Commentary: What Explains the Decline in r^*? Rising Income Inequality versus Demographic Shifts

Fatih Guvenen

In this thought-provoking paper, Atif Mian, Ludwig Straub and Amir Sufi address an important question: What explains the decline in the natural rate of interest, r^*, since the 1980s in the United States and around the world? As is well known, the decline was not limited to r^* but was observed across a broad range of interest rates, from short term T-bill rates to 10-year government bond rates. These low rates both limit the ability of central banks to respond to a future recession and raise concerns about a series of unintended effects that can create new distortions or worsen inequality. To name a few examples, low rates typically boost stock prices and housing values (disproportionately owned by high-wealth households) thereby amplifying wealth inequality; they reduce the income flow from retirement savings for millions of retirees (whose portfolios skew toward safer assets) thereby worsening income inequality; and they lower the profitability of banks, potentially leading to more risk taking in search of higher returns, among others. These wide-ranging potential effects all highlight the urgency of the question that this paper seeks to answer. And the fact that many of these side effects concern inequality makes the paper a great fit with the theme of this year's Jackson Hole symposium on "Macroeconomic Policy in an Uneven Economy."

The present paper studies two potential explanations for the decline in r^*. The first hypothesis ("rising inequality") proceeds in three steps. First, the well-documented rise in U.S. income inequality since the 1970s increased the income share of top earners (e.g., top 10%) at the expense of the bottom 90%. Second, this income shift, combined with the higher savings rate of high-income households, leads to an increase in the savings of top earners both in absolute terms and as a share of national income. Third, this higher savings demand by top earners, in turn, puts downward pressure on interest rates. This paper is an empirical analysis of the first two steps, and another paper by the same authors (Mian et al. 2021) proposes a model to address the last step.

The second hypothesis ("baby boomers' transition") also has three steps. First, the baby boom cohorts (born between 1945-64) went through two important phases of life during this time period: they entered the middle-age phase (45-64) starting in the 1990s and began to retire starting in the 2010s. Second, if savings rates vary sufficiently by age and the income share of each age group changes (in the right direction) during this time, savings inequality by age rises. And third, to the extent that this change raises the savings demand of middle-age savers and/or lowers the borrowing demand of the young, this can lower the market clearing interest rate.

The main results of the paper are summarized in two shift-share decompositions (Tables 1 and 2). I will discuss the first hypothesis in greater detail, so let me first review the findings regarding the second hypothesis (baby boomers) and get it out of the way. The main takeaway from Table 2 is that there is no clear age pattern that stands out from the shift-share decomposition. In particular, while the savings rate is higher for middle-age households (ages 45-64) than the young (18-44) and the retirees (65+), this rate has declined over time, narrowing the gap in savings rates. So, even though income shares have shifted from young toward the middle-age group, the canceling effects lead to a muted change in savings of each group. The authors also present a several other cuts of the data, which all reinforce the conclusion that the baby boomers' transition did not lead to a substantial change in the composition of savings across age groups. Overall, I find this analysis to be fairly persuasive and a valuable contribution to the literature.

Turning to the first hypothesis (Table 1), the decomposition reveals clearer patterns. First, the savings rate of top (10%) earners was high and relatively than stable over this period, which, combined with the rise in their income share, led to a large rise in the savings of this group (by about 720 billion in 2019 U.S. dollars or 3.3% of national income). As for the bottom 90%, not only their income share fell—to mirror the rise for the top 10%—but their savings rate also fell significantly (by 9.4 percentage points for households below median income), resulting in a large decline in the savings of this group (by about 3.7% of national income). To be clear, I am using the word "savings" here to include negative rates, and in fact the bottom half of U.S. households saw their savings rate fall from 2.6% to –6.8%, showing a significant rise in borrowing during this period. I will return to this point when I discuss the link from savings to r^* in a moment. Finally, putting the two pieces together, the decomposition implies that aggregate savings actually declined since the 1970s, consistent with what U.S. aggregate data shows (Chart 1).

So, overall, the shift-share decomposition shows a sizable widening in savings dispersion across income groups in the last 40 years. This result is both plausible and consistent with the well-documented rise in household debt over this period, from mortgage and credit card debt to student loans.

This brings us to the last step of the rising inequality hypothesis—that the rise in savings inequality leads to a decline in r^*. As I mentioned earlier, this link is not established in this paper but in Mian et al. (2021), so in a sense, it may be considered out of the scope of my discussion. At the same time, without this link established, the paper does not answer the question stated in the title, so this step is an integral part of the authors' thesis about r^*. Therefore, in the rest of my comments, I will discuss the determination of r^* in equilibrium and the mechanism proposed by the authors.

I will not attempt to describe the details of the model proposed in Mian et al. (2021) because I cannot do justice to it in a limited space, and delving into its details would take me far out of the scope of this discussion. However, let me say that I really like this paper and the

Chart 1
U.S. Aggregate Savings Rate and U.S. Wage Inequality, 1970–2019

A. U.S. Savings Rates

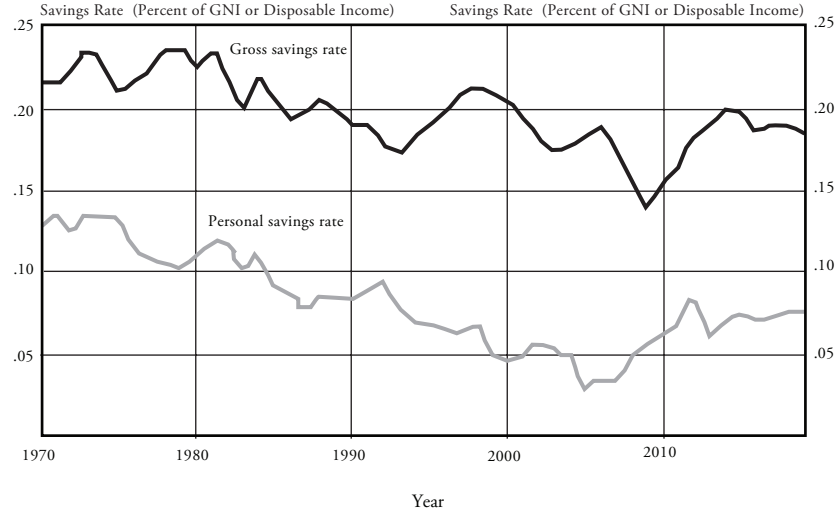

B. Inequality: Above and Below Median

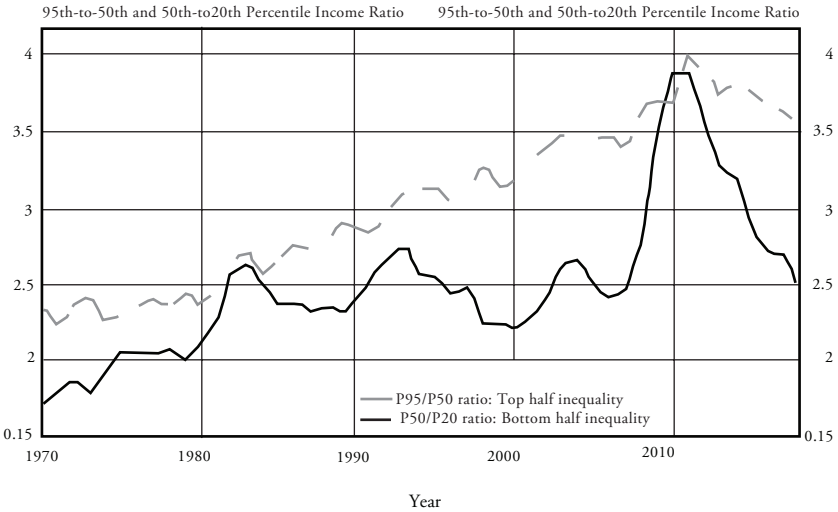

Sources: Personal savings rate is from the U.S. Bureau of Economic Analysis (PSAVERT), downloaded from Federal Reserve Bank of St. Louis' FRED database. Gross savings rate is from the OECD National Accounts files. Data on wage inequality is from Heathcote et al. (2020), who compute them from the U.S. Current Population Survey.

model the authors propose, which delivers a novel economic mechanism. So, without going into too much detail, let me briefly describe the mechanism that delivers a lower r^*.

The model is a deterministic perpetual-youth framework with two types of agents who trade a single risk-free asset (borrowing/saving) with each other subject to a debt constraint. There is a fixed endowment of real assets ("trees") that yield an exogenous income stream, so there is no saving at the aggregate level. The two agents are endowed with vastly different amounts of the real asset, generating large income inequality.

The key feature that the authors add is a preference for wealth—it enters the utility function—which can be justified by a warm-glow bequest motive. While this is a fairly common ingredient in life cycle models, the paper adds a twist: if wealth is a luxury good in this specification, wealthier agents have a higher marginal propensity to save out of lifetime income and the saving supply schedule slopes *downward.* So, when income inequality rises, the rich want to save a larger fraction of the extra income than the rise in the demand for borrowing by the poor. Since there is no saving at the aggregate level, the interest rate has to fall to clear the market.

There is quite a bit in this mechanism that I find plausible: there is certainly empirical evidence supporting a higher savings rate for the very rich (and the authors cite some of them), and interpreting the two groups as the top earners and the rest, it is true that the latter group is a net debtor in financial markets (i.e., excluding housing). That said, to focus on this new mechanism, the paper abstracts from some features that have been central to the heterogeneous-agent macro/inequality literature since the 1990s. Those features activate a different set of mechanisms that push the interest rate up when inequality rises. I view the two sets of mechanisms as complementing each other, and the relative strength of each depends on some key empirical details, which I discuss next.

I will begin with a few comments on savings rates, since they are central to the discussion. The first key fact in my view is the decline in aggregate savings rates during this time, which we saw in the shift share decomposition. To get a sense about the magnitudes involved,

in the top panel of Chart 1, I plot two measures of aggregate savings rates based on aggregate data, which provides a more accurate measurement than micro survey data can. Both measures of savings rates—U.S. gross savings as a fraction of national income and personal savings as a fraction disposable income—show a steady decline since the late 1970s, with a slight reversal after the Great Recession. The former measure includes the government sector whereas the latter only includes households. Using the same two time periods as in the paper (1970-80 and 1995-2019), the two measures fell by 3.8 percentage points (pp) and 6.0 pp, respectively. These are very large figures: $830 billion and $944 billion, respectively, using 2019 data for the denominators.

This matters for two reasons. First, it implies that the decline in r^* cannot be explained easily by "too much" savings (demand) in the economy. Whereas aggregate demand for savings increases in Mian et al. (2021), and interest rates fall to keep it at zero, we see a declining trend in savings in the data. Moreover, both savings rates start to recover after the Great Recession, which is also when the U.S. wage inequality started to decline after rising for decades (bottom panel of Chart 1). So, at least since the 1970s, U.S. wage inequality and aggregate savings have moved in *opposite* directions. I don't view this necessarily as definitive evidence against the rising inequality hypothesis, but I think it clearly suggests that any downward pressure on r^* from rising inequality has to come from the equilibrium effects of *compositional* changes in savings across income groups, while the aggregate savings rate itself declines.

My second takeaway from declining aggregate savings is that it means the rising borrowing of the bottom 90% is likely to be just as important for understanding the effects on r^* as the rise in the savings of the top 10%. In other words, an equally valid way to summarize the trends studied in this paper would be to say: the large rise in wage inequality reduced the income share of 90% of U.S. households, which in turn led to a significant fall in their savings, driving down the aggregate savings rate, despite a rise in the savings of the top 10%. So, I think understanding the factors that affect the saving/borrowing demand of the bottom 90% in face of rising inequality is critical for understanding what the effect on r^* will be.

The heterogeneous-agent macro literature emphasizes the importance of income uncertainty in driving borrowing and saving decisions to smooth consumption. A key idea is that the marginal propensity to save out of "current resources" rises with the level of resources. Current resources is income plus wealth but the latter is adjusted for borrowing capacity, so someone with zero wealth but, say, $20,000 in unused borrowing capacity will have income plus $20,000 in resources. In this setup, households' response to an income fall depends on three main factors: how persistent the income fall is, how much resources they have on hand, and how strong their preference is for consumption smoothing (loosely speaking, risk aversion).[1]

To explain how a rise in inequality affects the savings decision, it is convenient to divide the full range of current resources into three: very low, low-to-medium, and very high (the sizes of each depending on specific parameterization). Households in the lowest range behave completely myopically: they spend every dollar of current income and also borrow as much as they can to smooth current consumption until they hit the borrowing constraint. In the low-to-medium range, how savings respond to an income fall becomes more complicated but generally speaking, households dissave (or borrow) to smooth current consumption unless the income fall has a persistence close to a unit root. However, this borrowing demand gets less potent with higher persistence. In fact, when income shocks are more persistent than a unit root, households see the income fall as the harbinger of worse things to come and respond by increasing their savings, letting current consumption decline more than current income. Turning to the very high resource range, while the same trade-offs for low-to-medium group are present, the precautionary response is weaker in magnitude because these households are not in immediate danger of hitting their borrowing limits.

So how much does this uncertainty channel matter for r^*? This is certainly a quantitative question, and among other things, it depends on the distribution of the population across these three ranges. In infinite horizon models, households typically have enough time to move out of the bottom range, so the model does not generate large differences in savings rates by income. But models that feature a lifecycle structure, large wealth inequality, or differences in risk aversion can deliver significant

differences in savings rates and large borrowing responses to a negative income shock that is fairly persistent. An earlier literature in the 1990s and 2000s considered two-agent models with a subset of these features that generated such behavior.

In one example I know best (Guvenen 2009), I proposed a two-agent model with income uncertainty in which one group (that turns out to be the wealthy in equilibrium) could trade in a risk free asset with the other group but could also invest in the stock of a firm that produces aggregate output. In addition, this first group was assumed to have higher elasticity of intertemporal substitution (EIS) than the other. This limited stock market participation model generates substantial wealth inequality (despite assuming same labor income for both types) as well as delivering other macro and asset market facts, including a high and volatile equity premium.

So how is r^* determined? Low-wealth households have a very strong consumption smoothing demand, which means their saving demand is income-elastic but interest-inelastic. This is because these households want to borrow heavily when their income falls to prevent consumption from falling, and this demand does not vary too much with the interest rate. High-wealth households are at the other end of the spectrum: the income elasticity of their savings is lower (because they are more willing to let their consumption fall and rise) and their interest elasticity is higher both because of their preferences (higher EIS) and also because they have another asset—equity—they can invest in. So, now let's consider a scenario where only the low-wealth group experiences an income fall that is not permanent. Their strong desire to borrow meets the reluctance of the high-wealth group who could earn the equity return instead of lending to the low-wealth group. To convince them to lend, the equilibrium interest rate, r^*, has to rise.

One can consider variations of this scenario where the same forces are at play but their strength differ. For example, if inequality increases in a symmetric fashion so that the income of high-wealth households rise (while the income of low-wealth households fall), this rise will tamper the reluctance of the high wealth group to lend, so the upward pressure on the interest rate will be smaller than in the first scenario.

As this example illustrates, how inequality rises matters for the effects on r^*: Was it mainly driven by an acceleration in income growth for top earners or a stagnation/fall in the incomes of everyone else, or both? The answer is in Chart 2, which plots the percentiles of the U.S. wage distribution since 1967. Percentiles are normalized to zero in the initial year to emphasize the changes over time. I also plot a 1.5% annual growth trend line (dashed black line) to represent a conservative estimate of the average U.S. trend wage growth before 1973.

As seen here, even the 95th percentile of the wage distribution grew *at or below* the pre-1970s trend line while all lower percentiles saw a deceleration in growth, and the bottom half of the wage distribution was stretched out by lower percentiles seeing steep declines. This may sound surprising because of how often we hear that the top incomes have risen extremely fast during this time. But those statistics are very often stated as *share* of aggregate income, and the decline in the rest of the economy gives the impression that wages have risen very fast at the top. Chart 2 leaves no doubt that rising inequality was driven by a stagnation or decline in wages for the bulk of the U.S. population. This is along the lines of the first scenario I discussed in the limited participation model.

To summarize my comments so far, I think there is more than one channel through which rising inequality can affect r^* and in the particular cases discussed here, they affect r^* in opposite ways. It seems to me that a model that captures both elements would provide a fuller picture and give us a more definitive evaluation of the rising inequality hypothesis.

Another issue that seems quite important for sorting out these different theories concerns households' expectations during this period. In particular, during the 1970s and 1980s, did households believe that they were living through a trend change that would go on for decades? Or did they have myopic expectations and viewed each year as brining another income decline that happened to go in the same direction year after year? Or was there some learning of the new trend, and if so, how fast was that? As I discussed above, the (perceived) persistence of an income shock is critical for borrowing and saving behavior in heterogenous agent models with economic uncertainty. These expectations also matter for the framework proposed by the authors.

Chart 2
Percentiles of U.S. Wage Distribution, 1967-2018

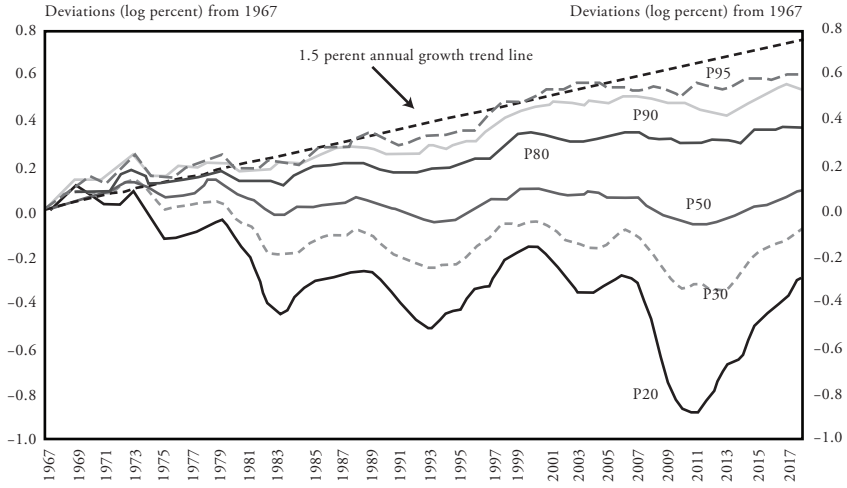

Note: Data on wage percentiles are from Heathcote et al. (2020), who compute them from the Current Population Survey.

Before concluding, I want to briefly mention two other issues that I think are quite important but discussing them would take me beyond the scope of the discussion. The first issue is that safe assets (and close substitutes) make only a small fraction of households' portfolios, which is especially true at the top end (Wolff 2021 Table 6). In addition, many high-wealth households actively invest a portion of their wealth abroad, which further broadens the set of alternatives for their savings, which in turn is likely to limit the impact on r^*.

Second, given the integration of global financial markets and the massive amount of financial flows across countries, it seems most appropriate to me to think about r^* as being determined in the global financial markets. This perspective is also consistent with the fact that interest rates have been declining globally during this period. I think it is much easier to make the "too much" saving argument at the global level because the aggregate saving rate of the world economy has risen significantly (Chart 3), from an average rate of 21% in the 1980s to about 27% in 2019, which plays in favor of the main hypothesis in this paper.

Chart 3
Global Trends: Rising Savings Rates, Slowing Population Growth and GDP Growth

World Gross Savings Rate

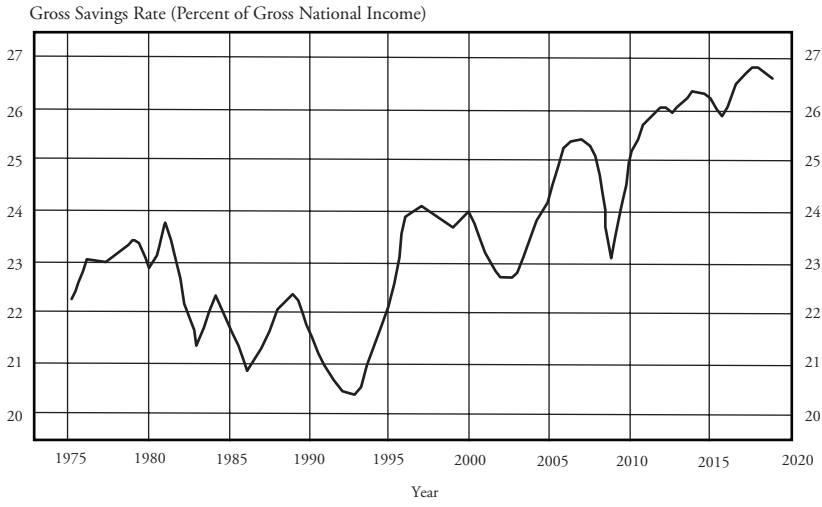

World Annual Population Growth

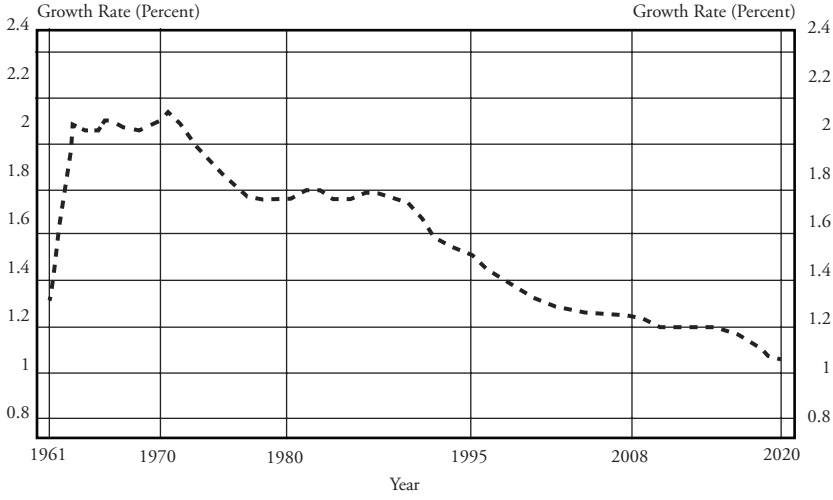

Chart 3 continued

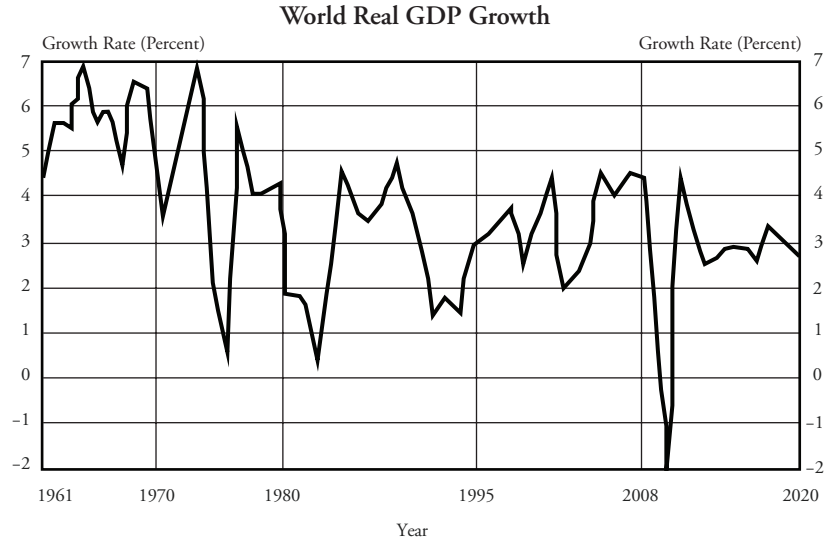

Source: Data used in this chart is the World Bank national accounts data.

As for the contribution of rising inequality to excess savings at the global level, I think that is an open question. While income inequality has been rising in some countries, it has been flat or has decline in some others. Inequality between countries has also been declining. So, to what extent has rising global savings been driven by rising inequality as opposed to aging world population and slowing trend growth (Chart 3) in GDP and productivity? With their extensive expertise in this area, the authors are very well positioned to answer this question.

Author's note: The author is grateful to Fabrizio Perri and Cristina Arellano for insightful discussions and to Cristian Aguilera for excellent research assistance. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Endnote

[1]I am leaving out some technical details that are not critical for this discussion. Further details can be found in Deaton (1991); Aiyagari (1994, 1995), among others.

References

Aiyagari, Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*.

Aiyagari, S. Rao. 1995. "Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting," *Journal of Political Economy*, 103(6), 1158-1175.

Deaton, Angus. 1991. "Saving and Liquidity Constraints," *Econometrica*, September, 59(5), 1221-1248.

Guvenen, Fatih. 2009. "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, November, 77(6), 1711-1750.

Heathcote, Jonathan, Fabrizio Perri and Giovanni L. Violante. 2020. "The Rise of U.S. Earnings Inequality: Does the Cycle Drive the Trend?" *Review of Economic Dynamics*, 18(6), 2715-2754.

Mian, Atif R., Ludwig Straub and Amir Sufi. 2021. "Indebted Demand," *Quarterly Journal of Economics*.

Wolff, Edward. 2021. "Households Wealth Trends in the United States, 1962 to 2019: Median Wealth Rebounds ... But not Enough," Working Paper Series 28383, National Bureau of Economic Research.