Comment on Rachel (2025)

Adrien Auclert*

February 2026

Monetary policy makers care about a lot of numbers, but there is one number they arguably care about more than any other, and that number is r^*. Adding r^* to their inflation target gives them the nominal interest rate they should be setting to hit that inflation target in the long run; comparing current estimates of r^* to the current real interest rate tells them how tight or loose current monetary policy is. Fiscal policy makers care a lot about r^* too, since it determines the long-run sustainability of a given fiscal trajectory, and how urgent any fiscal consolidation might be.

This excellent and timely paper by Lukasz Rachel gives us historical estimates of r^* and scenarios for the future. Those will be very welcome by monetary and fiscal policy makers alike. Even more importantly, in my opinion, the paper gives us a framework for thinking about the determinants of r^*, both qualitative and quantitative. The framework is of enormous value because we know from an earlier quantitative literature that forces such as productivity growth, markups, automation, population aging, inequality, government debt, taxes, and the generosity of social security should affect r^*, but it is not always easy a priori to understand which force acts in which direction, or the extent to which, say, a one percentage point increase in the debt-to-GDP ratio should raise r^*.

The paper is also a timely update on Lukasz's earlier Brookings paper with Larry Summers (Rachel and Summers, 2019). That paper argued, based on its estimates of low and declining r^* at the time, that "the industrialized world, taken as a whole, [was] currently—and for the foreseeable future [would] remain—highly prone to secular stagnation".[1] In the intervening time, long bond yields have riven by almost two percentage points in the U.S., leading commentators including Summers himself to argue that we have entered a "new normal" where secular stagnation is no longer a concern.[2] One of the benefits of the paper is that it uses the very model that Summers had used to argue in favor of secular stagnation, updated to the most recent macroeconomic data, to evaluate these kinds of arguments. Ironically, the paper does not find a strong case for rising r^*: its central scenario is a continued decline, although r^* could increase significantly if "several upside risks crystallize all at once".

*Stanford University, CEPR and NBER. Email: aauclert@stanford.edu.

[1] Once r^* turns sufficiently negative, the central bank can no longer keep output at potential even by setting nominal interest rates at their effective lower bound. This situation can result in a steady state with below-potential output known as secular stagnation (see for instance Eggertsson et al., 2019).

[2] See for instance "Larry Summers isn't worried about secular stagnation anymore", Slow Boring Substack, Milan Singh, April 18, 2024, https://www.slowboring.com/p/larry-summers-isnt-worried-about

Figure 1: Supply and demand determinants of r^*

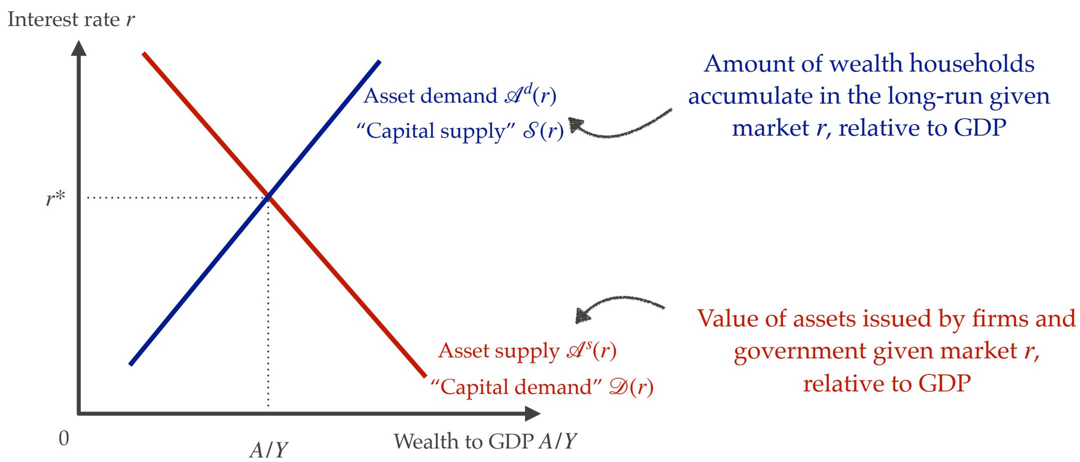

I think this is a fantastic paper: timely, clear, well-written, with important findings. What I love the most about it is that it uses an approach that is clearly right: conceptualizing r^* as the price that equalizes asset supply and asset demand in the long-run. This kind of framework allows us to use classic price theory to get at both the qualitative and the quantitative effects of various forces on r^*. It turns our attention to matching not just price (r^*) but also quantity (in this case, the wealth-to-GDP ratio); it also highlights the sufficient statistics that are relevant for making quantitative predictions. Of course, I am a little biased in saying this, since this framework is very similar to the one I used in a recent paper on this topic (Auclert et al., 2025). Our two papers implement this framework in very different ways, but our conclusions end up being quite similar!

My discussion will highlight these key points of agreement, as well as some sources of disagreement, and provide some thoughts on where the literature could go forward from here.

1 Supply and demand determinants of long-run r^*

Many of the leading macroeconomic models of r^* feature a similar structure: in the long run, the interest rate r^* is the one that equilibrates the supply and the demand for assets. Figure 1 visualizes what this means. Given the market interest rate (r), households decide the amount of wealth they would like to accumulate in the long-run, relative to GDP: I call this relationship the asset demand curve, $\mathcal{A}^d(r)$. This curve is typically upward sloping (higher interest rates make households want to postpone consumption, which they do by accumulating wealth) but it can be downward sloping, especially at low interest rates, where the income effect (higher interest rates mean retirees can afford to save less) can dominate the substitution effect. Critically, "interesting" models of r^* have an imperfectly elastic asset demand curve. This requirement rules out the presence of any infinitely-lived agent with standard macro preferences $\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$, since any such agent accumulates infinite long-run wealth as soon as the interest rate r exceeds $\frac{1}{\beta}(1+g)^{\sigma} - 1$, where g is the growth rate of income, making the economy-wide asset demand curve $\mathcal{A}^d(r)$ infinitely elastic at

Figure 2: Equilibrium effect of asset supply and demand shifts

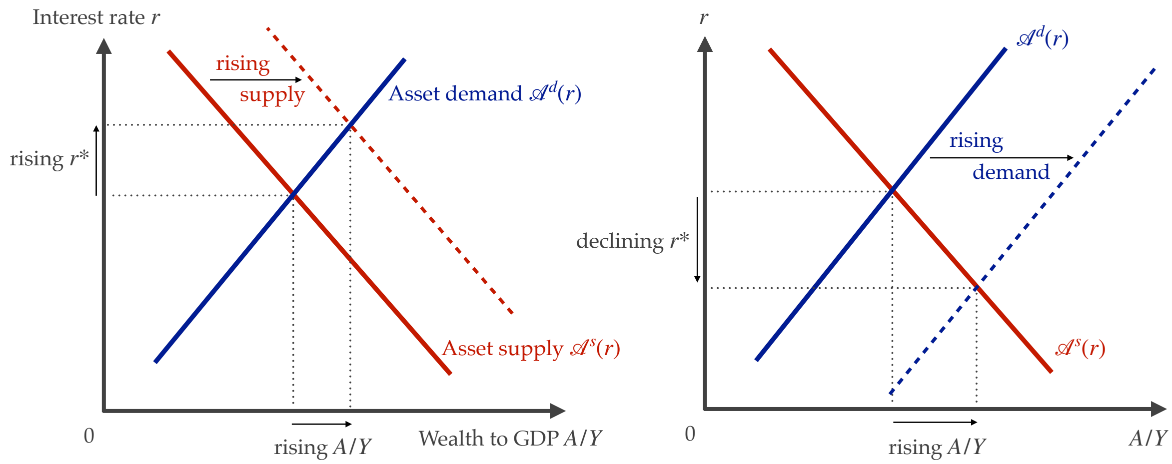

that rate.

Similarly, given the market interest rate r, the value of assets issued by firms and the government achieves a certain level relative to GDP: I call this relationship the asset supply curve, $\mathcal{A}^s(r)$. This curve is typically downward sloping, since neoclassical firms choose to accumulate less capital when the user cost of capital is higher; and also because stock market valuations are lower when interest rates are higher, and because the government may decide to reduce its stock of government in response to rising market rates.

One semantical point of departure with the Rachel (2025) paper is that he calls "capital supply" the asset demand curve, and "capital demand" the asset supply curve. This convention makes sense in simple models, where asset supply is made up entirely of the capital stock, and the relationship between asset supply and the interest rate comes from firms's demand for capital given that interest rate; in turn, households must accumulate wealth to provide that capital, such that their asset demand is the capital supply. However, I think the convention becomes less tenable in models, such as ours, where asset supply is also made up of stock market values and government bonds. So I will follow the asset demand/asset supply convention—a convention dating back to at least to Caballero (2006) in the literature—rather the author's for the rest of this discussion.[3]

The reason why this framework is such a nice way to understand the determinants of r^* is that it allows us to immediately make sense of the direction through which potentially complex forces act. For instance, higher government debt-to-GDP ratios, but also automation/rising capital intensity or lower risk premia (which boost firm valuations) all increase asset supply: we should expect them to raise the long-run level of wealth-to-GDP as well as long-run r^* (figure 2, left panel). By contrast, population aging or a decline in the generosity of social security (both of which make people save more), rising inequality, or an increasing foreign demand for assets all raise asset

[3] Another point in favor of the Rachel (2025) convention is that capital demand is downward sloping r, while capital supply is upward sloping in r, as in standard price theory. This is convenient, but also a little confusing, since asset prices actually move inversely to the interest rate r.

demand, so these forces raise the long-run level of wealth-to-GDP but lower long-run r^* (figure 2, right panel). The effect of some forces become clear this way: rising markups, for instance, have very large positive effects on valuations (rising supply), which offsets any effect from reduced capital accumulation; therefore, at least to the extent that the corresponding profits are capitalized in the stock market, markups should raise r^*. Some forces act in more complex ways, affecting both asset demand and supply. Take rising productivity growth: if cohorts perceive higher life-cycle growth, they will borrow more (or save less, i.e. falling demand), at the same time, firm valuations will get a boost (rising supply). The net effect is a rise in r^*, but an ambiguous effect on the level of wealth-to-GDP. And so on. Virtually any force can be understood in this way.

Given this framework, it is clear that any understanding of the past evolution or the future direction of r^* relies on correctly identifying the shifters of asset demand and supply. The paper proposes a fully structural approach to do this.

2 This paper: a structural approach

Specifically, the paper constructs the asset supply and asset demand schedules $\mathcal{A}^s(r)$ and $\mathcal{A}^d(r)$ by fully specifying the primitives underlying both and then solving for the value of each of these schedules at each possible value of r. The model is simple enough that $\mathcal{A}^s(r)$ and $\mathcal{A}^d(r)$ have an almost closed-form solution, making this procedure numerically straightforward.

On the asset supply side, the model features standard neoclassical production, markups resulting from monopolistic competition, and government debt. This results in a simple functional form for $\mathcal{A}^s(r)$ as a function of primitives such as capital intensity α, depreciation of capital δ, markups φ, and the government debt-to-GDP ratio. The paper proposes simple strategies to back out these primitive parameters from the data. One could quibble with the details of how this is done, but for the most part it is a fairly straightforward procedure. The only real stance to take is the level and the pace of increase of markups over time, which is subject to considerable debate in the literature; given the structure of the model, capital intensity then follows from $1-\alpha=\varphi\times\frac{wL}{Y}$ where $\frac{wL}{Y}$ is the labor share in the data.

On the asset demand side, the model is what one could call a "double heart attack" model. As discussed above, one must step outside of the world of standard infinitely-lived consumers to generate an inelastic asset demand curve. Overlapping generation models with finite lifetimes are one way to achieve this, but they can get quite complicated when modeling many periods of life, with agents solving a very different problem at each age, as they go through their life cycle with evolving income prospects and mortality risk. Blanchard (1985) introduced the single heart attack model as an alternative: agents face a constant probability of having a heart attack in a period, after which they die; a constant fraction of agents is born each period to replace the dying. This setting has tractable aggregation because agents have linear consumption rules, but it also builds in life-cycle effects, with agents accumulating wealth over time as they age. However, agents have effectively only working period of life, which doesn't allow for the possibility of a retirement

Figure 3: Asset demand wins

(a) Rachel (2025)

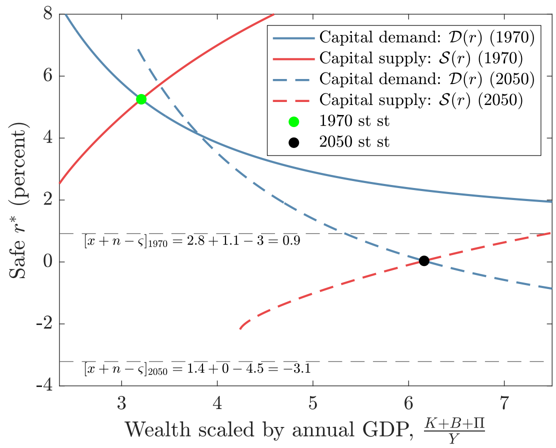

(b) Auclert et al. (2025)

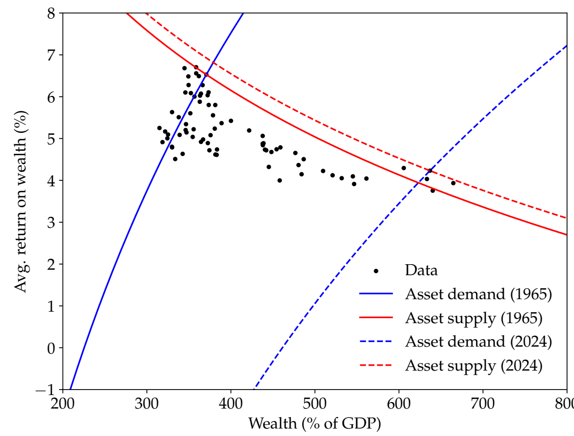

phase. To remedy this and think about issues such as social security, Gertler (1999) introduced the *double* heart attack model: after a first heart attack, agents move to a retirement phase of life, in which they face another (possibly higher) constant probability of death. This setting remains highly tractable, with nice aggregation properties, which is why it is heavily used in the macro-demographics literature to derive analytical results (see for instance Carvalho et al., 2016).

The resulting model has very few parameters, most of which (the growth rate of newborns, the length of working life, the length of retirement, the average income of workers and retirees) can be taken directly from the data. As in standard life-cycle models, asset accumulation responds to cuts in social security income or to increases in the length of time in retirement; one drawback, however, is that it is not possible to separate the effect of a shift in an increase in the retirement age (a policy parameter) from an actual increase in life expectancy.

The main finding from the paper is that, between 1970 and 2050, factors that increase asset demand/capital supply at given interest rates, such as longer retirement, falling population growth, and declining income growth, are much more potent than factors that increase asset supply/-capital demand, such as automation, rising government debt, and rising markups. This is nicely visualized in figure 6 of the paper, which I reproduce here as my figure 3(a). The outcome is a decline in r^*, by about 530 basis points (from 5.3% in 1970 to 0% today). This baseline specification refutes Summers's "new normal" hypothesis: in the transition between 1970 and 2050, interest rates fall continuously, and they do not get back to their pre-2008 level in the decades to come.

The paper's Figure 8 breaks down the overall shifts into contributions from each of the factors. It is clear there that longer retirement, slowing population growth and falling growth are the three major drivers of the shift in asset demand. Since the first two are slow-moving demographic drivers that cannot be reversed, this finding implies that reversing the prediction of falling r^* would require a large increase in productivity growth, perhaps once that could result from the development of artificial intelligence.

3 An alternative reduced-form approach

The structural approach followed by the paper relies on trusting that the model provides the "true" nonlinear mapping between macroeconomic fundamentals and asset demand and supply. In practice, however, there is a lot of uncertainty about the right underlying model. The double heart attack model is convenient but also highly counterfactual at the microeconomic level, implying for instance the simultaneous presence of very young retirees (who had a heart attack in their 20s) and extremely old workers (who have yet to get a heart attack at age 200). At that level, the model—just like most macro models—is therefore clearly unrealistic. Nevertheless, it *may* give similar predictions as a more complex model that would have more micro-level realism. How can we evaluate whether it does?

One way to answer this question is to take a more reduced-form approach, deriving sufficient statistics for predictions for r^* that would be valid in a class of models, including more realistic models of household behavior. Indeed, the price theory framework I spelled out in section 1 is useful not just to understand qualitative directions, but quantitative magnitudes as well, at least to a first order. To see this, write $\mathcal{A}^s(r,\eta)$ for the asset supply curve and $\mathcal{A}^d(r,\eta)$ for the asset demand curve with macro shifter η, then the relationship $\mathcal{A}^s(r,\eta) = \mathcal{A}^d(r,\eta)$ implies that the change in the level of the equilibrium interest rate dr^* and in the equilibrium log asset-to-GDP ratio $d\log\mathcal{A}$ resulting from a small change in macro fundamentals $d\eta$ is given by

$$dr^* = \frac{\frac{\partial \log \mathcal{A}^s}{\partial \eta} d\eta - \frac{\partial \log \mathcal{A}^d}{\partial \eta} d\eta}{\epsilon^s + \epsilon^d} \tag{1}$$

$$d\log\mathcal{A} = \frac{\epsilon^d \frac{\partial \log \mathcal{A}^s}{\partial \eta} d\eta + \epsilon^s \frac{\partial \log \mathcal{A}^d}{\partial \eta} d\eta}{\epsilon^s + \epsilon^d} \tag{2}$$

where the semielasticities of asset supply and demand are defined as

$$\epsilon^s \equiv -\frac{\partial \log \mathcal{A}^s}{\partial r} \quad \epsilon^d \equiv \frac{\partial \log \mathcal{A}^d}{\partial r}$$

with all derivatives evaluated at the original steady state. Equations (1) and (2) are helpful because they show that, to the first order, the equilibrium interest rate and quantity changes in the asset market following a change in macro fundamentals only require knowledge of the semielasticities of asset supply and demand ϵ^s and ϵ^d and the extent to which macro fundamentals shift the asset supply and the asset demand curves and given interest rates, $\frac{\partial \log \mathcal{A}^s}{\partial \eta} d\eta$ and $\frac{\partial \log \mathcal{A}^d}{\partial \eta} d\eta$. Similarly, by linearity of the first-order approximation, the effect of any combination of shifters $\eta_1, \eta_2 \ldots \eta_N$ is then the sum of each effect in isolation, which gives us a simple way to decompose the total effect into a sum of components. Any difference between model predictions for given fundamental shift $d\eta$ can then be traced back to differences in the elasticities of supply and demand ϵ^s, ϵ^d or to differences in asset supply and demand shifters $\frac{\partial \log \mathcal{A}^s}{\partial \eta} d\eta$, $\frac{\partial \log \mathcal{A}^d}{\partial \eta} d\eta$.

If we have data on the historical evolution of both r^* and $\log \mathcal{A}$, and if we know ϵ^s and ϵ^d, we can back out the shifters of asset demand and supply that rationalize this evolution. Auclert et al. (2025) propose a procedure to do this.[4] Using the procedure and data from that paper, I back out the 1965 and 2024 asset supply and demand curves and visualize them in figure 3b). Notice the striking visual analogy between this figure and that from the Rachel (2025) paper in the left panel. The slopes of asset supply and demand are similar: $\epsilon^s = 13$ and $\epsilon^d = 12$ in the Rachel paper, $\epsilon^s = 20$ and $\epsilon^d = 8$ in my paper. In particular the sum $\epsilon^s + \epsilon^d$ is about the same, implying that the papers will make similar predictions for interest rates given shifters, even though they will have different implications for how much these shifters translate into equilibrium quantities. Both papers agree that asset demand shifts were much larger than asset supply shifts over the 1965-2024 period, which Auclert et al. (2025) summarize as "demand won the historical race between asset demand and supply".

On the other hand, some differences appear too: most notably, the 1970-2050 decline in the equilibrium interest rate is much more pronounced than the 1965-2024 decline in my paper. This is partly because we treat r^* very differently— r^* is an unobservable factor in the Rachel (2025) paper while I argue we can observe it in my paper—but, given that we have similar elasticities, must also boil down to much larger underlying shifters of asset demand and supply.[5]

Backing out underlying drivers. Further progress with the reduced-form approach can be made by noticing that first-order shifters of asset supply and demand can sometimes be straightforward to calculate. For instance, a rise in government debt by 1% of GDP raises asset supply by $\frac{1}{A/Y}$%, where A/Y is the asset to GDP ratio: this gives us $\frac{\partial \log \mathcal{A}^s}{\partial \eta} d\eta$ for such a shift. A little more involved, but just as simple to calculate, a decline in fertility affects the age structure of the population, and its effect on asset demand can then be studied by integrating the life cycle profile of assets using the new age distribution, in a shift-share type of calculation (see Auclert et al., 2021). Auclert et al. (2025) provide a full set of sufficient statistics for shifters of $\frac{\partial \log \mathcal{A}^d}{\partial \eta} d\eta$ in standard life-cycle models, and implement them using micro data on household behavior, together with information on the evolution of inequality and demographics.

4 Comparing results from the two approaches

Table 1 shows the change in r^* implied by our two models and its decomposition into various forces. I backed out the numbers on the left column from Figure 15 in the Rachel (2025) paper; this includes an interaction term because the model is nonlinear so the overall isn't just the sum of the

[4] Most models agree with the formulation of the supply side, and ϵ^s is then easy to calculate from model primitives and a few macroeconomic observables. Auclert et al. (2025) provide a formula for ϵ^d that hold in standard life-cycle models, as a function of the elasticity of intertemporal substitution in consumption and observable age profiles.

[5] The difference in endpoints does not make much of a difference given that the transitions in the Rachel (2025) model see very little action between 2025 and 2050.

basis point change in r^*	Rachel (2025)	Auclert et al. (2025)
Time period	1970-2050	1965-2024
Overall	-530	-247
Productivity growth	-215	-103
Demographics	-275	-110
Inequality	—	-92
Global saving glut	-25	-72
Risk premium	-160	9
Government debt	145	44
Markups	105	97
Social security and taxes	130	22
Interaction terms	-200	—

Table 1: Comparison of the effects on r^* of various forces

individual effects. The right column is from Auclert et al. (2025) over the 1970-2024 range, here the interaction term is zero by construction.

The two decompositions should of course not be expected to line up exactly. The time periods are slightly mismatched. The Rachel (2025) model is one of all advanced economies, while the Auclert et al. (2025) model is for the U.S. only: this matters, for instance, when evaluating the magnitude of the increase in government debt or the net foreign asset position. Nevertheless, the signs of almost all forces are aligned—a reflection of the fact that the underlying framework is essentially the same. The main notable exception is the risk premium, which lowers r^* by 160bps in the Rachel (2025) model and raises it by 9bps in the Auclert et al. (2025) model. This is due to a disagreement over the evolution of the risk premium itself over this time period: Rachel (2025) assumes that this risk premium rises over the period, following Moll et al. (2022), while the data in Auclert et al. (2025) suggests a decline in the risk premium instead.

For other forces, the signs are the same, and the magnitudes tend to align as well. As I mentioned above, since the sum $\epsilon^s + \epsilon^d$ is comparable in both models, those differences can largely be traced back to differences in the magnitude of the underlying asset supply or demand shift. Sometimes, this is just due to a mere difference in scope: for instance, the increase in net foreign assets or government debt depends on whether one takes Rachel (2025)'s advanced-economy perspective or Auclert et al. (2025)'s U.S. perspective.

The papers have more fundamental disagreements about the effects of productivity and demographics: these have much bigger effects in the Rachel (2025) than in the Auclert et al. (2025) model. For productivity, this comes from the fact that the Rachel (2025) paper loads productivity growth on time effects: falling productivity growth implies that incomes rise less steeply for each individual in the economy, and this results in much higher savings. In the Auclert et al. (2025) model, productivity growth loads on cohort effects instead: lower productivity growth means new cohorts enter the labor market with lower earnings, but the life-cycle profiles of income are unaffected. There are therefore no life-cycle savings responses to lower productivity growth, with the only remaining effect coming from savings composition across cohorts. The bottom line is

Figure 4: Micro implications of the double heart attack model

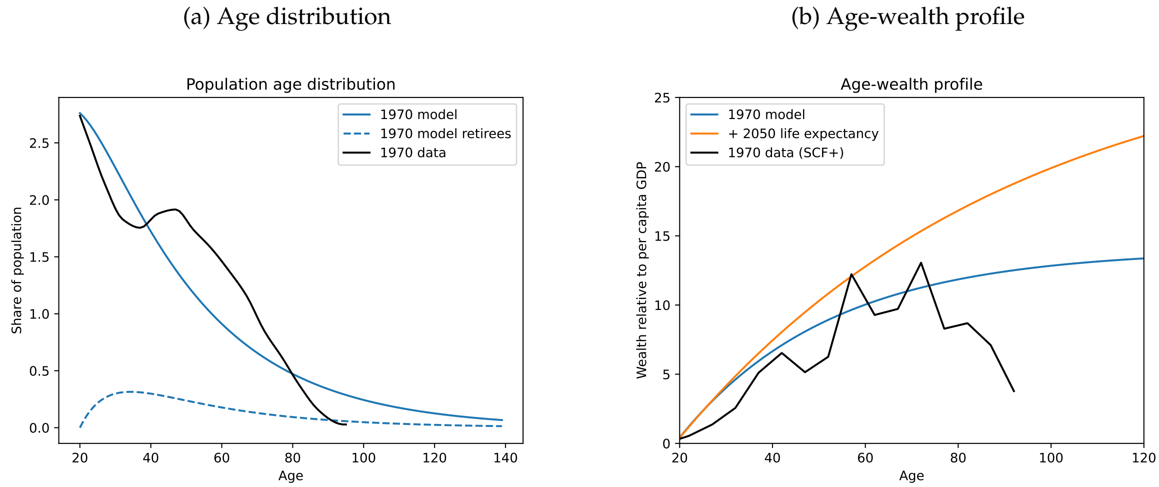

that, as pointed out by Blanchard (2023), the r^* effect of productivity growth depends a lot on whether this growth loads on time or cohort effects. Future literature should try to adjudicate this important question.

For demographics, my calculations in the next section suggest that the double heart attack model tends to overstate the savings responses to rising longevity, for reasons I will explain.

Another noteworthy factor is that interaction terms end up mattering a lot in the Rachel (2025) model, explaining almost the entire difference in aggregate predictions in the change in r^* between it and Auclert et al. (2025). There are two ways to view this result. One is that nonlinearities may matter quite a lot when thinking about these issues. The other one is that we should have less confidence about any result coming from nonlinear interactions, since they cannot be explained by taking an alternative reduced-form approach (at least to first order). By construction, this type of nonlinearity will be quite model specific, and future work can also help determine whether the linearities robustly go in the direction of the Rachel (2025) paper.

This discussion just illustrates one way in which results from papers on r^* can be compared to each other. While the point estimates differ, my view is that there is already quite a bit of convergence, and future work can narrow that gap further.

5 Aggregate savings implication of the double heart attack model

One particular asset demand shifter that I suspect the Rachel (2025) model of overpredicting is that of the effect of rising mortality on life-cycle saving. Figure 4 present my simulations of the micro part of the model, using the author's parameters, for the population age distribution in the left panel and the age-wealth profile in the right panel. As discussed above, the model features a lot of young retirees as well as extremely old retirees and workers. However, in terms of the fit to the overall population age distribution in 1970, it does a decent job: in both the model and the data, high recent fertility means the population is disproportionately young. Aggregating across all agents, we obtain an age-wealth profile of assets which is relatively aligned with the data from the Survey of Consumer Finance in 1970 (right panel). However, considering the 2050 steady state with much higher life expectancy implies a very large savings response in the model, with the (now more numerous) 120+ year olds accumulating an enormous amount of assets. This effect seems implausible. If the model embedded the type of bequest forces that are generally seen as necessary to match life-cycle asset profiles, these effects would be dramatically mitigated, as shown by Auclert et al. (2021).

6 Our fiscal future

Going forward, the model consider a number of very interesting alternative scenarios. I think most of these are plausible, but I would highlight one that seems quite optimistic. The model assumes that the level of advanced economy government debt to GDP rises very modestly between 2024 and 2050, from 82% to 90%. Contrast this with the CBO projections, with U.S. debt exploding out to the forecast horizon, with no fiscal consolidation in sight. This effect is directly related to the aging of the population: expenditure on programs such as Medicare and Medicaid grow with the elderly population. This creates a channel through which demographic change affects not just asset demand, but asset supply as well. In Auclert et al. (2025), I argue that this is actually the main reason to think that r^* may go up significantly in the future.

7 Conclusion

I have stressed some of the benefits of the reduced-form approach over the Rachel (2025) structural approach. But the structural approach has advantages too: notably, the tractability of the model allows the author to easily solve for transition dynamics rather than treating the system as if it were in steady state at all time, and to incorporate a more realistic treatment of expectations than the usual perfect-foresight assumption. Ultimately, reduced-form and structural approaches are highly complementary and will help us refine our views of the evolution of r^* going forward, providing those all-important numbers to monetary and fiscal policy makers.

References

Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie, "Demographics, Wealth, and Global Imbalances in the Twenty-First Century," Working Paper 29161, National Bureau of Economic Research, August 2021.

_ , _ , Matthew Rognlie, and Ludwig Straub, "The Race Between Asset Supply and Asset Demand," *Proceedings of the Jackson Hole Economic Policy Symposium*, November 2025.

Blanchard, Olivier J., "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, April 1985, *93* (2), 223–247.

_ , *Fiscal Policy Under Low Interest Rates*, MIT press, 2023.

Caballero, Ricardo J., "On the Macroeconomics of Asset Shortages," in Andreas Beyer and Lucrezia Reichlin, eds., *The Role of Money: Money and Monetary Policy in the Twenty-First Century*, Fourth European Central Banking Conference November 2006.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio, "Demographics and Real Interest Rates: Inspecting the Mechanism," *European Economic Review*, September 2016, *88*, 208–226.

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins, "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics*, January 2019, *11* (1), 1–48.

Gertler, Mark, "Government Debt and Social Security in a Life-Cycle Economy," *Carnegie-Rochester Conference Series on Public Policy*, June 1999, *50*, 61–110.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 2022, *90* (6), 2645–2683.

Rachel, Lukasz, "What Next for r*? A Capital Market Equilibrium Perspective on the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 2025, *2025.*

_ **and Lawrence H. Summers**, "On Secular Stagnation in the Industrialized World," *Brookings Papers on Economic Activity*, 2019, pp. 1–54.