on capital. Baker, Delong, and Krugman (2005) have since argued along similar grounds, noting that in some models (such as the Solow model) population growth plays a role in determining real interest rates. If labor and capital are complements, then slower population growth should reduce the marginal product of capital, as firms have fewer workers to get the best out of their machines. As a result, slower population growth should mean the rate of return on capital falls, pushing down real interest rates. Given the lack of empirical estimates for α, we assume there is some role for population growth in driving real interest rates but a one-to-one mapping is likely to be an upper bound (i.e., $\alpha \leq 1$).

It is difficult to account for much (if any) of the pre-crisis fall in real rates by just appealing to past changes in growth, because global growth was fairly steady in the pre-crisis decades—averaging 3 to 4 percent per year. However, the crisis itself may have triggered a broader reassessment of trend growth expectations. And greater pessimism about future growth could be playing an important role in driving the decline in real rates we have seen most recently.

Broadly speaking, there are three factors that might lead trend growth to weaken over the future: (i) a reduction in labor supply growth; (ii) a slower rate of catchup in emerging markets; and (iii) weaker growth at the technological frontier. So how pessimistic should we be about each?

3.1 Labor Supply

Growth in global labor supply peaked at just over 2 percent in the 1980s as the demographic dividend from the post-war baby boom (and falling mortality rates in emerging market economies) fed through to the labor market. Since then, the pace of global

Figure 2. Global Population Growth and Labor Supply Growth

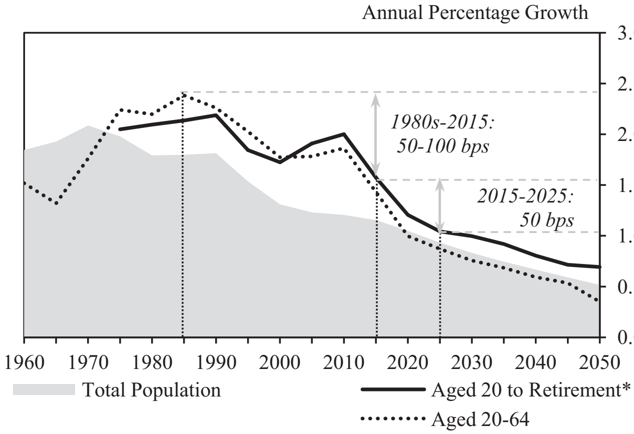

Notes: The solid black line shows growth in the global working-age population, aged twenty to retirement, using historical and forecast figures from the United Nation's Population Projections. Retirement ages are calculated using OECD data on average effective retirement ages. Over the future, global retirement ages are assumed to grow by one and a half years every decade to keep pace with expected rises in longevity.

labor force growth has slowed by a third, reducing its contribution to global GDP growth by around 1 pp (figure 2).

The age structure of the population means further falls are likely: global population growth slowed sharply in the mid-1990s, and that effect is now feeding through to global labor supply (figure 2). Expectations of slower labor supply growth could reduce global growth by around 0.5 pp over the next decade. As noted earlier, the mapping from labor force growth to real rates is highly uncertain (α in equation (1) is unknown). If we assume a mapping of one-to-one as an upper bound between labor force growth and real rates (i.e., $\alpha \leq 1$), this would suggest that expectations of slower labor force growth could account for up to 50 bps of the fall in real rates we have seen.

Slower labor supply growth could also have wider effects on economic growth and real rates via productivity spillovers—boosting productivity by alleviating resource constraints or reducing productivity by reducing the returns to innovation (see Kuznets's seminal

1960 paper for a discussion). However, given that the effects can go both ways, we do not explicitly account for these additional spillovers here.

3.2 Catchup Growth

Another driver of global growth is productivity catchup. As countries accumulate more capital and improve efficiency by adopting the latest technologies from overseas, productivity per worker rises. In the early 2000s, rapid catchup growth among Asian economies boosted global growth significantly, offsetting the impact of slowing labor force growth.

Given that the pace of catchup has accelerated in recent years, it seems implausible that slower convergence has been a key driver of the steady decline in real interest rates we have seen over the past few decades. That said, the robust rate of catchup growth seen in the early 2000s is relatively unusual when compared with a broader sweep of history (figure 3). Between 1980 and 2010, GDP per capita growth in the United States (widely used as a proxy for the technological frontier) was actually faster than the average across the rest of the world in fifteen out of thirty years—so the rest of the world spent just as long falling further behind the frontier as catching up.

The difficult question is what will happen to the pace of catchup going forward. In order for the recent positive trend to continue, many emerging markets will need to overcome the middle-income trap and continue to avoid geopolitical and financial crises. There are some grounds for optimism on that front. The unusually rapid growth of many emerging markets earlier this century shows that it is possible for sustained periods of catchup to take place. In addition, the rising importance of digital technologies in driving innovation, combined with the spread of the Internet and other communication technologies (e.g., distance learning), means it is now easier for ideas and skills to be shared across borders more quickly. On the other hand, the mixed performance of the 1980s and 1990s, combined with ongoing concerns about the stability of China's financial system and the rise of emerging market indebtedness, suggests that a more gloomy outlook is possible (at least in the near term).

We take a neutral view and assume that the contribution of catchup growth to global growth remains stable at its average rate

Figure 3. GDP per Capita by Region Relative to the United States

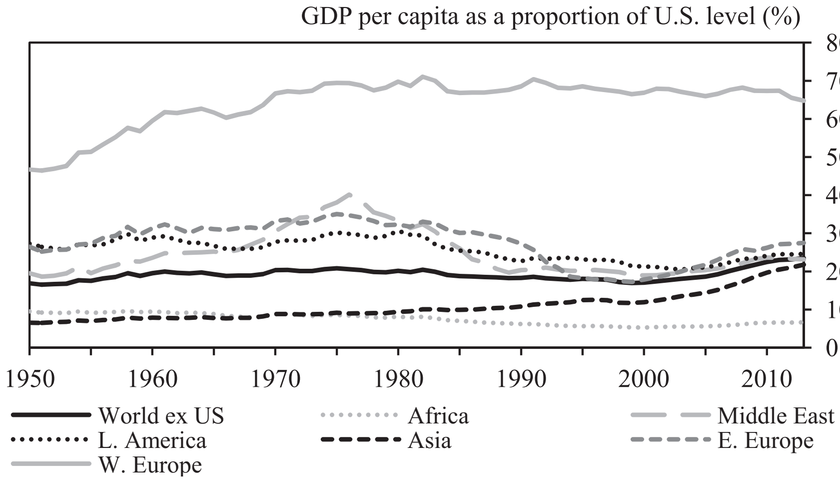

Note: Figures are calculated from the Conference Board's Total Economy Database (2013).

of the past twenty years—so not as fast as the early 2000s, but not as slow as the 1990s either. This equates to catchup growth continuing to add 1 pp to global growth per year and means we do not think the decline in global real interest rates is driven by slowing expectations of the pace of convergence.

3.3 Growth at the Technological Frontier

The other driver of global trend growth is the pace of growth at the technological frontier (proxied by productivity growth in the United States). Gordon (2014a, 2014b) has championed the view that several structural headwinds will hold back U.S. growth in the future, including further falls in the pace of educational attainment, rising inequality, and fiscal drag. Gordon suggests that these factors could drag down trend growth at the frontier by up to 1 pp—either by directly affecting the supply side or via weaker demand leading to hysteresis. Having interrogated Gordon's analysis, the argument on the educational plateau seems justified—the number of years of schooling per worker cannot go on rising indefinitely and has already started to slow. Moreover, Gordon's figures are backed up by others, e.g., Jorgenson and Vu (2010). But Gordon's estimates of the impact from inequality and fiscal policy seem high given uncertainties over fiscal multipliers and the overlap between the inequality and education arguments (rising inequality makes it more difficult for poorer households to afford college). In our view it seems more likely that growth at the frontier could be 0.5 pp weaker in the decade ahead due to these headwinds, i.e., half Gordon's estimate—Rachel and Smith (2015) contains the detailed analysis looking at the magnitude of each headwind that supports this claim. Given that the multiplier between productivity growth and real rates probably lies between 1:1 and 1:2 (i.e., $0.5 \leq \sigma \leq 1$ in equation (1)), expectations of slower frontier growth of around 0.5 pp could be pulling down on real rates by 50–100 bps.

The other major uncertainty is over the pace of innovation. Gordon (2014b) argues that the recent weakness in U.S. productivity growth is a longer-lived phenomenon stretching back to the 1970s and hence will continue going forward. Others, such as Brynjolfsson and McAfee (2014), see the recent slowdown as a blip—growing pains as a result of disruptive new digital technologies that will soon give way to rapid productivity gains. Some, such as Kurzweil (2005) even argue we are about to enter a phase of unprecedented growth, fueled by artificial intelligence and robotics. Our reading of the above arguments is that Gordon's characterization of recent history and the near future is the most compelling. Measured U.S. productivity growth has been weak since the 1970s; it was lifted temporarily by the information and communication technology (ICT) boom, but has since fallen back to its sluggish underlying rate. In the absence of clear advances in technology, it seems reasonable to assume this trend will continue in the near term—particularly given the recent weakness in productivity globally. Consequently, we assume that the pace of innovation will continue at its recent sluggish rate—consistent with the experience of the past thirty years, save for the ICT boom. However, a word of caution is warranted here. The further we peer into the future, the more likely a positive technology shock is to occur, so there are substantial upside risks to our forecast—Brynjolfsson and McAfee (2014) or even Kurzweil (2005) could eventually be proved right.

Figure 4 and the accompanying table bring together the above insights into one place—showing the contributions to each of the three factors to global growth. Looking over the past, it is reasonably clear that growth has not changed materially since the 1980s—averaging 3 to 4 percent per year. In fact, growth in the years before the crisis was actually a little higher than in the 1980s. Consequently, it is difficult to account for much (if any) of the pre-crisis fall in real rates by appealing to significant changes in the pace of global trend growth. Yet while growth may not have fallen much over the past, there are reasons to think global trend growth will slow in the future. Although there is a great deal of uncertainty, if we add up all the factors analyzed above, we think we can come up with a reasonable case for why global growth could slow by up to 1 pp over the next decade or so—largely relating to slower labor supply growth and structural headwinds at the technological frontier. Some of these factors could have been predicted before the financial crisis, but arguably the crisis was the trigger that caused financial markets to focus attention on the issue. Indeed, the 1 pp decline in global growth we predict is similar in scale to the downward revisions to medium-term growth forecasts that the International Monetary Fund (IMF) and private-sector forecasters have made since 2008. Depending on the mapping, these weaker growth prospects could account for around 100 bps of the fall in real rates we have seen post-crisis.

4. Desired Saving and Investment and Real Interest Rates

The previous section explored the link between global trend growth and the global neutral rate of interest. But as we have seen, changes in global growth can only explain part of the secular decline in global real rates over the past thirty years—mainly in the post-crisis period. Other factors must also be responsible for driving the long-term decline in the global neutral rate.[9] Since the real interest rate is the price of future consumption expressed in terms of consumption today, shifts in time preferences that describe how households

[9]Hamilton et al. (2015) arrive at a similar conclusion.

Figure 4. Global Growth Accounting

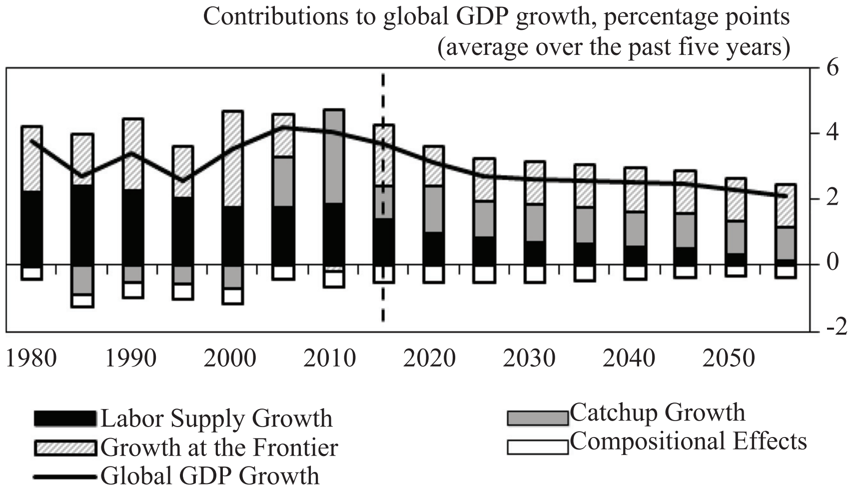

	1980 to 2015	2015 to 2030
Change in Global Growth	0.0 pp	–1.0 pp
Labor Supply Growth	–0.8	–0.5
Catchup Growth:	+1.0	0.0
Growth at the Frontier:	–0.2	–0.5
Educational Plateau	–0.2	(0 to –0.2)
Inequality	0.0	(0 to –0.6)
Fiscal	+0.2	(0 to –0.2)
Technological Progress	–0.2	(–)

Notes: Global growth is expressed in constant PPP-weighted 1990 dollars. Figures are expressed as averages over five-year intervals so as to smooth out cyclical variations and to provide a simple metric for the underlying trend. The white "compositional effect" bars in the figure show the impact on average global per capita incomes of having high population growth in low-income countries. Figures are constructed using data from the U.S. Conference Board's Total Economy Database, IMF, UN, and the authors' own calculations. Figures in parentheses show the range of uncertainty of the impact of each headwind on productivity growth at the frontier. Our central assumption is that the effects of the education plateau are towards the upper end of the range, while the fiscal and inequality effects are towards the lower end. Global growth is expected to slow by up to 1 pp weaker over the next decade due to slower labor supply growth (–0.5 pp) and slowing productivity growth at the frontier (–0.5 pp). The multiplier from changes in growth to real rates is likely to range from 1:1 to 1:2 for productivity growth and 1:0 to 1:1 for labor supply growth. This implies that slowing growth expectations could have reduced real rates by between 50 and 150 bps. We quote the midpoint of this range in our calculations above.

spread consumption over their life cycle will also move real rates around. The simple growth model we introduced in the previous section is too parsimonious to permit the analysis of the kind of preference shifts we are concerned with here: the secular trends considered below would show up, in a reduced-form way, as changes in the value of the parameters of the Euler equation, and would not be very informative about the workings of various mechanisms or their magnitude. So instead of a formal model, in this section we utilize a saving-investment (S-I) framework to shed light on these phenomena.

The basic idea behind this framework is that, given growth expectations, the neutral rate will depend on agents' preferences for desired saving and desired investment. Intuitively, desired saving will tend to rise as real rates increase (the saving schedule should slope upwards), because higher rates generate higher returns on saving and yield higher future consumption. By contrast, desired investment will tend to fall as real rates rise (the investment schedule should slope downwards) because the real rate is a key component of the user cost of capital, so as real rates rise it becomes more costly to invest. The focus of our analysis is on changes in *desired*, rather than *actual* saving and investment. For the world as a whole—as for any closed economy—actual saving and investment will always be equal by identity. But the sensitivity of desired saving and investment (the slopes of the curves) and the forces that shift them (preference shifts) will be key in determining the actual (equilibrium) level of saving and investment and the observed real interest rate.

An important source of uncertainty when using this framework for quantitative analysis is that we do not know the sensitivities of desired saving and investment to real rates: the slopes of the saving and investment schedules are unobservable, so we need to rely on empirical estimates. The key difficulty when estimating these slopes is endogeneity: interest rates and S-I ratios may be driven by common factors. For example, a more optimistic demand outlook would raise investment and interest rates simultaneously. This is why studies that estimate elasticities using time-series correlations can produce a wide range of estimates. To make our exercise robust, we take an average of available estimates from the literature and then later conduct sensitivity analysis to show the impact of varying these assumptions.

Table 1. Estimates of the Elasticity of Saving with Respect to the Real Rate

Author of Study	Elasticity
Blinder (1975)	0.03
Boskin (1978)	0.4
Carlino (1982)	0
Carlino and DeFina (1983)	0
Gylfason (1981)	0.3
Heien (1972)	1.8
Howrey and Hymans (1978)	0
Summers (1982)	1.3
Taylor (1971)	0.8
Wright (1967)	0.2
Average	0.5

Source: DeFina (1984).

For the elasticity of saving, this suggests an elasticity of 0.5 (table 1), although the range of estimates is admittedly very wide. For the elasticity of investment, we rely on more recent studies that aim to overcome the endogeneity problem by using structural models or by employing cross-sectional data (such as Ellis and Price 2003, Gilchrist and Zakrajsek 2007, and Guiso et al. 2002). These tend to find that long-run elasticities are between –0.5 and –1. We assume an elasticity of –0.7; this makes investment more sensitive to interest rates than saving.

Together, these assumptions form the basis of the slopes of the curves shown in figure 5. We think the slopes have been calibrated based on a fairly neutral reading of the range of estimates in the literature, but we recognize the wide bands of uncertainty. If we are wrong about one of the slopes of the curves—say the investment schedule is shallower—then it becomes more likely that shifts in the investment curve (rather than the saving curve) have been responsible for more of the fall in real rates we have seen. The slopes of the curves thus matter in terms of the relative weight one puts on different explanations for the fall in real rates, but should not necessarily affect our ability to account for the scale of the fall overall.

Figure 5. The Saving-Investment Framework

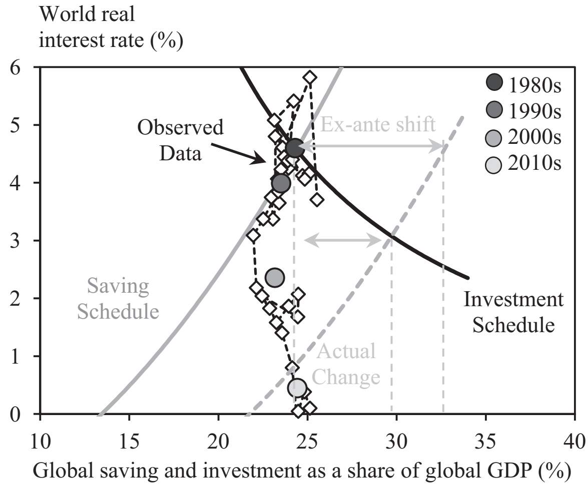

Notes: Constructed using data from the IMF, King and Low (2014), and authors' calculations. Global saving and investment rates are reported by the IMF. The world real rate is taken from King and Low (2014). Specific calibrations of neoclassical or overlapping-generations models can deliver very different results.

A more detailed discussion of the sensitivity of the results to these assumptions is provided at the end of this section.

We choose the average level of real rates during the 1980s as the reference point of our analysis. This choice is driven by the observation that the initial decline in real rates from the peak of around 6 percent in the early 1980s was likely driven by the disinflation policies during Volcker's Chairmanship of the Federal Reserve Board (IMF 2014b). The average of the 1980s also equates to the level of real rates in the first half of the 1990s, when such monetary policy regime shifts had largely been completed.

We focus on ex ante shifts in the schedules, i.e., shifts that are independent of the moves in the real rate itself. We can then determine the impact on the real rate by comparing the new and old equilibria. Precisely because the real rate adjusts, the ex post change in actual saving or investment will tend to be smaller than the ex ante, or desired, change. Figure 5 illustrates this with an example: a preference shift increases desired saving for any given interest rate.

This shifts the saving schedule to the right. But as the desired level of investment is unchanged (for a given interest rate), this shift would push down on the interest rate until desired investment is equal to desired saving. As a result, the actual increase in saving is smaller than the shift in desired saving.

One striking feature of the data is that despite the 450 bps fall in global real rates, global saving and investment have remained fairly stable as a share of global GDP over the past thirty years (diamonds in figure 5). This vertical pattern could suggest that either saving or investment is insensitive to changes in real rates (one of the curves is vertical). While mindful of this possibility, we assume that the slopes of the curves match empirical estimates in the literature, which implies that *both* curves must have shifted. Various factors have been put forward to explain such shifts. Our approach is to run through them and try to quantify the size of each effect on real rates. We begin by focusing on trends that have affected the saving schedule: changes in the demographic structure of the global population, rising inequality, and a preference shift by emerging market governments towards higher saving (the emerging market saving glut). We then analyze trends that have mostly affected desired investment: the fall in the relative price of capital, shifts in public investment, and the changes in the spread between broad rate of return to capital and the risk-free rate.

We view the estimates below as largely independent of the growth effects we have identified in the last section, and hence claim that our results should not suffer from double-counting. Although we do not have a comprehensive general equilibrium model to prove this point, we rely on the economic logic that is commonly accepted in the macroeconomic literature: namely that in macroeconomic equilibrium, the response of the real interest rate will cushion the impact on economic growth. Up to the most recent crisis, the real interest rate was falling, but it remained above zero. Thus, for most of the period under study, the equilibrium rate had sufficient room to respond, and precisely because of that response the secular forces we discuss should not have had much impact on the growth rate up until recently. Take inequality as an example: given a fixed real rate of interest, rising inequality could have lowered aggregate demand and pushed down on the growth rate. But as long as there was room for the interest rate to fall, this direct effect would have been offset

Figure 6. Cross-Country Saving Rates and Dependency Ratios

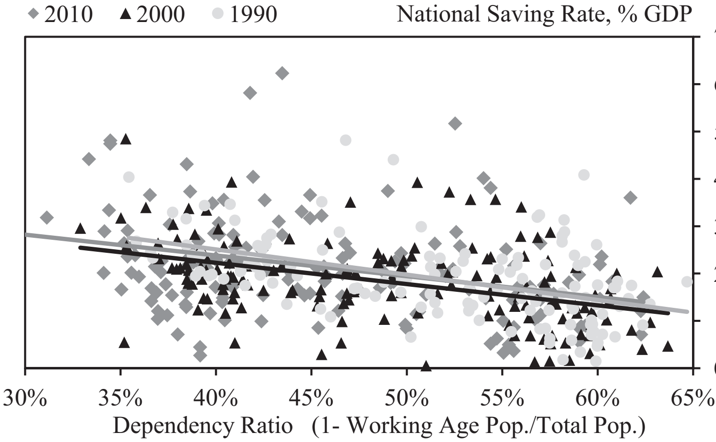

Sources: IMF's *World Economic Outlook* and UN Population Statistics.

Notes: The scatter plot shows cross-country values of dependency ratios and saving rates for three different decades. Working-age population is defined simply as those aged twenty to sixty-five.

in equilibrium. Confirming this intuition, the moving average of the global growth rate appears to have been relatively stable over much of the time period we analyze, as highlighted in section 3.

4.1 Demographics

The key mechanism through which demographics plays a role in our analysis is the life cycle of saving. Consumption is fairly stable over the life cycle, but income is hump shaped, so people of working age are those who tend to save the most. Consequently, the greater the proportion of the population that is of working age, the higher the desired level of saving in aggregate is likely to be.

This simple intuition is discernible in cross-country data on national saving rates (figure 6). There is a significant negative relationship between the dependency ratio (defined as the proportion of the population not of working age) and national saving rates: every 1 pp fall in the dependency ratio translates to around a 0.5 pp rise in national saving rates. This relationship is stable through time,

Figure 7. Shift of the Desired Saving Schedule Due to Changes in the Composition of the Global Population

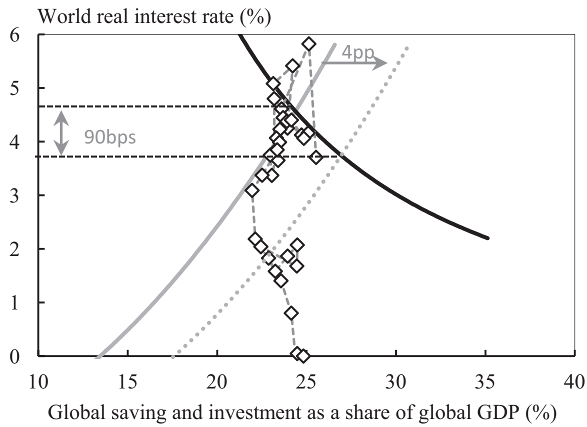

Sources: IMF and authors' calculations.

suggesting that it is robust and can be used to calibrate the ex ante shift of desired saving caused by demographic changes at the global level.

Over the past thirty years the proportion of dependents has fallen from around 50 percent of the global population to 42 percent. The main driver of this decline has been a fall in the proportion of young dependents—reflecting the slowdown in demographic growth discussed earlier. This effect has more than offset the gradual rise in the proportion of old-age dependents linked to aging societies. Using the estimated cross-country relationship depicted in figure 6, the 8 pp fall in the global dependency ratio translates to a 4 pp rise in desired saving, for a given real interest rate.

The 4 pp rise in desired saving can be illustrated by a rightward shift in the saving schedule (figure 7). The effect on the global real rate can then be easily read off the y-axis of the chart by comparing the two intersection points. This suggests that the effect of the fall in the dependency ratio has been to lower real rates by around 90 bps.[10]

[10] Aside from changes in the global dependency ratio, changes in average age of the working-age population may also shift desired saving. This is because over

Looking ahead, the dependency ratio is likely to stop falling or even rise as a growing share of the world's population enters retirement age. At face value this suggests that the effects of demographics on desired saving and hence real rates could reverse. But the extent of this reversal is uncertain. Two factors that may limit the extent of the reversal are increases in retirement ages and longevity. Effective retirement ages across countries have been increasing gently over the past fifteen years. This trend largely reflects changes in old-age participation rates rather than official retirement ages, which have been fairly static. If this trend continues, the average retirement age across the OECD could reach sixty-seven by 2030—enough to halt any uptick in the dependency ratio. In addition, an increase in longevity over and above this increase in retirement age may mean that people of working age choose to save a larger share of their income to fund a longer retirement—pushing up on saving rates (Carvalho, Ferrero, and Nechio 2016). These two factors suggest that there could be little to no reversal of the impact of demographic forces on global real rates in the years ahead.

On the other hand, the shifting composition of dependents from young to old could mean that the reversal is larger than suggested by simple dependency metrics. Old-age dependents can have much lower net saving rates than young dependents due to the distinct pattern of consumption over the life cycle. In advanced economies (such as the United States), consumption tends to drift up in retirement, particularly in the last few years of life. This is driven largely by greater consumption of health care (hospitalization, emergency procedures, etc.).[11] Thus in advanced countries, a rising share of retirees will have a disproportionately large *negative* effect on desired saving—implying a faster turnaround in the impact of demographic trends on desired saving globally compared with the simple dependency ratio metrics. Indeed, studies that look at the "support ratio"—which can be thought of as a measure of the dependency ratio that corrects for the differential spending patterns highlighted

their working lives, individuals' saving rates tend to rise up to a certain age and then decline beyond that. Our analysis suggests that this effect is second order compared with the 4 pp shift due to the dependency ratio discussed above. See Rachel and Smith (2015) for a fuller discussion.

[11] This stands in contrast to a more flat pattern in emerging countries, such as China—see Lee and Mason (2011) for more details.

above—find that the impact of demographics on desired saving looks set to reverse (Lee and Mason 2011).

On balance, we think that the evidence based on support ratios provides the more reliable steer, and so we expect the demographic effect on rates to gradually reverse. The current projections of dependency ratios alone would suggest that demographics will be neutral on global real rates going forward, but given that older people in advanced economies are big spenders, desired saving is likely to decrease. Quantitatively, the current forecasts for the support ratios in advanced and emerging economies suggest that the reversal of the demographic effect over the next twenty years is likely to be about half that of the downward drag on real rates over the past twenty years.

4.2 Inequality

Changes in the distribution of income can affect desired saving because the rich and poor tend to save different proportions of their income. To the extent that the rich save more, rising inequality will result in lower consumption, higher desired saving, and hence a lower equilibrium real rate.[12] Empirical evidence supports the notion that the rich do save more, although the range of estimates available is relatively limited and primarily covers the United States.[13] The seminal contribution is by Dynan, Skinner, and Zeldes (2004), who show that average saving rates and marginal propensities to save tend to rise with the level of income. This is confirmed by more recent evidence: for example, Cynamon and Fazzari (2016) show that the richest 5 percent save much more than the rest (with saving rates around three times as high), and Saez and Zucman (2014) give a long-run perspective on the high saving rates of the wealthy.

Two global trends have taken hold since the 1980s: inequality *between* countries has fallen because developing countries, particularly those in Asia, have been catching up with their Western peers; but at the same time income inequality *within* countries has been rising (Piketty 2014). The relationship between national saving rates

[12] This general mechanism has been incorporated in recent models, e.g., in Eggertsson and Mehrotra (2014) and Kumhof, Ranciere, and Winant (2015).

[13] Crossley and O'Dea (2010) provide a UK perspective.

and a country's level of development is not clear cut—for example, many Asian economies have very high saving rates despite having relatively low incomes. As such, it is less clear what impact the reduction in inequality between countries has had on global saving rates. By contrast, the relationship between individual saving rates and individual income levels within countries is better established. We therefore focus on calibrating the impact of this latter effect on desired saving rates globally.

To isolate the effect of rising inequality within countries on desired saving, we perform a thought experiment using U.S. data. First, we take the saving rates from Dynan, Skinner, and Zeldes (2004) for different income quintiles. Then we combine those saving rates with data from the U.S. Census Bureau showing how income shares across the population have changed since the 1980s. Over this period, the richest fifth of the population—who are also the keenest savers—have seen their share of national income rise by around 7 pp. On average, this group saves an extra third of their income compared with the rest of the population, so this shift in the income distribution translates to a rise in desired saving of around 3 pp. Changes in the income distribution among the four lower quintiles of the population reduce this figure to a net rise of around 2 pp in aggregate. This acts to drive the real interest rate downward by 45 bps.[14]

Our assessment of the effect of inequality relies primarily on U.S. data, which could overestimate the size of the effect at the global level, since the rise in inequality in the United States has probably been a little larger than average. But there are also other reasons why inequality may have had a bigger effect on real rates than postulated above.

First, over the past thirty years, the share of global income earned from owning capital has risen as the labor share has fallen. Mechanically, this should push up on desired saving if the propensity to save out of capital income is higher than from labor income. A lack of empirical data on saving rates from these different sources of income prevents detailed analysis of this channel, but we note it as an upside risk.[15]

[14] We provide further cross-checks on this figure in Rachel and Smith (2015).

[15] See Rachel and Smith (2015) for further discussion.

Second, inequality between countries has been falling as real incomes in the bottom two-thirds of the global income distribution have risen rapidly. The impact of this income shift on global saving depends on relative saving rates between lower-income countries that are catching up and advanced economies. National saving rates in advanced economies and emerging economies were virtually identical up until the year 2000, so the effect on global desired saving of emerging markets catching up would have been negligible until then. But since the turn of the century, saving rates in emerging markets have actually increased *above* those in advanced economies. This means that faster income growth in emerging market economies may have raised global desired saving.[16] We are cautious of including this effect separately, as we think it is at least partly driven by cyclical factors and partly related to the emerging market saving glut story, which we discuss next. To avoid double-counting, we exclude it from our analysis.[17] Instead, we focus on the future of inequality within countries and how that affects global saving.

The future of within-country inequality will ultimately depend on policy. Piketty and Saez (2014) point out that "inequality does not follow a deterministic process. In a sense, both Marx and Kuznets were wrong. There are powerful forces pushing alternately in the direction of rising or shrinking inequality. Which one dominates depends on the institutions and policies that societies choose to adopt." So any forecast of inequality will necessarily rely on heroic assumptions about inherently unpredictable political processes. That said, it may still be useful to consider the main *economic* forces that determine income inequality in the long run, taking policies and political processes and institution as given.

The most widely used model for analyzing *labor income inequality* is based on the idea of a race between education and technology (Goldin and Katz 2007). The basic intuition is that more education leads to a rise in the supply of skilled labor, while technological change leads to a rise in the demand for skilled labor. If there is a relative shortage of high-skilled labor, because technological progress

[16] Buiter, Rahbari, and Seydl (2015) makes a similar point.

[17] Taken literally, the rise in emerging market economies' share in world GDP would mechanically push up on desired saving by around 2 percent of world GDP. The effect is strongest since 2007. The pickup is largely driven by China.

races ahead of educational attainment, then inequality is likely to increase, as those with the sought-after skills will tend to see their earnings rise relative to the rest of the population. In this context, the recent rise in the cost of education, if not tackled by policy, could potentially represent a rise in inequality of opportunity that may limit educational attainment and hence increase income inequality, probably with a long lag. This mechanism could be further strengthened if the latest technological advances not only increase the demand for skilled workers but also replace low-skilled jobs—Frey and Osborne (2013) predict that 47 percent of U.S. employment may be subject to computerization over the next twenty years.

Other factors may also play a role. For example, further globalization could make "winner-takes-all markets" more common, raising the share of income accruing to the global "superstars." Piketty (2014) also suggests that if the growth rate of labor income declines as global growth falls back (g), but the rate of return on capital (r) is maintained at its historic rate, then inequality is likely to rise further. On balance, absent a major policy shift, we judge that labor income inequality is more likely to continue rising than to fall back in the years ahead, but the future path of inequality is very uncertain. Hence our treatment is cautious, assuming only a very gentle increase in inequality going forward.

4.3 The Emerging Market Saving Glut

Following the Asian crisis in 1998, many emerging markets significantly increased their foreign exchange reserves as a precautionary measure against the future risk of destabilizing capital outflows. In tandem, the era of high oil prices prompted an increase in saving among oil producers. Bernanke (2005) suggests that these forces represented a preference shift by governments (in Asia) and a shift in circumstances (for oil exporters) that were largely exogenous to the global system. These preference shifts resulted in an increase in desired saving in those countries. To the extent that this increase was not matched by a rise in desired investment, it led to a net increase in global saving.

On average, the current account surplus of Asian economies and oil exporters—indicative of the net amount of financial capital that those countries send abroad—has been 1 percent of world GDP since

Figure 8. Global Imbalances

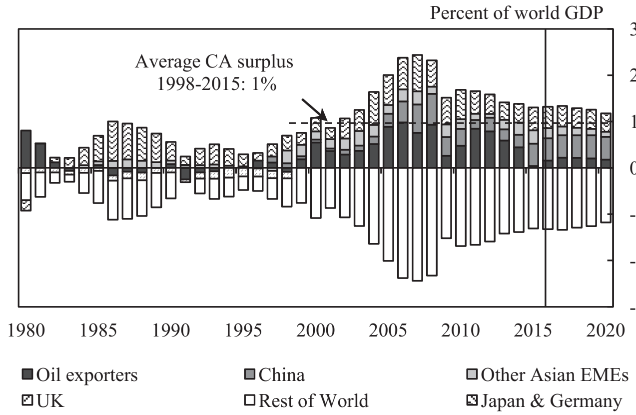

Sources: IMF and authors' calculations.

the late 1990s, around 1 pp higher than the roughly balanced current account pre-1998 (figure 8). Using the increase in emerging markets' current account surplus as a guide suggests that the desired saving schedule has shifted to the right by 1 pp as a result of the emerging market saving glut, which lowers the global real rate by round 25 bps. This is only around half of the effect of inequality, and a quarter of the effect of demographics.

Bernanke (2015b) discusses the future of the emerging market saving glut and concludes that the outlook is mixed. On the one hand, three factors suggest that these imbalances may have run their course: (i) some of the emerging market economies, particularly China, are rebalancing their economies away from exports toward domestic demand; (ii) the buffer stock of FX reserves that emerging markets hold is already sufficiently large, and the buildup of foreign currency reserves is slowing and in some cases now falling; and (iii) oil prices have fallen, so we might expect the excess savings from oil producers to decline further from pre-crisis peaks. On the other hand, there are also some new potential sources of the "saving glut." For example, Bernanke notes that Germany has the highest current account surplus in the world, and there is a concern that this will persist, exerting further downward pressure on global real rates. But to us it is unclear whether Germany's surplus will act as an additional force or has already been captured by other trends discussed in this paper, notably demographics. Overall, we think that the IMF forecast for global imbalances—as shown in figure 8—is a reasonable baseline forecast, which suggests a very gradual unwinding of the emerging market saving glut going forward.[18]

Taken together, shifts in desired saving linked to the three trends above can account for around 160 bps of the fall in global real rates since the 1980s. If this had been the whole story, we would have expected to see a steady rise in actual saving rates globally. But global saving and investment ratios have been remarkably stable over the past thirty years—as noted earlier. This suggests that the desired investment schedule has also shifted. Here we focus on three trends that could potentially explain such a shift: (i) the secular decline in the relative price of capital goods; (ii) a preference shift away from public investment projects; and (iii) an increase in the spread between the risk-free rate and the return on capital.

4.4 The Falling Relative Price of Capital Goods

Perhaps one of the most pervasive trends that may have affected desired investment expenditure is the 30 percent decline in the relative price of capital goods since the 1980s.[19] Cheaper capital means that a given investment project costs less to pursue, so investment volumes can be maintained by committing a smaller share of nominal GDP. But cheaper capital also incentivizes additional investment projects, given the lower cost. The overall impact on capital expenditure is the sum of the two effects—its sign depends on the elasticity of substitution between capital and labor.

If capital and labor are easily substitutable, a fall in the relative price of capital goods will induce a lot of additional investment

[18]Others present compelling alternative forecasts. For example, Speller, Thwaites, and Wright (2011) perform long-run simulations of gross capital flows out to 2050 and predict persistently high current account surpluses in emerging markets, with the Chinese current account shrinking but India's rising.

[19]IMF (2014b) has examined changes in the relative price of investment in the advanced economies since 1980. The Fund documents a downward trend in the relative price of investment that then levels off in the early twenty-first century. In explaining this movement, the Fund points to the work of Gordon (1990), who emphasizes the role of research and development that is embodied in cheaper, more efficient investment goods.

projects, potentially by enough to counter the effect of falling prices and hence maintain investment as a share of nominal GDP. But most empirical work points to the elasticity being smaller than one.[20] The IMF (2014b), for example, asserts that any increase in the volume of investment caused by a decline in the price of capital goods has been insufficient to offset the negative impact on real interest rates. Thwaites (2015) surveys the literature and arrives at a similar conclusion. An elasticity smaller than one means that a fall in the relative price of capital goods will tend to be associated with a shift of the investment schedule to the left (desired investment expenditure is lower for a given interest rate). To calibrate the size of this shift, we rely on the model developed in Thwaites (2015). A 30 percent decline in the relative price of investment lowers the steady-state nominal investment-to-GDP ratio by around 1 pp in that model. The fall in the relative price of capital goods also has an additional effect, which is to pivot the investment schedule (so that it becomes steeper), as any given amount of real investment now requires less of today's output to be sacrificed. In other words, the opportunity and financial cost of investment become a less important factor in making investment decisions. Desired investment becomes less sensitive to interest rates by roughly the same amount as the fall in the relative price of capital goods.[21]

The 1 pp shift in the investment schedule, together with a 30 percent drop in the elasticity of investment with respect to the interest rate, delivers around a 50 bps fall in the real rate in the saving-investment diagram. This is similar to the peak-to-trough fall in the interest rate along the transition path to the new steady state in Thwaites (2015).[22]

The future of capital goods prices is still being hotly debated, and no clear consensus has yet emerged. Eichengreen (2015) argues that further falls are not guaranteed:

[20] Piketty (2014) is a notable exception.

[21] More formally, in the formulation of the user cost of capital, the weighted average cost of capital enters multiplicatively with the relative price of capital goods.

[22] The paper argues that the initially positive response of the interest rate matches what happened in the real world in the 1970s, when real rates were increasing. Since we are trying to explain the fall in real rates since the 1980s, it is appropriate to compare the decline in the real interest rate since the peak.

> Evidently, R&D is not embodied more easily and fully in investment goods than consumption goods in all times and places. The presumption behind "the Baumol effect"—that consumption goods, and in particular that portion provided by the service sector, are difficult to mechanize and therefore become relatively more expensive over time—may not hold in the future as it has in the recent past. Even if the post-1980 decline in the relative price of investment goods is part of the explanation for the concurrent decline in real interest rates, there is no ruling out that it may be reversed in the future.

Others have also noted that the relative price of capital goods has stabilized more recently, taking this as evidence that the ICT revolution has run its course.

On the other side of this debate, researchers at the Federal Reserve (Byrne, Oliner, and Sichel 2013) believe that the price of ICT equipment has been persistently mismeasured. In their view, statisticians struggle to capture the higher capability of the latest technologies such as quad-core processors. Byrne et al.'s quality-adjusted estimates suggest that microprocessor prices continue to fall at a rate of around 30 percent a year. Furthermore, Thwaites (2015) argues that the effect on real rates can build for a long time even after capital goods prices stop falling.

Overall, it seems reasonable to assume some further contribution from the decline in the relative price of capital goods to lower rates, albeit at a diminished magnitude compared with the past.

4.5 Lower Public Investment

Public investment has been on a declining trend as a share of global GDP since the 1980s (see IMF 2014a). This could be because political views have become more polarized, thus making it difficult to agree upon and implement large-scale public investment projects. Or it may be because voters have shifted their preferences away from large governments. Either way, the result of this shift has been to lower the global investment-to-GDP ratio by around 1 pp between 1980 and 2007. Since 2007, public investment in emerging economies—particularly China—has accelerated rapidly, unwinding the long-term decline seen in the preceding decades. However, we think that much of the post-2007 pickup is a cyclical response to weakening demand during the global financial crisis. We therefore expect this to reverse and the downward secular trend to eventually reassert itself. Consequently, we think that lower public investment has shifted the desired investment curve to the left by around 1 pp, lowering real rates by around 20 bps—a relatively small effect.

An alternative interpretation of the recent movements is that higher public investment in emerging market economies is currently pushing up on the global real rate (relative to pre-crisis), and if that unwinds, the global real rate will fall further. The difference between these two explanations comes down to whether the shift away from public investment has already affected global equilibrium rates or whether this is still to come. In either case, given that the size of the effect is relatively small, this channel is not a major driver of movements in the global real rate.

4.6 Spread between the Risk-Free Rate and the Rate of Return on Capital

So far our analysis has abstracted from the fact that the interest rate that matters for firms' investment decisions is the rate of return on capital, not the risk-free rate. Strictly speaking, when analyzing desired investment, the rate of return on capital rather than the risk-free rate should appear on the vertical axis of the S-I diagram. This distinction would not be important if the spread between the risk-free rate and the return on capital had been constant over time—the desired investment schedule shown in this paper would represent a simple vertical transformation of the "correct" schedule. However, there is some evidence that the spread has risen over time, which has implications for desired levels of investment. A rise in the spread shifts the desired investment schedule vertically down—because in order to keep desired investment unchanged, the risk-free rate must fall by exactly the same amount as the spread has increased, all else equal. However, in general equilibrium all else is not equal, and a lower risk-free rates induces people to save less—suggesting that the eventual decline in the risk-free rate may be a little smaller than the rise in the spread.

The IMF constructs a weighted measure of the spread across these measures for the world as a whole (figure 9). This shows that the rate of return on capital has fallen since the early 1990s, but

Figure 9. The Global Risk-Free Rate and Rate of Return on Capital

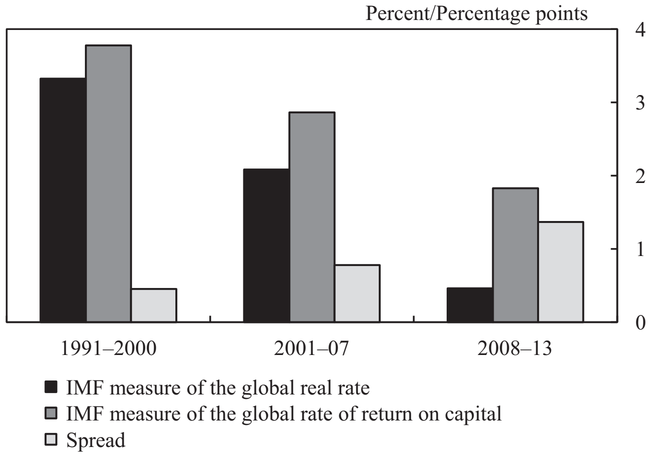

Sources: IMF (2014b) and authors' calculations.

not by as much as the risk-free rate—the spread has increased by around 100 bps. Market-by-market analysis supports this conclusion (Rachel and Smith 2015). In our framework, such a rise in the spread causes the real interest rate to drop by about 70 bps.

Together, the three investment trends highlighted above (the falling price of capital goods, lower public investment, and rising credit spreads) account for around 140 bps of the fall in real rates seen since the 1980s—a similar order of magnitude to the effect from the three saving trends highlighted earlier.

Figure 10 brings together all of these trends using the saving and investment framework outlined earlier. Our analysis suggests that the desired saving schedule has shifted out materially due to demographic forces (90 bps of the fall in real rates), higher inequality within countries (45 bps), and a preference shift towards higher saving by emerging market governments following the Asian crisis (25 bps). In addition, desired investment rates appear to have fallen as a result of the decline in the relative price of capital goods (accounting for 50 bps of the fall in real rates), a preference shift away from public investment projects (20 bps), and an increase in the

Figure 10. Quantifying Shifts in Desired Saving and Investment

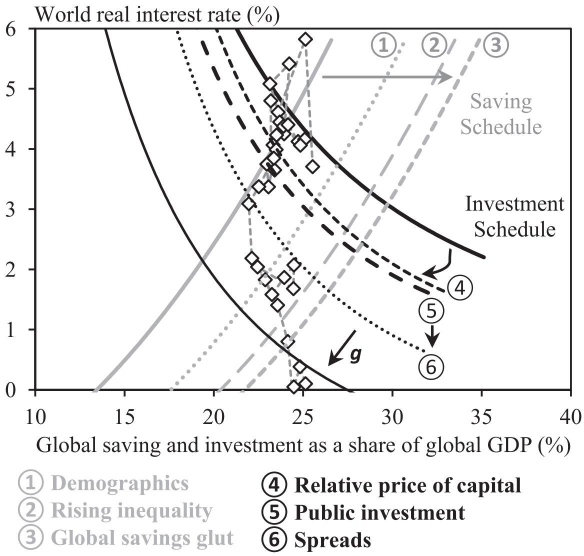

spread between the risk-free rate and the return on capital (70 bps). Together these effects can account for 300 bps of the fall in global real rates. We also include an illustrative shift of the desired investment schedule to account for weakening global growth prospects identified in the previous section (labeled "g" in the diagram).

This saving-investment framework provides a broad description of the relative sizes of the different forces at play. Taken at face value, shifts in preferences appear to explain around 300 bps of the decline in real rates since the 1980s, on top of the 100 bps explained by the deterioration in the outlook for trend growth seen more recently. In other words, we think we can account for most of the decline in global real rates using evidence independent of the decline itself.

Around 50 bps of the fall in real rates remains unaccounted for. This could reflect a number of factors, and each of these underscores the possible uncertainties that surround our point estimates. First, we might be missing certain secular trends from our analysis. For example, rising short-termism (Gutiérrez and Philippon 2016) or the decline in capital intensity of production (Summers 2013) could also be pushing down on real rates. Second, some of the trends we have quantified could be having bigger effects than we have estimated—for example, some studies that focus on an individual factor have found larger effects (e.g., for inequality and demographics). Third, the unexplained component could reflect global headwinds from the financial crisis, such as deleveraging or heightened risk aversion, which are temporarily pushing down on real rates. Fourth, the market measures of real rates we are using, which are derived from government bond yields, may be distorted. Post-crisis regulatory changes may have increased demand for safe government assets by financial institutions, while central bank quantitative easing (QE) has been temporarily boosting the demand for government bonds.

As mentioned earlier, significant uncertainty also surrounds the slopes of the saving and investment curves. The vertical pattern of the intersection between real interest rates and saving-investment ratios over the past thirty years may suggest that one or both of the schedules is steeper than we have assumed. While that is a possibility, we note that a near-vertical slope would imply that agents' behavior is invariant to changes in interest rates, and so would be at odds with empirical evidence suggesting that interest rates do affect behavior and the macroeconomy (Christiano, Eichenbaum, and Evans 1999; Ramey 2016). Moreover, it would likely lead to significant volatility in the observed interest rates, as small shifts would have a very large impact on the equilibrium real rate. This would be at odds with the relatively smooth downward trend in the long-term real rate over time. On the other hand, if both of the curves were significantly flatter than we assume, we would not have observed the large decline in the real interest rate, even in the presence of strong underlying secular trends: small adjustments in the interest rate would have realigned desired saving and investment. To check the robustness of our results, we vary the elasticities of the desired saving and investment curves in a number of ways. Table 2 reports four sensitivity tests compared with our benchmark calibration. The key point here is that while the slope of the curves affects the allocation of saving and investment trends in driving the change in real rates, when both sets of forces are viewed together they produce a broadly similar total impact on real rates across calibrations. Put differently, the conclusion that around 3 pp of the downward move in the real interest rate is due to shifts in desired saving and

Table 2. Sensitivity of the Results to Changing Elasticities of Saving and Investment Curves

Calibration	**Baseline**	**Saving More Sensitive to Interest Rates; Investment Less Sensitive**	**Saving Less Sensitive to Interest Rates; Investment More Responsive**	**Both Saving and Investment More Responsive (Less Steep)**	**Both Saving and Investment Less Responsive (Steeper)**
Elasticities	$\varepsilon_S = 0.5$; $\varepsilon_I = -0.7$	$\varepsilon_S = 1$; $\varepsilon_I = -0.35$	$\varepsilon_S = 0.25$; $\varepsilon_I = -1$	$\varepsilon_S = 1$; $\varepsilon_I = -1$	$\varepsilon_S = 0.25$; $\varepsilon_I = -0.35$
Shifts in Desired Saving	1.6 pp	1.7 pp	1.3 pp	1 pp	2.5 pp
Shifts in Desired Investment	1.4 pp	0.8 pp	1.6 pp	1 pp	1.7 pp
Total	3 pp	2.5 pp	2.9 pp	2 pp	4.2 pp

Note: ε_S and ε_I are estimates of the elasticities of saving and investment to the real interest rate, respectively.

Figure 11. Secular Drivers of Global Real Interest Rates

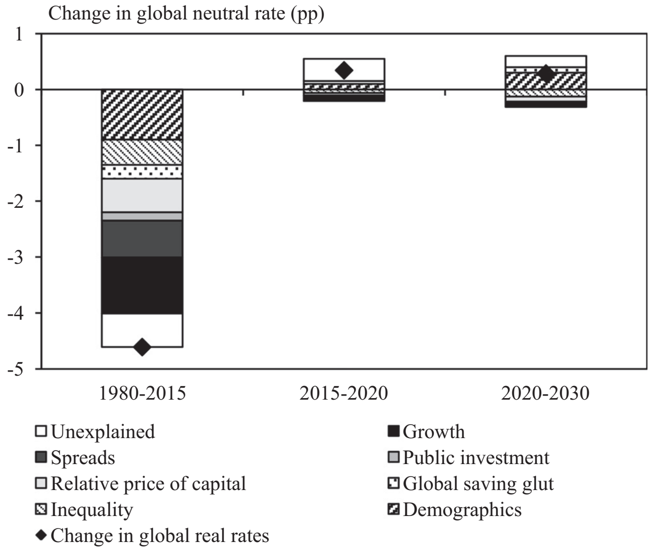

investment seems to be broadly robust to assumptions about the elasticities of the saving and investment schedules.

Clearly, the confidence interval around all of the above estimates is very wide. Nevertheless, quantifying the potential size of the impact from each of the different secular trends serves a useful purpose—it not only helps us explain movements in real rates over the past but also allows us to opine on how such trends (and hence real rates) are likely to evolve going forward. Figure 11 provides a summary of our findings about past movements in real rates, together with our best judgments about the direction of travel based on the discussions in this paper. The main message is that the trends we have analyzed are likely to persist at their current level: we do not predict a big further drag, or a rapid unwinding of any of these forces. Some are likely to drag a little further (global growth is set to decline further out, and we assume this will feed into slightly lower rates in anticipation; the relative price of capital is likely to continue to fall, albeit at a slower pace; and inequality may continue to rise), but this will be broadly offset by a rebound in other forces (particularly demographics). What happens to the unexplained component depends on what is driving it. In figure 11 we illustrate the effect of assuming that it is largely cyclical. Despite this, our predictions still imply that the global neutral rate will remain low, perhaps settling at around 1 percent in the medium to long term.

5. Conclusion

Since the 1980s, market measures of long-term risk-free real interest rates have declined by around 450 bps across both emerging and developed economies, with a discernible common trend suggesting that global factors are at work. The causes of the fall are likely numerous and diverse. In this paper we have attempted to quantify the impact that several secular forces could have had on real rates, by affecting global growth, desired saving, and desired investment. Although there is great uncertainty, our estimates suggest that these secular forces can account for about 400 bps of the 450 bps decline in the global long-term neutral rate seen since the 1980s. Moreover, most of these secular trends look likely to persist, suggesting that the global neutral rate may settle at around 1 percent over the medium to long run.

The policy implications of permanently low real interest rates are extensive. In the face of adverse shocks, central banks are likely to run up against the zero lower bound on nominal interest rates more often, requiring the use of unconventional policy instruments such as quantitative easing (QE). For large adverse shocks, fiscal policy may need to bear more of the burden of business-cycle management. Low rates may also fuel search-for-yield behavior, posing challenges for macroprudential and microprudential policymakers. More generally, the possibility of the global neutral rate remaining at persistently low levels should motivate a real debate across the policy spectrum on the best approach to stabilize the cycle.

The fact that the evolution of the global neutral rate remains highly relevant for policy and that our analysis has only briefly touched on a vast territory means there are many exciting areas for future research. One extension to the analysis in this paper would be to use a regional perspective to shed light on the fall in global real rates. As discussed in more detail in Rachel and Smith (2015), developed and emerging economies have exhibited different trends in saving and investment over time—future work could usefully explore the reasons for these differences to pinpoint the drivers of global real rates more precisely. Future analysis could also acknowledge the central role of the monetary and financial sector in the global economy of today, by looking at monetary trends that could have affected the global real rate, such as a structural change in global liquidity. This could complement the analysis of real trends highlighted in this paper. Further work could also look to extend our analysis further back in history, in an attempt to explain not only why real rates have fallen since the 1980s but also why they rose in the decades before the 1980s. Finally, development of structural models capable of handling secular, medium- to long-run trends may shed further light on the size of the effects highlighted here, and help clarify how the various effects interact—allowing more in-depth policy experiments and welfare analysis. To the extent that our headline prediction of a persistently low long-run neutral rate comes to pass, such research will remain highly relevant for years to come.

References

Baker, D., J. D. Delong, and P. R. Krugman. 2005. "Asset Returns and Economic Growth." *Brookings Papers on Economics Activity* 2005 (1): 289–330.

Bernanke, B. S. 2005. "The Global Saving Glut and the U.S. Current Account Deficit." Remarks at the Sandridge Lecture, Virginia Association of Economists, Richmond, Virginia, March 10.

———. 2015a. "Why Are Interest Rates So Low?" Brookings Blog Entry, April 1.

———. 2015b. "Why Are Interest Rates So Low? Part 3: The Global Saving Glut." Brookings Blog Entry, March 30.

Bomfim, A. 1997. "The Equilibrium Fed Funds Rate and the Indicator Properties of Term-Structure Spreads." *Economic Inquiry* 35 (4): 830–46.

Broadbent, B. 2014. "Monetary Policy, Asset Prices and Distribution." Speech given at the Society of Business Economists Annual Conference, October 23.

Brynjolfsson, E., and A. McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in Times of Brilliant Technologies.* W.W. Norton & Company.

Buiter, W., E. Rahbari, and J. Seydl. 2015. "Preventing Secular Stagnation from Becoming the 'New' Normal: Why the Time of One-Armed Economic Policy Is Gone." Citigroup Research.

Byrne, D M., S. D. Oliner, and D. E. Sichel. 2013. "Is the Information Technology Revolution Over?" Finance and Economics Discussion Series No. 2013-36, Board of Governors of the Federal Reserve System.

Caballero, R. J., and E. Farhi. 2013. "A Model of the Safe Asset Mechanism (SAM): Safety Traps and Economic Policy." NBER Working Paper No. 18737.

Carvalho, C., A. Ferrero, and F. Nechio. 2016. "Demographics and Real Interest Rates: Inspecting the Mechanism." Working Paper No. 5, Federal Reserve Bank of San Francisco.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" *Handbook of Macroeconomics*, Vol. 1, Part A, 65–148 (chapter 2). Elsevier.

Crossley, T., and C. O'Dea. 2010. "The Wealth and Saving of UK Families on the Eve of the Crisis." Report, Institute for Fiscal Studies.

Cynamon, B. Z., and S. M. Fazzari. 2016. "Inequality, the Great Recession, and Slow Recovery." *Cambridge Journal of Economics* 40 (2): 373–99.

DeFina, R. H. 1984. "The Link between Savings and Interest Rates: A Key Element in the Tax Policy Debate." *Business Review* (Federal Reserve Bank of Philadelphia) (November/December): 15–21.

Dynan, K. E., J. Skinner, and S. P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112 (2): 397–444.

Eggertsson, G. B., and N. R. Mehrotra. 2014. "A Model of Secular Stagnation." NBER Working Paper No. 20574.

Eichengreen, B. 2015. "Secular Stagnation: The Long View." NBER Working Paper No. 20836.

Ellis, C., and S. Price. 2003. "UK Business Investment: Long-Run Elasticities and Short-Run Dynamics." Bank of England Working Paper No. 196.

Fischer, S. 2016. "Why Are Interest Rates So Low? Causes and Implications." Speech at the Economic Club of New York, New York City, October 17.

Frey, C. B., and M. A. Osborne. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" Paper presented at Oxford University's Machines and Employment Workshop, March 6.

Gilchrist, S., and E. Zakrajsek. 2007. "Investment and the Cost of Capital: New Evidence from the Corporate Bond Market." NBER Working Paper No. 13174.

Goldin, C., and L. F. Katz. 2007. "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005." NBER Working Paper No. 12984.

Gordon, R. J. 1990. *The Measurement of Durable Goods Prices.* University of Chicago Press and National Bureau of Economic Research.

———. 2014a. "The Demise of US Economic Growth: Restatement, Rebuttal and Reflections." NBER Working Paper No. 19895.

———. 2014b. "US Economic Growth Is Over: The Short Run Meets the Long Run." In *Think Tank 20: Growth, Convergence and Income Distribution: The Road from the Brisbane G-20 Summit*, 185–92. Washington, DC: Brookings Institution.

Guiso, L., A. K. Kashyap, F. Panetta, and D. Terlizzese. 2002. "How Interest Sensitive is Investment? Very (when the data are well measured)." Conference Paper, European Central Bank.

Gutiérrez, G., and T. Philippon. 2016. "Investment-less Growth: An Empirical Investigation." NBER Working Paper No. 22897.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West. 2015. "The Equilibrium Real Funds Rate: Past, Present and Future." Paper written for the U.S. Monetary Policy Forum, New York City, February 27.

Hansen, A. 1934. "Capital Goods and the Restoration of Purchasing Power." *Proceedings of the Academy of Political Science* 16 (1): 11–19.

———. 1938. "Economic Progress and Declining Population Growth." Presidential address delivered at the 51st annual meeting of the American Economic Association, Detroit, Michigan, December 28.

Havranek, T., R. Horvath, Z. Irsova, and M. Rusnak. 2015. "Cross-Country Heterogeneity in Intertemporal Substitution." *Journal of International Economics* 96 (1): 100–118.

International Monetary Fund. 2014a. "Is It Time for an Infrastructure Push? The Macroeconomic Effects of Public Investment." In *World Economic Outlook: Legacies, Clouds, Uncertainties* (October), 91–112 (chapter 3). Washington, DC: International Monetary Fund.

———. 2014b. "Perspectives on Global Real Interest Rates." In *World Economic Outlook: Recovery Strengthens, Remains Uneven* (April), 75–114 (chapter 3). Washington, DC: International Monetary Fund.

Jorgenson, D. W., and K. M. Vu. 2010. "Potential Growth of the World Economy." *Journal of Policy Modeling* 32 (5): 615–31.

King, M., and D. Low. 2014. "Measuring the 'World' Real Interest Rate." NBER Working Paper No. 19887.

Kumhof, M., R. Ranciere, and P. Winant. 2015. "Inequality, Leverage, and Crises." *American Economic Review* 105 (3): 1217–45.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology.* The Penguin Group.

Kuznets, S. 1960. "Population Change and Aggregate Output." In *Demographic and Economic Change in Developed Countries*, 324–39. Columbia University Press and NBER.

Laubach, T., and J. C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85 (4): 1063–70.

Lee, R., and A. Mason, eds. 2011. *Population Aging and the Generational Economy: Global Perspective.* Edward Elgar.

Mendes, R. R. 2014. "The Neutral Rate of Interest in Canada." Discussion Paper No. 2014-5, Bank of Canada.

Mishkin, F. S. 1984. "Are Real Interest Rates Equal Across Countries? An Empirical Investigation of International Parity Conditions." *Journal of Finance* 39 (5): 1345–57.

Piketty, T. 2014. *Capital in the Twenty-First Century.* Harvard University Press.

Piketty, T., and E. Saez. 2014. "Inequality in the Long Run." *Science* 344 (6186): 838–43.

Rachel, L., and T. D. Smith. 2015. "Secular Drivers of the Global Real Interest Rates." Staff Working Paper No. 571, Bank of England.

Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics,* Vol. 2, ed. J. B. Taylor and H. Uhlig, 71–162 (chapter 2). Elsevier.

Saez, E., and G. Zucman. 2014. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." NBER Working Paper No. 20625.

Speller, W., G. Thwaites, and M. Wright. 2011. "The Future of International Capital Flows." Financial Stability Paper No. 12, Bank of England.

Summers, L. H. 2013. "Economic Possibilities for our Children." Remarks delivered at the Fifth Annual Martin Feldstein Lecture, held at the NBER Summer Institute, Cambridge, Massachusetts, July 24. Published in the *NBER Reporter* 2013 No. 4.

———. 2014a. "Bold Reform Is the Only Answer to Secular Stagnation." *Financial Times*, September 7.

———. 2014b. "U.S. Economic Prospects: Secular Stagnation, Hysteresis and the Zero Lower Bound." *Business Economics* 49 (2): 65–73.

Thwaites, G. 2015. "Why Are Real Interest Rates So Low? Secular Stagnation and the Relative Price of Investment Goods." Discussion Paper, Centre for Macroeconomics.