Discussion of "Are Low Real Interest Rates Here to Stay?"*

Pierre-Olivier Gourinchas
University of California at Berkeley

1. The Trillion Dollar Question

The decline in global real interest rates since the early 1980s is nothing short of spectacular. In the United States, the entire yield curve on government securities has shifted down by more than 10 percent between 1983 and 2013. A similar pattern, although perhaps not as pronounced, has been documented in most advanced and many emerging economies. According to the authors, long-term global real interest rates have declined by more than 450 basis points (bps) since the early 1990s. Coupled with low and stable worldwide inflation, this overwhelmingly suggests a decline in the global "neutral" or "natural" rate of interest, colloquially known as r^*, of a similar magnitude.[1]

It is difficult to overstate the macroeconomic importance of such a phenomenon. Among other things, a decline in r^* of this magnitude seriously limits the ability of monetary authorities to stabilize aggregate demand, as policy rates cannot be lowered below the "effective" lower bound. It also poses specific challenges to financial stability. For instance, excessive cuts in the policy rate can lower banks' net interest margin, with adverse effects on the health of the banking sector. A low-r^* world also creates powerful incentives for financial intermediaries to engineer "pseudo" safe assets, and for investors to load up excessively on higher-risk/higher-yield assets.

*This discussion was prepared for the November 2016 IJCB conference held at the Federal Reserve Bank of San Francisco. Part of this discussion builds upon joint work with Hélène Rey.

[1]The paper presents a useful discussion of the subtle differences between the global neutral rate and country-specific r^*, but these subtleties are not relevant for the purpose of this discussion since the purpose of the paper is on the determinants of the global real rate, and the authors argue that this global neutral rate is a long-run anchor for r^*.

The trillion-dollar question, then, is to understand the sources of this decline. Not surprisingly, given the importance of the topic, potential explanations abound—from a decline in long-run productivity growth to demographic forces due to an aging global population; from rising income inequality to a decline in the relative price of investment goods; or from increased foreign exchange reserve accumulation by emerging market central banks to a shift in investors' desired portfolio weights towards safe assets, among others—so much so that one might be tempted to conclude that the decline in global real rates is, if anything, *over-explained*.

Sifting carefully through these different explanations and sorting out which are relevant and which are not is no easy task. Yet it is the one that the authors, somewhat heroically in my view, set for themselves in this paper. To do so, they propose a very detailed and nuanced quantitative assessment of the importance of each of the large number of factors mentioned above, and some others.

The final decomposition looks both sensible and interesting: of the 450 bps decline in long-term real yields, the authors claim to account for 400 bps—100 bps due to a decline in future trend growth, 160 bps due to increased desired savings, and 140 bps from decreased desired investment. Two determinants appear particularly relevant in the authors' decomposition: the impact of demographic forces on desired savings (90 bps) and that of a rising spread between risk-free rates and the return to capital, dampening desired investment (70 bps).

I am very sympathetic to that final point, in part because it is very much in line with my own work on the topic! For instance, in Caballero, Farhi, and Gourinchas (2017), we propose a decomposition of that same spread into a risk premium component, and a component due to increased markups ("rents") or to increased capital-augmenting technological progress ("automation"). That paper concluded that while increased rents or automation cannot be ruled out, they alone cannot account for the increased spread between returns to capital and safe real interest rates. Instead, a substantial part of the increase in that spread reflects increased compensation for risk, indicating a substantial increase in the demand for, or a significant decline in the supply of, safe assets relative to risky assets.

However, while my general views are very sympathetic with the general conclusions of the paper, I am much less convinced by the methodology that the authors have adopted in their paper. The next section reviews my main concerns in that respect. Next, I present an altogether different approach, based on the intertemporal budget constraint, and my recent work with Hélène Rey (Gourinchas and Rey 2017). That framework, using long-run data, finds also an important role for the financial cycle, with real interest rates remaining low for an extended period of time following global financial crises.

2. Accounting for Changes in the Natural Interest Rate

At its core, an economic model is a mapping from shocks $\bar{\varepsilon} = \{\varepsilon_i\}$ to observables $Y = \{y_j\}$. That mapping is controlled by a number of structural parameters Θ and can be summarized as $Y = f(\bar{\varepsilon}; \Theta)$.[2]

The first step consists in estimating the fundamental parameters $\hat{\Theta}$. An impulse response can then be obtained by setting some component i of ε to 1 and all other shocks to 0. An accounting decomposition can be obtained once the fundamental shocks $\hat{\bar{\varepsilon}} = \{\hat{\varepsilon}_i\}$ are also estimated. The contribution of shock i can then be obtained by constructing $Y_i = f(\hat{\bar{\varepsilon}}_i; \bar{\Theta})$, where $\hat{\bar{\varepsilon}}_i = \{\hat{\varepsilon}_i, 0\}$ ignores all other shocks beside shock i.[3]

While the general principles are well understood, their application is often difficult. For instance, proper identification of the structural parameters Θ is often difficult to achieve. Precise estimation of the structural shocks can also be difficult.

In the context of this paper, such a decomposition exercise would start with a model incorporating all the potential sources of decline in the neutral rate of interest listed above: productivity growth, demographic forces, income inequality, the relative price of investment goods, foreign exchange reserve accumulation, etc. Notice that some of these sources, such as income inequality or the relative price of investment goods, are themselves endogenous to, inter alia, the structure of the tax system; the degree of competition in factor

[2] That definition is quite generic and does not impose that the model be static: y_j can include current, lagged, and even future observations of macroeconomic variable j, and similarly ε_i can include current, lagged, and future values of shock i.

[3] With non-linearities, these decompositions do not necessarily add up to 100 percent since the full response also incorporates interaction terms.

markets; productivity in the human accumulation sector; unobserved abilities; or productivity in the investment good sector, etc. The full model would incorporate a large number of potential shocks and a large number of structural parameters. Unfortunately, the complexities and identification challenges of estimating such a model lie well beyond our technical abilities.

Instead, this paper cobbles together estimates from two different approaches. The first approach is a Euler equation approach:

$$r = \frac{1}{\sigma} g + \rho, \quad (1)$$

where r is the real interest rate, g is the growth rate of technology, σ is the intertemporal elasticity of substitution (IES), and ρ is the rate of time preference. This equation obtains from the optimal intertemporal allocation of consumption by households. The interpretation is straightforward: faster productivity growth means more output (and consumption) tomorrow relative to today. Households want to smooth consumption over time, hence they would like to borrow and consume more today. This pushes up the real interest rate, the more so the lower is the IES (low σ). For a given IES, this maps changes in g into changes in r.

The second approach is an $S - I$ diagram. It consists of identifying shifts in desired savings S and desired investment I arising from the different mechanisms mentioned above. For given estimates of the interest rate elasticity of the saving and investment curves, it maps the changes in desired savings and investment into changes in the equilibrium real interest rate.

Unfortunately, while this approach is more tractable than the full-blown estimation of a structural model, it leads to a number of conceptual difficulties.

To start with, many of the "channels" discussed can have complex effects on the real interest rate, so the issue of identification remains unsettled. Consider for instance the link between the dependency ratio (defined as the fraction of the population not of working age) and savings. The empirical evidence documented in the paper indicates a weak negative relationship between the saving rates and dependency ratios in the cross-section, which is stable over time (figure 6). The authors assume that this relationship also holds in the aggregate time series. Since the dependency ratio has declined globally, they conclude that this contributed to increased desired aggregate savings. Given an assumed elasticity of investment of -0.7, this accounts for 90 bps of the 160 bps declined in interest rates due to increased savings. This is one of the largest single contributors to the decline in real interest rates. If true, this would suggest that future increases in the dependency ratio would *raise* real interest rates as aggregate savings decline. Yet, it is well known that some countries with a relatively high dependency ratio due to rapid population aging, such as Japan or Germany, are also among the countries with the highest saving rate.

The answer may well lie in the fact that the dependency ratio can vary when the proportion of either young dependents or old dependents varies. But the implications for aggregate savings may be vastly different. Historically, countries with low dependency ratios are countries that successfully made their demographic transition. Fewer children also mean higher savings, since in many developing countries children were effectively a means of saving for old age. As an illustration, China's one-child policy may well have contributed to an increase in Chinese savings. As the proportion of young dependents has decreased globally, savings have increased. But the implications may well be very different if the change in the dependency ratio comes from the fraction of old dependents. An increase in population aging means more retirees relative to the working-age population. It is likely to be associated with an *increase*, not a decrease, in savings. Once this heterogeneity is taken into account, it is not so clear that the global dependency ratio is a good summary statistic for the desired savings shift due to demographic forces. The authors' conclusion from this reduced-form exercise seems relatively fragile.

Assuming that desired shifts in savings and investment are properly measured, translating these into shifts in equilibrium real interest rates requires reliable estimates of the savings and investment interest rate elasticities. The authors assume an elasticity of savings of 0.5, but mention a range from 0.03 to 1.8. From a theoretical and empirical point of view, it is well known that the relationship between savings and interest rates can be quite fickle. It is in fact quite telling that the most recent paper cited by the authors dates back to 1983: the literature seems to have concluded that the elasticity of aggregate savings to the real interest rate is not a well-defined parameter and one should perhaps avoid trying to estimate it. Yet it is a critical component of the exercise. The interest rate elasticity of investment is also imprecisely estimated, between –0.5 and –1, although at least both the theoretical and empirical literature agree that it should be negative.

Finally, it is not clear how the $S-I$ framework is an alternative to the Euler equation approach. After all, the Ramsey-Cass-Koopman model embeds an $S-I$ diagram too. A decline in productivity growth lowers the marginal product of capital, which reduces the desired capital stock and reduces investment. In equilibrium the real interest rate needs to decline to reduce saving, so that $S = I$ obtains. In other words, the $S-I$ diagram is precisely what delivers $r = g/\sigma + \rho$. It is not clear to me whether we are not double-counting when we are considering the effect of, e.g., demographics or rising inequality on trend productivity via the Euler equation and the $S - I$ diagram. One of the advantages of the structural model $Y = f(\bar{\varepsilon}, \Theta)$ is precisely that it avoids such a double-counting.

Finally, the Euler equation is a weak peg to hang the empirical estimates on. A large and abundant empirical literature has documented very weak support for the aggregate Euler equation and the absence of a strong relationship between real interest rates and growth.

The bottom line is that, while the structural approach may be infeasible, the approach followed in this paper, cobbling and adding together many reduced-form estimates of "desired" shifts in savings and investment, combined with an Euler equation, does not provide solid empirical estimates of the relative contributions of the various channels. Unfortunately, no matter how complex the problem is, we cannot afford to avoid using some structural approach. In doing so, the trade-off is between parsimony and tractability. The next section presents the results from such a framework, borrowed from Gourinchas and Rey (2017).

3. An Alternative Framework to Understand Secular Movements in r

In Gourinchas and Rey (2017), we propose an accounting framework based on the movements in the ratio of total consumption expenditures C to total wealth W over long periods of time. Figure 1 plots this ratio for a world that comprises the United States, the United

Figure 1. The Global Consumption-Wealth Ratio

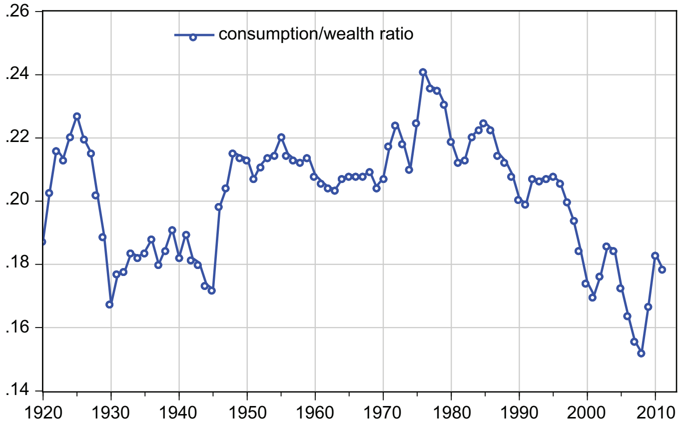

Source: Jordà, Schularick, and Taylor (2016), Piketty and Zucman (2014).

Note: The figure reports the ratio of aggregate annual private consumption expenditures to total private wealth (land, housing, financial assets) for the United States, the United Kingdom, Germany, and France.

Kingdom, Germany, and France between 1920 and 2011. Wealth consists of total private wealth, including land, housing, and financial assets as estimated by Piketty and Zucman (2014). The ratio, which can be interpreted as an average propensity to consume out of wealth, exhibits low-frequency fluctuation, between a low of 0.155 in 2009 and a high of 0.24 in 1975.

Under the mild assumption that this ratio is stationary, an assumption consistent with most theories of consumption, we can write the following decomposition:

$$\begin{aligned}\ln(C_t/W_t) \equiv cw_t &= \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t r_{t+s} + \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t rp_{t+s} - \sum_{s=0}^{\infty} \rho^s \mathbb{E}_t g^C_{t+s} \\ &= cw^{rf}_t + cw^{rp}_t + cw^c_t. \qquad (2)\end{aligned}$$

In this equation, the log consumption-to-wealth ratio cw_t is the sum of three components: the present discounted sum of future risk-free rates r_{t+s}, the present discounted value of future risk premia

Figure 2. Decomposing *cw*

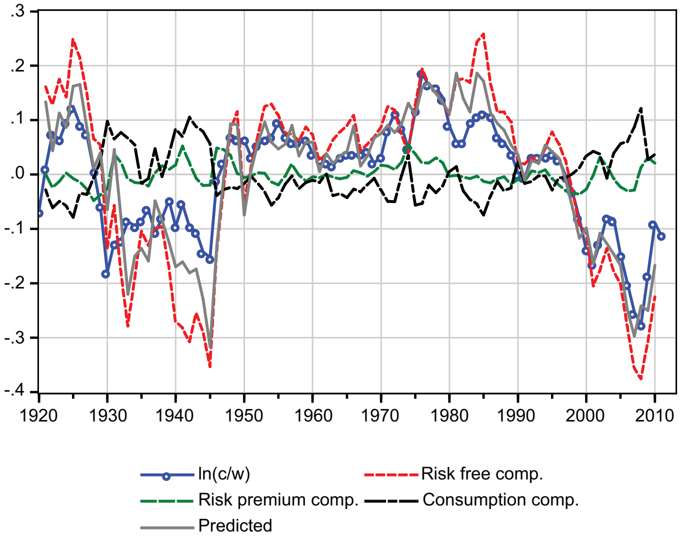

Note: The figure decomposes $\ln(C/W)$ into a risk-free component (cw^{rf}), an excess return component (cw^{rp}), and a consumption growth component (cw^{c}).

rp_{t+s}, and the present discounted sum of future aggregate consumption growth rates g^C_{t+s}. The discount rate ρ is a constant that depends on the steady-state value of C/W. The interpretation of this formula is quite straightforward: if C/W is stationary, a low C/W value today predicts a higher value in the future. The increase in C/W must come from either an increase in the numerator, i.e., high future consumption growth g^C_{t+s}, or a decline in the denominator, i.e., low returns on wealth, which implies either low future risk-free rates r_{t+s} or low future risk premia rp_{t+s} or both.

Using a vector autoregression approach, Gourinchas and Rey (2017) estimate the three components on the right-hand side of this equation. Figure 2 reports the estimated components, together with the overall fit of the regression. Because equation (2) is the approximation of an accounting identity (the global budget constraint), we expect the overall fit to be high, as is indeed the case. More importantly, the figure reveals that the risk-free component cw^{rf} accounts for most of the variations in cw while other components (future consumption growth or future risk premia) play relatively minor roles. In other words, a low consumption-to-wealth ratio today indicates that the present discounted value of future risk-free rates is below average. Over long periods of time, because of the effective lower bound, the real interest rate is an upper bound on the natural rate r^*. It follows that a low value of cw is associated with low future r^*. In Gourinchas and Rey (2017), we discuss how this result suggests that the slowdown in productivity growth or demographic forces are unlikely to be the main drivers of the decline in r^*. The basic insight is that low future productivity growth or a slowdown in population growth would both reduce the future growth rate of aggregate consumption growth g^C_{t+s}. According to equation (2), this would tend to *raise* cw today, not depress it, unless the equilibrium response of real interest rates to a growth slowdown, i.e., the intertemporal elasticity of substitution $1/\sigma$ in equation (1), is sufficiently low. In that case, cw^{rf} and cw^c would be negatively correlated and their ratio, an estimate of σ, would be implausibly low.

The upshot of the exercise is that either we live in a world of very low IES, so that global interest rates are extremely responsive to changes in productivity growth (more than one for one), or the movements in r originate elsewhere. Our favorite hypothesis is that of financial boom-bust cycles casting a long shadow on future real rates. In particular, the two episodes of persistently low cw seem to occur in the run-up to and the aftermath of global financial cycles: a Great Depression episode from from 1925 to 1945, and the global financial crisis from 2000 onwards. In both cases, the decline in cw arises first from a rapid run-up in asset and housing prices that increases wealth (the denominator) that is not matched by a corresponding increase in consumption. The financial crisis, when it occurs, should have a corrective effect, restoring cw back to its trend. However, we observe instead that cw remains depressed. One explanation is that the financial collapse triggers increased demand for safe assets and economy-wide deleveraging dynamics: as households, corporates, and governments all simultaneously try to reduce their borrowing levels, safe interest rates collapse. Our findings also suggest that though risk premia seem to increase, as expected if the demand for safe assets increases, this is not sufficient to offset the decline in risk-free rates, and the overall expected return on wealth remains depressed post-crisis.

Figure 3. Predicting Global Risk-Free Rates

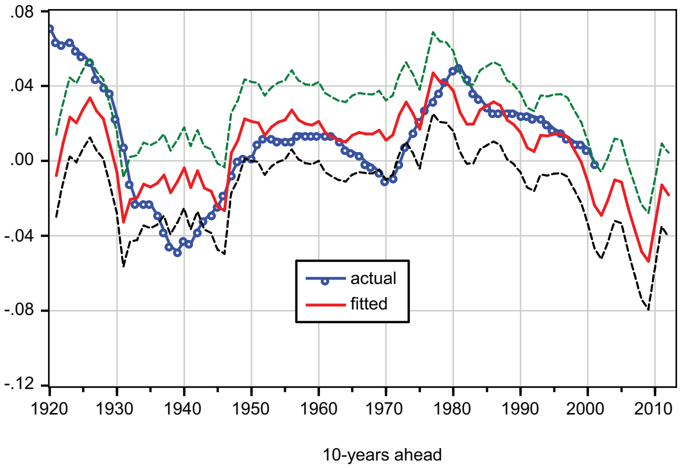

Note: The figure forecasts the ten-year average future short risk-free rate using $\ln(C/W)$. The graph includes two standard deviation bands.

Finally, our decomposition allows us to answer directly the question posed by the authors: Are low real interest rates here to stay? Figure 3 reports the predicted values based on a regression of a ten-year average of the real risk-free rate on the initial value of cw. The line with circles reports this average future interest rate between 1920 and 2001 (since our data cover the period 1920–2011). The solid line reports the estimated future risk-free interest rate, with two standard deviation confidence bands. The in-sample fit is very high. Extrapolating past 2001 until 2011, it reveals that global real risk-free rates are expected to remain low for an extended period of time. As of the last data point in our sample, 2011, the average short-term real risk-free rate over the subsequent ten years, 2011–21, is expected to be –2 percent.

4. Conclusion

The paper presents an ambitious attempt to account quantitatively for the different economic forces behind the recent decline in real interest rates. This is a difficult empirical exercise. Without a formal structural framework, identification is weak, and it is difficult to isolate cleanly the effect of various forces. Yet the exercise remains useful as a starting point for more structural explorations. I have presented the results from one such exercise, based on Gourinchas and Rey (2017). It provides a complementary set of results and emphasizes the historical role of deleveraging dynamics in depressing real interest rates.

References

Caballero, R. J., E. Farhi, and P.-O. Gourinchas. 2017. "Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares." *American Economic Review* 107 (5): 614–20.

Gourinchas, P.-O., and H. Rey. 2017. "Real Interest Rates, Imbalances and the Curse of Regional Safe Asset Providers at the Zero Lower Bound." In *The Future of the International Monetary and Financial Architecture*. Proceedings of the ECB Forum on Central Banking, sponsored by the European Central Bank, held in Sintra, Portugal, June 27–29, 2016.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2017. "Macrofinancial History and the New Business Cycle Facts." *NBER Macroeconomics Annual 2016*, ed. M. Eichenbaum and J. A. Parker, 213–63. National Bureau of Economic Research.

Piketty, T., and G. Zucman. 2014. "Wealth and Inheritance in the Long Run." CEPR Discussion Paper No. 10072 (July).