Life-Cycle Inequality: the Blacks and Whites Differential*

Giacomo De Giorgi†
GSEM-University of Geneva
BREAD, CEPR, IPA

Luca Gambetti‡
Universitat Autònoma de Barcelona, Barcelona GSE
Università di Torino, and Collegio Carlo Alberto

Costanza Naguib§
University of Bern and GSEM

Abstract

With 40 years of PSID data from 1981 to 2017, we document persistent racial differentials in life-cycle consumption dynamics. Starting from similar positions in the consumption distribution Blacks end up in lower percentiles than Whites. This difference is particularly large for those starting at the top of the consumption distribution. A much lower wealth accumulation for blacks fully accounts for the observed differentials. The findings are in line with a very stylized model of life-cycle consumption where Blacks have lower life-expectancy than Whites, as it is in US data for the relevant cohorts.

JEL classification: E21, E63, D12, C3

Keywords: Consumption, Income, Earnings persistence, quintile transitions

*We would like to thank Costas Meghir, Robert Moffitt, and Luigi Pistaferri for the insightful discussions. Thanks to workshop participants at the First Dondena Workshop, and the BGSE Summer Forum on Inequality, and seminar participants at the University of Warwick, University of Zurich.

†Contact: Giacomo De Giorgi. E-mail: giacomo.degiorgi@gmail.com. De Giorgi gratefully acknowledges the financial support of the EU, through the Marie Curie CIG grant FP7-631510, and the SNF no. 100018-182243.

‡Luca Gambetti acknowledges the financial support from the Spanish Ministry of Science and Innovation, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), the financial support of the Spanish Ministry of Science, Innovation and Universities through grant PGC2018-094364-B-I00, and the Barcelona Graduate School Research Network.

§Contact: Costanza Naguib. E-mail: costanza.naguib@vwi.unibe.ch

1 Introduction

Economic inequality is one of the major challenges for policy-makers and economists. In recent years there has been a burgeoning literature in this area. Starting from Piketty and Saez [2003], several authors have documented and investigated the rising income, and wage inequality, in the US and other countries (see Autor et al. [2008], Bonhomme and Robin [2009], Primiceri and Van Rens [2009], Heathcote et al. [2010], Atkinson et al. [2011], Auten et al. [2013], Attanasio and Pistaferri [2014, 2016], Blundell [2014],Chetty et al. [2014a,b]). A related literature has also stressed the importance of focusing on consumption inequality in order to draw conclusions about households' well-being, consumption being closely related to permanent income (see for example Blundell and Preston [1998], Meyer and Sullivan [2003], Krueger and Perri [2006], Blundell et al. [2008], Attanasio et al. [2014], Aguiar and Bils [2015], Attanasio and Pistaferri [2016], Blundell et al. [2016]).

A stream of literature, closer in spirit to the current paper, has investigated the existing differences in both earnings and consumption levels between Black and White individuals in the US : suggesting that these differentials mostly arise from both the quantity and quality of schooling (see Blau and Graham [1990], Blau and Beller [1992], Card and Krueger [1992, 1993], Oaxaca and Ransom [1994], Chay and Lee [2000], Heckman et al. [2000], Peoples and Talley [2001], Charles et al. [2009], Heywood and Parent [2012], Bayer and Charles [2018]).

The current paper focuses on a less-studied area: the differential in consumption life-cycle dynamics between Blacks and Whites in the US over the past four decades.

We use PSID data from 1981 to 2017 to document persistent racial differentials in life-cycle consumption dynamics across the distribution. More specifically, we document large differences between Black and White in terms of mobility along the consumption distribution. Blacks, independently of the initial percentiles, tend, over the life-cycle, to end up in lower percentiles than white individuals. When controlling for individual characteristics, such as education, age, sector/occupation, these differences are much less important for the bottom end of the distribution. However, even conditional on those observables Blacks starting from the higher percentiles of the distribution tend to fall to lower percentiles than Whites.

We also find that differences in the life-cycle consumption dynamics cannot be accounted for by different income dynamics. Rather they appear to be attributable to a different wealth accumulation process. We show that Blacks, on average, accumulate much less wealth and savings than Whites. When controlling for savings, differences in consumption dynamics completely vanish suggesting that a lower amount of savings and wealth accumulation, thus a lower insurance capacity, is the main cause of the observed differences in the life-cycle dynamics.

We show that a very stylized life-cycle consumption model is able to quantitatively account for the observed differences in the wealth accumulation pattern. The key ingredients are the differences, about 8 years, in life-expectancy and in the gross interest rate of 2-3pp. Such a simple model is consistent with the observed wealth ratio of three between Whites and Blacks.

Our paper is close to Chetty et al. [2020] on the racial differences in the degree of parent-children income persistence, where they find that the main drivers of such differences are geographical segregation and lower marriage rates among the Blacks, which lead to having often only one income in the household. Ganong et al. [2020] use bank data matched with voter registry and firm-wide wage changes data in order to estimate the transmission of unexpected income shocks into consumption by race. They, too, link race differentials in the degree of insurance against shocks to the different amount of wealth held by Blacks and Whites. However, differently from our work, they exclusively focus on short-term consumption changes in reaction to income shocks. Our paper focuses on life-cycle racial differences in consumption persistence.

The remainder of the paper is organized as follows: Section 2 presents the data. Section 3 provides evidence of income and consumption persistence for Blacks and Whites via rank-rank regressions. Section 4 describes Black/White differentials in savings and wealth accumulation. Section 5 describes inequality in life expectancy and its implications for savings, with the help of a toy model. Finally, Section 6 presents a counterfactual analysis and Section 7 concludes. Appendix A provides further information on the construction of the dataset. Appendix B presents the application of the partial insurance model by Blundell et al. [2008] to our case. Appendix C presents some robustness checks.

2 Data

We use data from the Panel Study of Income Dynamics (PSID), a longitudinal survey conducted by the University of Michigan. The PSID began in 1968 with two samples: the Survey of Economic Opportunity (SEO) sample focused on low income families, while the Survey Research Center (SRC) sample interviewed a nationally representative selection of families. Members of these households became PSID "sample members" and were surveyed annually until 1997 (each yearly survey is called a "wave"), after which they were surveyed biannually. Furthermore, all lineal descendants of original sample members become sample members themselves and were independently followed and surveyed once they started their own families. Due to budgetary constraints, in 1997 the PSID dropped approximately 25% of its sample households, with reductions made mainly to the SEO subsample.

The PSID collects a wide range of variables, including information on demographics, income, and consumption. Most data is collected at the household level, though information for PSID-defined household "heads" and "wives" is also gathered. A limited selection of questions are asked about other family members. Typically, a family head is the male in a married pair with primary financial responsibility for the family. A wife is the female counterpart of the married couple. Females only qualify as heads in single adult households (single males can also be heads, of course). If a female head marries a man, he becomes the new head and the woman's classification changes to 'wife.'

2.1 Building the Dataset

To create our dataset for the analysis, we append together all waves from 1968-2017. The full PSID dataset contains 1,856,953 individual-year observations. We limit our sample to the SEO and SRC samples, eliminating individuals from the Immigrant and Latino surveys (two other surveys conducted by the PSID that we do not use due to limited data availability). We also include only current heads, since they are the individuals with the richest and most consistent set of observables over time. As there is one head per household, our analysis is therefore effectively at the household level.

We also create a consistent race indicator for all individuals. The PSID asked heads to identify their race in every wave. For all heads, we assign race as the mode value of race from all reported years. Due to the limited sample size of some reported races, we only keep individuals identifying themselves as Black or White. Our full sample, using all waves of data, includes 457,286 individual-year observations. In our main regression analysis, however, we define a base year of 1981 - so only individuals present in the 1981 wave and beyond are included. This brings us to 342,679 individual-year observations (from waves 1981-2017). The choice of 1981 as the base year is merely dictated by sample size considerations. We have also repeated the analysis using different base years.

2.2 Family Income

The PSID consistently asks respondents to report their household's total monetary income, defined as the sum of the taxable income of the head and wife, the total transfers of the head and the wife, the taxable income of other family unit members, and the transfer income of other family unit members. Beginning with the 1994 wave, the measure also includes total family Social Security income. For prior years, when Social Security was not already included in family income, we added in separate measures of Social Security income to family income. Before 1994, family income in the PSID data is bottom-coded. Any negative or zero values are recorded to \$1. Because this practice occurs for many years, we apply the same rule to the remaining years of data. To convert nominal incomes to real terms, we divide the nominal measure by the Consumer Price Index (CPI). In order to create a per capita measure, we then divide total family income by an Adult Equivalent scale, given by:

$$AE = 1 + 0.7(A - 1) + 0.5K \tag{1}$$

where A is the number of adults in the household and K is the number of children in the household. This scale assigns a value of 1 to the first household member, of 0.7 to each other adult in the household and 0.5 to each child. [1]

Our measure of real adjusted family income (TFA) is

$$TFA_i = \left(\frac{Nominal\ Family\ Income \times 100}{CPI \times AEscale} \right). \tag{2}$$

We multiply Family income by 100 to preserve the scale of the variable given that CPI is equal to 100 in the base year.

2.3 Consumption Imputation

For all years besides 1973, 1988, and 1989, the PSID asks respondents to report the monetary value of their family's consumption of food at home, food away from home, and food stamps. Therefore, household expenditures on food consumption are consistently recorded throughout the entire period. Then, beginning in 1999, households are also asked to detail their spending on a wide array of goods, such as utilities, transportation costs, and healthcare. Unfortunately, spending on clothing, vacations, entertainment, and other similar discretionary spending is only available since 2005.

Since much consumption spending information is not available prior to 1999, we use an approach developed by Blundell et al. [2008] that imputes household consumption using a demand function derived from other variables consistently present in the PSID. Specifically, the method uses spending on food, socio-demographic information (state, age, number of children, maximum education, marital status, disability, etc. . .), and price controls to predict non-food consumption (defined as total expenditures on rent equivalents,

[1] This scale, which is sometimes called the "Oxford scale", has been first proposed by the OECD in 1982. We also probe the robustness of the results to the chosen scale.

home insurance, electricity, heat, water/sewage, miscellaneous utilities, car insurance, gas, parking, bus/train, cabs, other transport, school fees, other school costs, childcare, health insurance, hospital care, doctors, and drugs). The idea is to estimate the relationship between consumption variables and the consistently reported demographic variables and food payments in later years, and use this relationship to predict consumption expenditure in earlier years.

Let n_{it} be our non-food consumption measure, defined as total expenditures on rent equivalents, home insurance, electricity, heat, water/sewage, miscellaneous utilities, car insurance, gas, parking, bus/train, cabs, other transport, school fees, other school costs, childcare, health insurance, hospital care, doctors, and drugs. Since consumption can sometimes take on the value of zero, instead of taking the log of consumption we consider the Inverse Hyperbolic Sine transformation, defined as:

$$IHS(n_{it}) = \ln\left(n_{it} + \sqrt{n_{it}^2 + 1}\right) \tag{3}$$

The model consists of the following regression, estimated using data from waves 1999-2017:

$$IHS(n_{it}) = Z_{it}\beta + p_t\gamma + g((f_{it};\theta)|\mu_{it}) \tag{4}$$

Z represents an array of dummy variables for our various demographic covariates: race, state, age, number of children, maximum education, employment, marital status, home-owner status, self-employment, and disability. We also include continuous covariates for total number of hours worked and total number of family members. We add p for price controls: the yearly CPI and the CPIs for food at home, food away from home, and rent. The polynomial function *g(.)* includes food expenditures, *f*, and μ is an error term, with θ measuring the importance of the different types of food expenditure. f_{it} stands for food at home, food away from home and food stamps, for individual i in year t.

Once the demand function is estimated for years 1999-2017, we use the coefficients to predict non-food consumption in all waves, including the earlier years. Total imputed consumption is then found for each household-year by adding their actual food consumption to the imputed non-food consumption. First we recover imputed non-food consumption from the Inverse Hyperbolic Sine transformation, and then we add actual food consumption:

$$c_{hat} = exp\{Z_{it}\hat{\beta} + p_t\hat{\gamma} + g(f_{it};\hat{\theta})\} \tag{5}$$

$$c_{IHS} = \left(\frac{c_{hat}^2 - 1}{2 \times c_{hat}}\right) \tag{6}$$

$$Nominal\ Imputed\ Consumption = \hat{c} = f_{it} + c_{IHS} \tag{7}$$

Since all consumption expenditures are reported at the household level, we divide this value by the Adult Equivalent scale. We also use a CPI adjustment to convert consumption into real terms:

$$Real\ Imputed\ Consumption = TCP_i = \left(\frac{\hat{c} \times 100}{CPI \times AEscale}\right) \tag{8}$$

As we will see later, we also probe our results using actual consumption expenditure whenever possible.

3 Life-Cycle Consumption Dynamics

We turn now our attention to life-cycle dynamics. More specifically we try to depict racial differences in terms of dynamics over the life-cycle within the income and consumption distribution. We use rank-rank regression to assess in which part of the distribution individuals starting from a given percentile end up at different time intervals, ten, twenty, thirty and about forty years. In all the results and Figures presented in this Section, each individual in a given year is assigned to the consumption percentile according to her position in the overall consumption distribution of that particular year. Our aim here is to describe the short and long-term dynamics along the national distribution.

3.1 Rank-Rank Regressions

In order to obtain insights on the differences in the degree of income and consumption persistence of Blacks and Whites along the overall distribution, we perform a rank-rank analysis in the spirit of Chetty et al. [2020]. The idea is to estimate the mean consumption percentile in which an individual ends up given that she was in a given consumption percentile in the base year. For example, we may consider individuals who were in the 10th lowest consumption percentile in 1981 and look at where they end up on average in, say, 2011. The rank-rank analysis allows us to answer such questions, separately by race, and to provide a simple graphic intuition of the results. We aim at assessing whether Black and White individuals end up in different consumption percentiles, either starting from the top, middle, or the bottom of the distribution after a few decades from the base year (1981). We perform this analysis separately by race, both in an unconditional and in a conditional version, i.e. controlling for age, age squared, gender, occupation, and education.

In Figure 1, we report the results of the unconditional (left column) and conditional (right column) rank-rank regressions for consumption. On the x-axis we have the percentile in 1981 while on the y-axis we report the percentile in 1991 (first row), 2001 (second row), 2011 (third row) and 2017 (fourth row) covering then 10-20-30 and 36 years transitions. In all graph we include the 45 degree line indicating a perfect correspondence of the ranking overtime.

Let us first discuss the results of the unconditional analysis. First, in all of the panels of the figure, the blue line (Whites) is above the red line (Blacks). Thus, for any possible percentile of origin in the consumption distribution, the average percentile of destination of Blacks is lower than that of Whites: black individuals tend to shift downward in the consumption distribution. The result suggests that Whites tend to be much more persistent at the top of the distribution relative to Blacks no matter the length of the period considered. Just to provide an example, from the upper left panel of Figure 1 we see that if a Black individual was around the 100th (top) percentile of consumption in 1981, then on average she will end up in the 60th percentile after ten years. While a White individual being in the 100th (top) percentile in 1981 will end up on average in the 80th percentile after ten years.

Second, the intersection with the 45 degree line for Blacks coincide with the 20th percentile while for White with the 50th. This means that 80% of Blacks are worse off than where they started already after a few years, while this is true only for 50% of Whites.

Third, the slope of the regression line provides information about the extent of mean reversion. A flat line would imply that on average the percentile of destination is the same no matter the percentile of origin. Especially for income, and to a lesser extent also for consumption, the regression line for Blacks is steeper than for Whites. This suggests that for white individuals there is a higher tendency to mean reverting. Black individuals not only tend to shift downward in the consumption distribution but also tend to be much less mobile than Whites.

We extend the above analysis in a conditional format, i.e. using residual consumption for the analysis. In practice, first we regress consumption, on a set of control variables (i.e. age, age squared, gender, education and industry of occupation), then the estimated residuals of this regression are ranked for each year of the analysis. The approach allows then for the control variables to have different effects in different years.

When we control for individual characteristics, racial differences in terms of consumption dynamics are substantially mitigated or even disappear at the bottom of the distribution. Indeed, the percentiles of destination for white and black individuals starting in 1981 in the left tail of the distribution are essentially the same.

On the contrary, the differences remain and in some cases appear to be slightly larger than in the unconditional analysis at the top of the distribution where the average percentile of destination for white individuals is much higher than the average percentile of black individuals. For instance, a white individual who is at the top percentile in 1981 falls on average to the 60th percentile while a black individual falls to the 40th percentile. In general both the blue and red line for consumption flatten substantially. This suggests a tendency towards the median of the within-race distribution independently on the initial percentile once substantial observable heterogeneity is accounted for. However this tendency is higher for Blacks (the red line flattens more), amplifying the difference at the top of the distribution. The result again confirms our previous finding: persistence at the top of the distribution is higher for Whites even after controlling for observable characteristics including education and occupation.

A plausible explanation could be a different dynamic in terms of income between Blacks and Whites at the top. So, to assess whether this is the case we add an additional control for income among the conditioning variables and we repeat the analysis. Figure 2 plots the results. In the left column we have the estimates without controlling for income, in the right column the estimates obtained controlling also for income. The results are essentially unchanged. Controlling for income does not make any difference in the estimated dynamics. It therefore appears quite unlikely that the different life-cycle dynamic between Blacks and Whites are determined by substantial differences in the income process.

To summarize, the rank-rank regressions show that, in general, Blacks exhibit a higher degree of downward mobility than Whites at any time interval. Accounting for individual characteristics mitigates substantially the differences in consumption at the bottom of the distribution but not at the top. This is true even when controlling for individual income.

(a) No controls (b) Controls

Rank in 1991

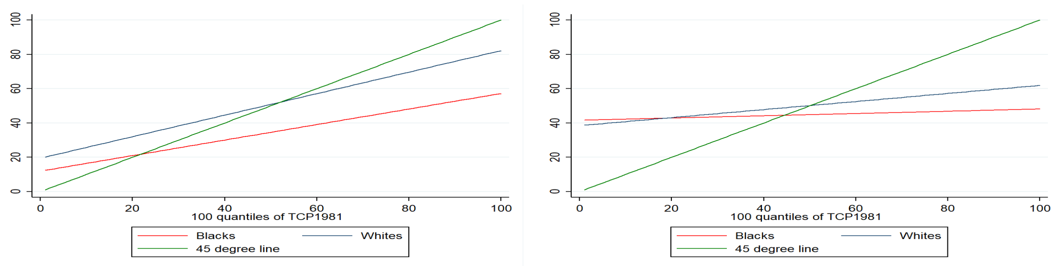

Rank in 2001

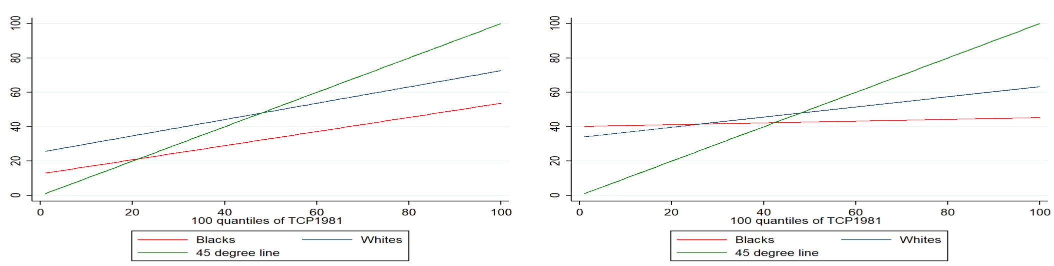

Rank in 2011

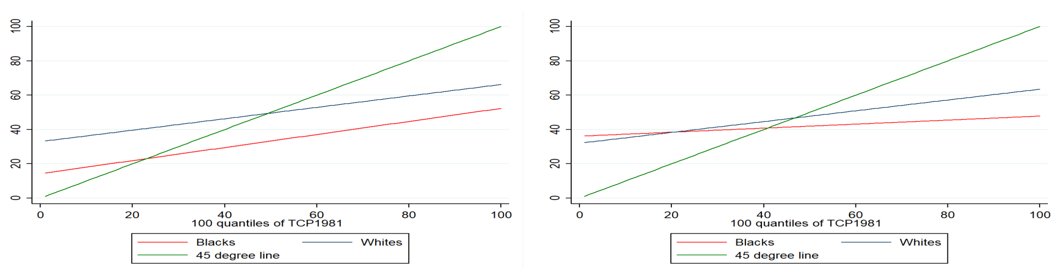

Rank in 2017

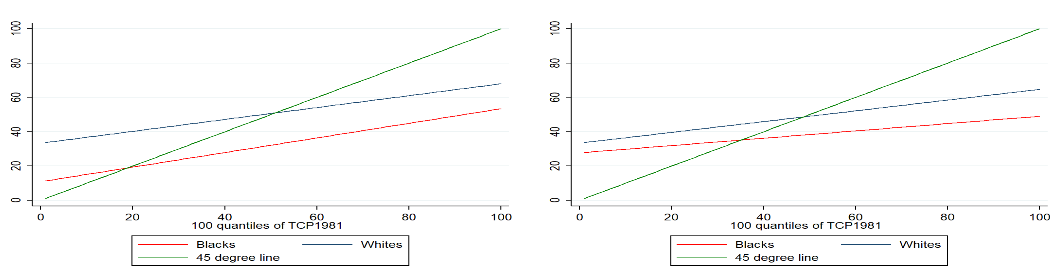

Figure 1: Average TCP consumption rank in a fixed year for an individual who was in each income or consumption percentile in 1981, by race. Left column: no control variables. Right column: residual TCP has been obtained by regressing TCP on a set of controls (age, age squared, gender, occupation, education) and taking the residual terms.

(a) Controls (b) Controls + TFA

Rank in 1991

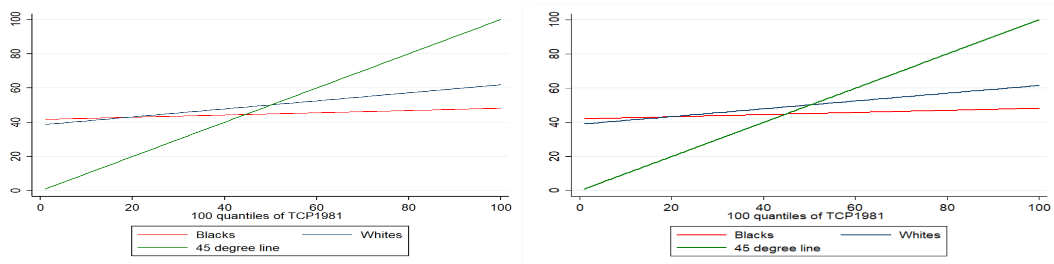

Rank in 2001

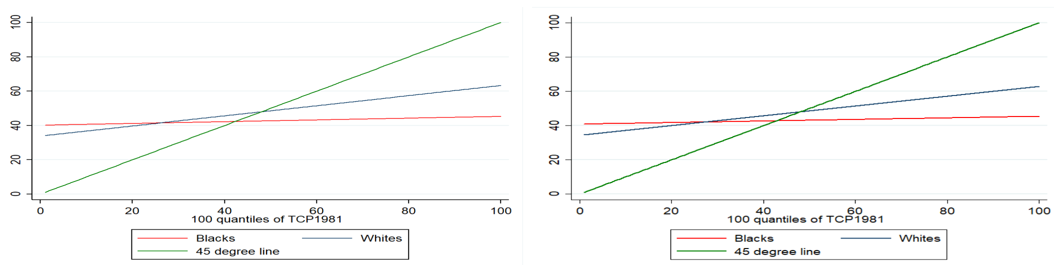

Rank in 2011

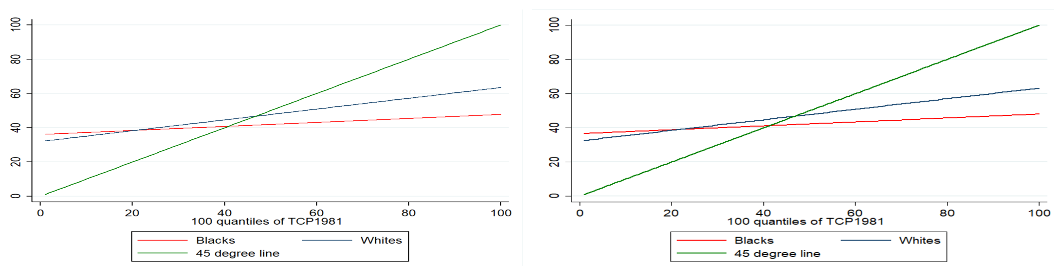

Rank in 2017

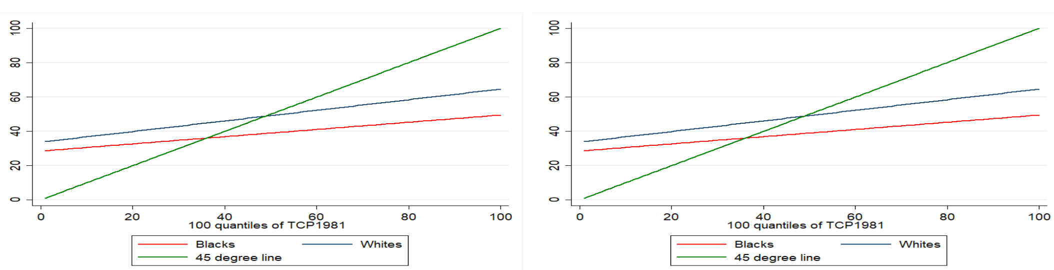

Figure 2: Average TCP consumption rank in a fixed year for an individual who was in each income or consumption percentile in 1981, by race. Left column: residual TCP has been obtained by regressing TCP on a set of controls (age, age squared, gender, occupation, education) and taking the residual terms. Right column: same controls as in the left column plus individual income.

4 Savings and Wealth Accumulation

In the previous section we have shown that different income dynamics cannot explain the differences in the dynamics of consumption at the top of the distribution between black and white individuals. Here we assess the plausibility of an alternative (and potentially complementary) explanation: the different process of wealth accumulation. The idea is that if wealth accumulation for Blacks is lower than for Whites, similar health and income shocks at the top of the consumption distribution will result in very different consumption movement. This is because Blacks are more exposed to shocks due to more limited buffer, and therefore the lack of self-insurance.

We compute wealth as the comprehensively as we can in the PSID, summing up seven asset types: imputed value of farm or business, imputed value of cash savings, imputed value of real estate other than home, imputed value of stocks, imputed value of vehicles, imputed value of other assets, value of home equity net of debt. This wealth measure is then divided by the Consumer Price Index (CPI), in order to obtain a measure of wealth in real terms. We focus on the quintiles of the distribution.

4.1 Evidence

Panel (a) of Figure 3 plots the extensive margin for wealth holding, i.e. the percentage of individuals having positive wealth across consumption quintiles. Differences between Blacks and Whites households are between 2-7 percentage points across all the consumption quintiles. From panel (b) it can be seen that, among people with a positive wealth, Blacks own far less of it than Whites, particularly at the top quintile. Wealth accumulated by Blacks is on average 3 times smaller than that accumulated by the Whites, with the exception of the top quintile where that ratio becomes almost 7.

In Figure 3, panel (c) and (d), we plot the intensive and extensive margin for stock holdings. A few remarks are in order. First, there are huge racial differences in the percentage of households holding stocks. For example, in the top 1981 TCP quintile only, more than 10% of Whites households hold stocks, whereas only around 1% of Blacks households do. Second, there are also relevant differences in the amount of stocks held. In the top 1981 TCP quintile, a White individual holds on average 180,000 US dollars in stocks, whereas a Black individual only holds around 40,000 US dollars. These differences in stock holdings are suggestive of large differential returns on assets and this difference is particularly relevant at the top TCP quintile. Given the very low probability of holding stocks in the bottom 3 quintiles (below 5%), especially for Blacks (below 2%), we

Total Wealth

Panel (a)

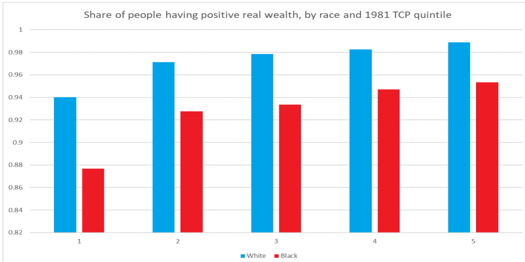

Panel (b)

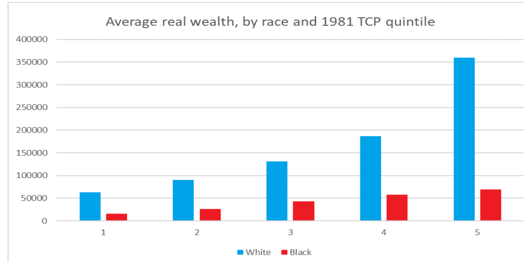

Stock Holdings

Panel (c)

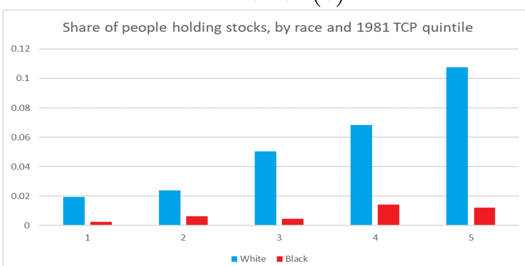

Panel (d)

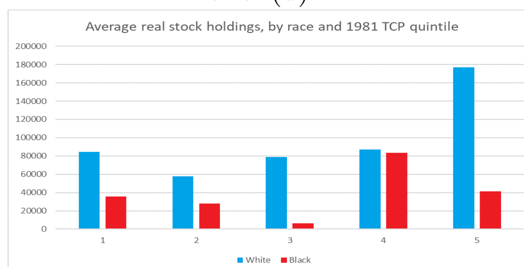

Annuities

Panel (e)

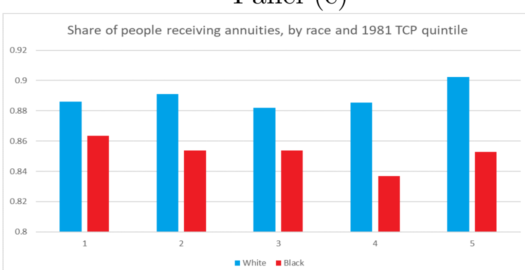

Panel (f)

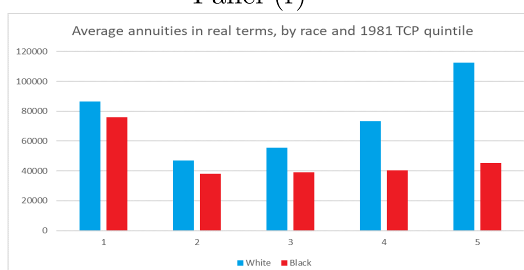

Figure 3: Total wealth: Panel (a) extensive margin: share of people having positive real wealth, by 1981 TCP quintile. Panel (b) intensive margin: wealth by 1981 TCP quintile (people with zero wealth have been excluded). Wealth iscomputed as the sum of seven asset types: imputed value of farm or business, imputed value of cash savings, imputed value of real estate other than home, imputed value of stocks, imputed value of vehicles, imputed value of other assets, value of home equity net of debt. Divided by the Consumer Price Index (CPI). Sample 1989, 1994 and 1999-2017. Stock holdings: average value of stock holdings in US dollars divided by CPI by 1981 TCP quintile. Panel (c) extensive margin. Panel (d) intensive margin (people with no stocks have been excluded). Sample: 1999-2017. Annuities: Panel (e) extensive margin for receiving annuities. Panel (f) intensive margin (people receiving no annuities have been excluded). Sample 1999-2017.

shouldn't be surprised of the intensive margin for the bottom 4 quintiles showing some non-monotonic relation.

In the PSID data the variable "annuities" is defined as follows. After a first filter question: "Did you receive any income in previous year from other retirement pay, pensions, or annuities?", then the sub-question "Then, how much of this was from annuities or IRAs?" gives the numeric value of our variable of interest. An annuity is commonly defined as a financial product that pays out a fixed stream of payments to an individual. These financial products are often used to ensure having a stable income stream during retirement age, as well as to avoid the risk of outliving one's savings. From Figure 3, we find evidence of significant racial differences, not only in the share of individuals receiving annuities, but also, among those receiving it, in the average yearly amount received. From Panel (e), the share of people receiving annuities is around 4pp higher among Whites in each consumption quintile. Further, from Panel (f), it emerges that among people who receive annuities, Whites in the top 1981 TCP quintile receive an average amount which is more than double than that received on average by Blacks in the same consumption quintile.

Summing up, we find substantial racial differences in the amount of savings and wealth accumulation. In particular, the ratio of wealth held by Whites versus Blacks increases along the consumption distribution, and in general appears to be at least 3-7 times higher for Whites than for the Blacks. Such differences, although large, are consistent with the predictions of the model discussed above given a racial differential in life expectancy of 8 years and a return on asset differential of about 3%.

4.2 Life-cycle Dynamics Controlling for Savings

The main conclusion from the previous sections is that Blacks are disproportionately exposed to downward mobility in the upper part of the distribution. This result is consistent with the prediction discussed in Section 4 that the lower amount of savings and wealth accumulation determines a substantially higher downfall risk for Blacks. To test whether savings can explain the differences in the estimated life-cycle dynamics we repeat the rank-rank analysis controlling also for savings we add as a control variable the value of savings. Given that complete information on savings has only been collected in the PSID since 1999, this analysis is only performed for the years 1999-2017. Hence, we use here 1999 as a reference year. This means that this analysis is performed on a subsample of our dataset, which contains 82'145 individual-year observations. We can here also use actual consumption as that exists from 1999 onwards. Figure 4 reports the results. When

(a) No control (b) Control

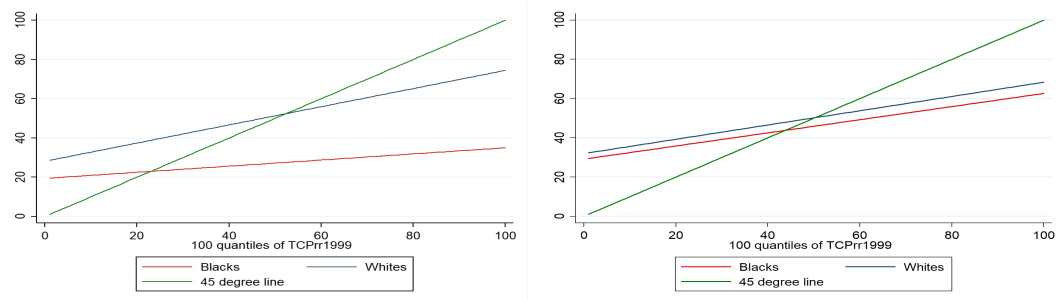

Figure 4: Rank-rank regression (1999-2017). Average residual TCP consumption rank in 2017 for an individual who was in each residual consumption percentile in 1999. Blue stands for Whites, red for Blacks. Residual consumption has been obtained regressing TCP on a set of controls (age, age squared, gender, occupation, education) and taking the residual terms. In Panel (b) we include savings as a control. Here savings are defined as home value equity and cash savings.

controlling for savings, the differences between Blacks and Whites in the rank-rank regressions virtually disappear, the two lines being extremely close (right panel). This difference is not due to the different sample used. Indeed, when performing the same analysis on the same sample without controlling for savings, differences are wide (left panel) and as wide as for the corresponding figures from 1981. Also in terms of persistence, the differences are substantially mitigated compared to the estimates reported in Figure 2. Due to a shorter life expectancy and a lower return on asset (partly due to lower access to the stock market), Blacks save, on average, much less than Whites and cumulate also much lower wealth. Further, Blacks seem to under-insure their health taking up lower premium plans (see Figure B.2 in Appendix). This translates into a much lower degree of insurance against shocks and results into a much higher downward mobility in the consumption distribution. Such results is also corroborated by the evidence provided in Appendix B.1 where we show that both transitory and permanent income shocks have a larger effects on consumption for Blacks than for Whites. Similarly, in Appendix B.2, we show how Blacks purchase less insurance and are more subject to health shocks than Whites all along the Consumption distribution (surprisingly more exposed to health risk at the top of the consumption distribution).

5 Inequality in Life Expectancy

When controlling for savings, the differences in consumption dynamics vanish. Thus the question is: why do Blacks accumulate less wealth than Whites? We believe that a shorter life expectancy and a lower rate of return are part of the explanation. We show, using an extremely stylized model, that a life expectancy differential like the one observed for the cohorts of interest and a plausibly lower rate of return deliver predictions in terms of savings and wealth accumulation in line with those observed in the data.

5.1 Estimated Life Expectancy

Our cohorts of interest in the PSID are born between 1917 (i.e. individuals who are 64 in 1981, the first year that we consider in our sample) and 1997 (i.e. individuals who are 20 in 2017, the latest year in our sample). According to the United States Life Tables prepared by Arias and Xu [2019], for those born in 1930 the life-expectancy differential Whites to Blacks is of 14 years overall, in 1940 is of 11 years, in 1950 is of 8 years, in 1960 and 1970 of 7 years, then in 1980 is 6.4 years, and in 1990 is 7 years. Overall given the distribution of year of birth in our data, a life-expectancy difference of 8 years is in line with the figures; and we will work under that benchmark of 8 years difference in life-expectancy at birth.[2] Importantly, the gradient of life-expectancy across education categories (as a proxy for permanent income) is much lower for Blacks than for Whites (see for example Hummer and Hernandez [2013]) which is consistent with a life-expectancy gap that doesn't close at high level of consumption.

We take the measure at birth as widely available and so to avoid making assumptions on the specific individual and household decision-making process. We note that this could be a reasonable approximation, as male differentials are substantially larger all the way to the 1980's and 1990's. Females' life-expectancy is higher than males' and this is true in particular for Blacks. We also note that the gap in life-expectancy is not closing at a fast rate: it shrunk by 45% between 1930 and 1950 and only by 20% in the last two decades, and due to current circumstances at the time of COVID-19, that gap is potentially getting larger.[3]

[2] One might want to start from differences in life-expectancy around age 18/20, when some of the financial decisions are taken and so to take into account concerns regarding low life-expectancy due to infant and child mortality. One might also want to consider male-female differential and its role in the household decision process.

[3] We know from recent CDC work that the mortality rates and overall deaths have been proportionally much larger among minorities in the US. With Blacks dying at a rate almost double that of Whites (`https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/racial-ethnic-minorities.html`)

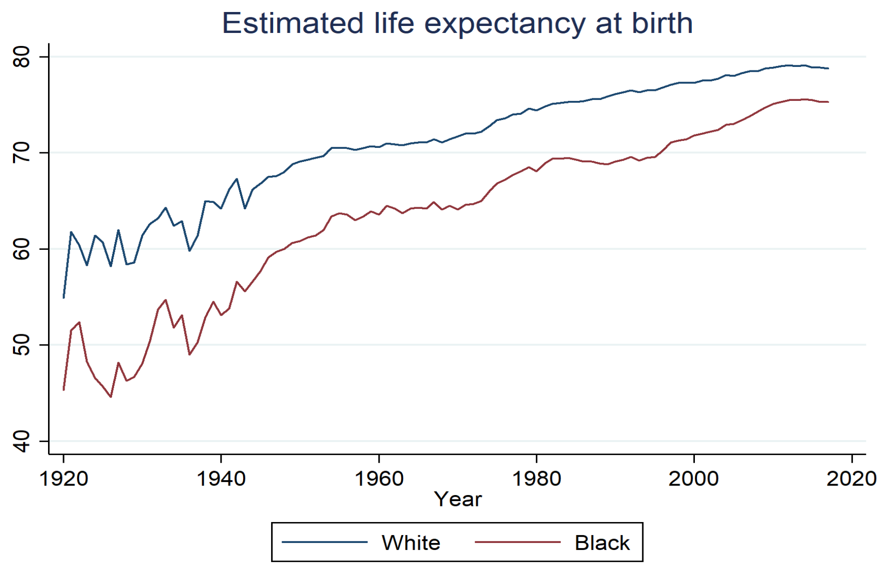

Figure 5: Estimated life expectancy at birth, by race. Data from the National Center for Health Statistics, Centers for Disease Control and Prevention Arias and Xu [2019]).

In Figure 5, we report the estimated life expectancy, by race, from 1920 to 2017. It is apparent from this Figure that life expectancy of Whites has been consistently between 15 and 5 years longer than that for Blacks, even if, as mentioned above, this difference has shrunk over time.

This marked difference in life-expectancy, of about 8 years, is a crucial piece in our analysis of consumption behavior over the life-cycle as we will show below.

5.2 A Simple Model

Here we provide a sense of the effects of life-expectancy differences in terms of saving rates and stock of savings between Blacks and Whites. At the same time we will add another crucial parameter to that decision-making process and therefore model all the mechanics through the interactions of two fundamental parameters in a life-cycle model of consumption: (i) life expectancy, and (ii) (gross) rate of return. We purposely use a basic off-the-shelf model of consumption without uncertainty. This is useful to establish a benchmark and see how far we get with the simplest model.

Let us introduce some notation. An individual maximizes lifetime consumption c subject to an inter-temporal budget constraint, we abstract from labor supply and fix the income per working period to $y_t = y_0$ where $t = 0$ indicates the first period of adult life, say 18 years of age. The agent works until retirement, i.e. for $L = 45$ periods, and lives for a total of T^j periods with $j = W, B$ and $T^W > T^B$ (where B stands for Blacks and W for Whites).

The allocation of consumption is then chosen according to the following maximization problem of the lifetime utility

$$\begin{aligned} \max_{c_t^j} \quad & \sum_{t=0}^{T^j} \beta^t U\left(c_t^j\right) \\ s.t. \quad & \sum_{t=0}^{T^j} \frac{c_t^j}{\left(R^j\right)^t} \leq \sum_{t=0}^{L} \frac{y_t}{\left(R^j\right)^t}. \end{aligned}$$

We fix Blacks' and Whites' incomes to be the same and to follow the same profile, this is because we are only interested in the role of life-expectancy and secondarily rate of returns. Blacks and Whites have the same working life of $L = 45$ years (start working at 18 and retire at 63 in line with the literature (see for example FRED data)). Whites die at age 80, while Blacks at age 72.

We assume a CES utility with $RRA = \theta$ and common discount factor β. It should be clear that we are abstracting from explicit differences in the discount factor β, curvature θ, and bequest motives. Importantly Altonji et al. [2000] state that

> *[... Several studies, including those mentioned above, have found large wealth differences even after controlling for differences between blacks and whites in average income and other factors. For example, Blau and Graham [1990] conclude that as little as one quarter of the wealth gap can be attributed to racial differences in income and demographic variables...] [...They tentatively suggest that the race difference in the wealth models is not driven primarily by inter vivos gifts and inheritances...]*

$$\begin{aligned} c_0^j \quad & = \frac{1 - \left(\beta R^{1-\theta}\right)^{\frac{1}{\theta}}}{1 - \left(\beta R^{1-\theta}\right)^{\frac{T^j}{\theta}}} \frac{1 - R^{-L}}{1 - R^{-1}} y_0 \\ c_t^j \quad & = (\beta R)^{\frac{t}{\theta}} c_0^j. \end{aligned}$$

We can then assess how savings and consumption profiles vary depending on our parameters of interest: difference in life-expectancy and gross interest rates.

It is important to note that in these models what matters in terms of consumption and savings evolution over the life-cycle is the discounted (gross) rate of return, so that aside for the initial level of consumption c_0 one cannot parse out R, and β. In Figure 6 below we present a series of scenarios characterized by the difference in life-expectancy $T^W - T^B = 0, 8, 12$ and gross returns on assets $R^W = 1.07$ for Whites, while we vary it for Blacks ($R^B = [1.02, 1.07]$); finally, we fix $\theta = 1.5$, and $\beta = .995$. For the difference in gross returns we base our scenarios on the existing literature on asset allocations (for example Badu et al. [1999] write: *...We find that Black households are significantly more risk averse in their choice of assets. Further, we find that Black households typically pay higher rates for several types of credit instruments, even though they self identify as conducting significantly more extensive searches in the financial markets...*). Similarly, Menchik and Jianakoplos [1997] suggests that blacks have a lower rate of return on assets, in particular because of the composition as we show in the previous sections. For the benchmark gross returns (1.07) we use the long-term figures suggested in Jordà et al. [2017].

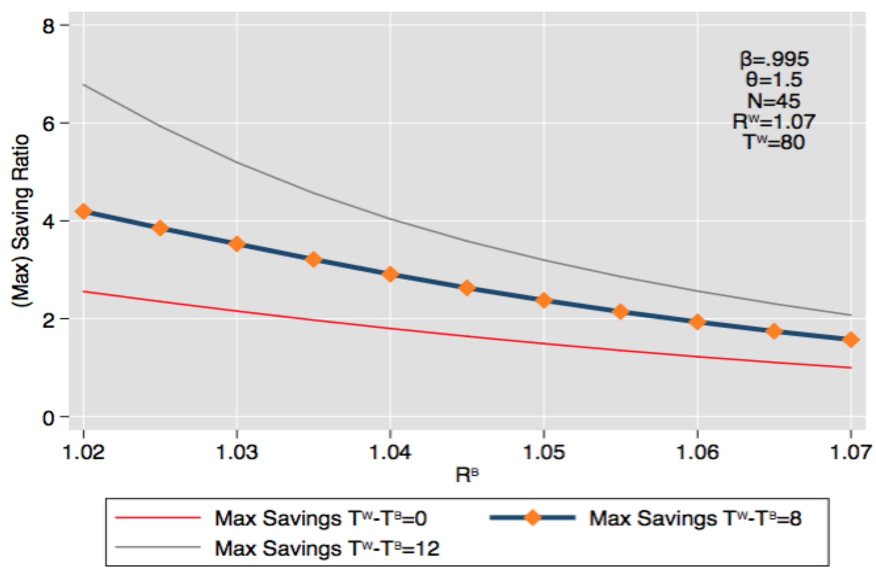

Figure 6: Model Simulations for different Life-expectancy and (gross) Interest Rates

What is immediately visible from Figure 6, where we show the (Whites over Blacks) ratio of the max savings is that the combination of difference in life-expectancy and returns on asset contribute substantially to the accumulation of savings in life. The larger the differences, the larger the savings gap. As we have shown, in the data, the wealth of White individuals is on average 3.5 times larger than the wealth of Blacks. That number is matched by this simple model when the life-expectancy differential is 8 years and the differential in the rate of return is around 3 percentage points.

This simple model appears able to fit an important starting point for the current paper. Given a life-expectancy difference of 8 years, as in our cohorts, and gross rate of return differential of 3 percentage points, we can explain (almost fully) the differential saving stocks between Blacks and Whites. As shown earlier, saving differences across race close the life-cycle consumption differences between Blacks and Whites. There are two fundamental reasons for that: 1. higher savings produce higher wealth to sustain consumption; 2. a larger buffer stock of savings avoids falls in consumption due to both temporary and permanent shocks. On this second channel we provide evidence in Section B.1 where we show that permanent and transitory shocks have a larger impact on Blacks than on Whites.

6 Counterfactual Analysis

In this Section, we present a simple counterfactual exercise, i.e. we estimate the impact of being Black on the amount of savings held, consumption, and income by giving to Blacks the same age, gender, education, industry distribution of Whites. In practice we estimate simple OLS of the outcome variables, savings/consumption/income, as a function of age, age squared, gender, education and industry/occupation. We then use the estimated parameters on actual Blacks as weights in the counterfactual exercise.

We restrict this analysis to households in the top consumption quintile, since it is the main focus of our research. In essence, the exercise allows us to make some headway on the determinants of the large differences in savings between Blacks and Whites, and in particular whether some unobservables including individual life-expectancy are crucial determinants of such differences. From this counterfactual analysis summarized in Figure 7 we deduce that the gap between the savings distributions of Whites and Blacks isn't closed by observables, in fact the role of the residuals appears extremely large as on average there is a very large gap between the true and counterfactual distribution. This once more is consistent with the role of unobservable to the econometrician life-expectancy (among other possible unobservables).

This counterfactual analysis is in line with a substantial part of the related literature on wealth inequality between Blacks and Whites as summarized by Altonji et al. [2000]. What seems quite interesting in these results is that while demographics close the income and, to a lesser extent, consumption gaps, the same is not true at all for savings. Once more this reaffirms the role of unobservables in such a difference between blacks and whites.

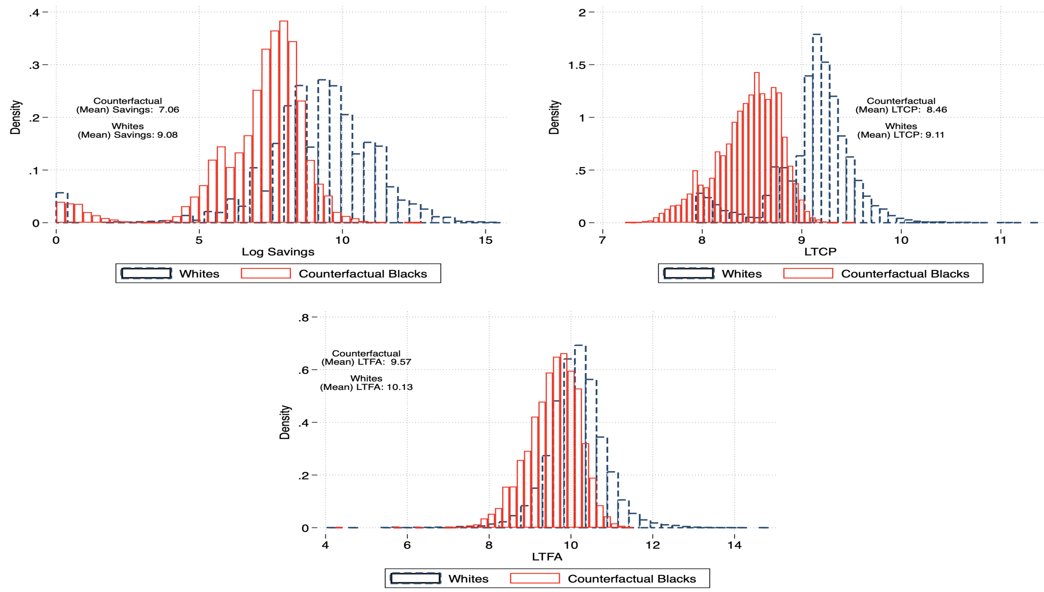

Figure 7: Impact of being Black on log savings (upper left panel), log TCP (upper right panel) and log TFA (bottom panel), based on the estimation of the difference between the actual savings/consumption/income distributions of the Whites and the counterfactual distribution of savings /consumption/income of the Blacks had they had the same characteristics of Whites in terms of age, education, gender, and occupation. Top TCP quintile in 1981 only. Data for 1999-2017.

7 Concluding Remarks

Our analysis strongly points towards the role of savings and asset accumulation as a key driver of racial differentials in the consumption dynamics. In particular, we show that Blacks at the top of the consumption distribution tend to fall in the ranking much more than Whites after a few years and that fall is persistent. At the same time while socio-demographics characteristics close the life-cycle dynamics between Blacks and Whites at the bottom of the consumption distribution, they do very little at the top. It is well known and confirmed here that Blacks and Whites differ substantially in their amount of savings and wealth, it is however novel that we show how those differences persist even when comparing Blacks and Whites with initially similar levels of consumption. The gap in savings then results into lower consumption dynamics and lower ability to insure against shocks, both permanent and transitory. The lack of insurance in face of both permanent and temporary shocks, such as income and health shocks, makes Blacks more vulnerable and in fact more prone to downfalls in the consumption distribution. While a standard analysis of mobility would probably show Blacks to be more mobile, the reality is that they are more mobile downwards and not upwards, both in the short and long run. Differential life-expectancy, about 8 years for our cohorts, seems to contribute substantially to such a life-cycle profile, and in particular, when paired with lower access to high-return assets. While understanding where such differences in life expectancy are coming from is beyond the scope of the current paper, policy actions to improve access to insurance and financial markets could be promoted by the policy-maker.

References

M. Aguiar and M. Bils. Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9):2725–56, 2015.

Joseph G. Altonji, Ulrich Doraszelski, and Lewis Segal. Black/white differences in wealth. *Economic Perspectives*, 24:1, 2000.

E. Arias and Jiaquan Xu. United states life tables. 68:7, 2019. 2017. National Vital Statistics Reports. Hyattsville, MD: National Center for Health Statistics.

K. J. Arrow. What has economics to say about racial discrimination? *Journal of economic perspectives*, 12(2):91–100, 1998.

A. B. Atkinson, T. Piketty, and E. Saez. Top incomes in the long run of history. *Journal of economic literature*, 49(1):3–71, 2011.

O. Attanasio and S. J. Davis. Relative wage movements and the distribution of consumption. *Journal of political Economy*, 104(6):1227–1262, 1996.

O. Attanasio and N. Pavoni. *Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption (mimeo)*. University College, London, 2006.

O. Attanasio and L. Pistaferri. Consumption inequality over the last half century: some evidence using the new psid consumption measure. *American Economic Review*, 104 (5):122–26, 2014.

O. Attanasio, E. Hurst, and L. Pistaferri. The evolution of income, consumption and leisure inequality in the united states 1980-2010. *Improving the measurement of consumer expenditures*, University of Chicago Press:100–140, 2014.

O. P. Attanasio and L. Pistaferri. Consumption inequality. *Journal of Economic Perspectives*, 30(2):3–28, 2016.

G. Auten, G. Gee, and N. Turner. Income inequality, mobility and turnover at the top in the us 1987-2010. *American Economic Review*, 103(3):168–72, 2013.

D. H. Autor, L. F. Katz, and M. S. Kearney. Trends in us wage inequality: Revising the revisionists. *The Review of economics and statistics*, 90(2):300–323, 2008.

Y. A. Badu, Kenneth N. Daniels, and Daniel P. Salandro. An empirical analysis of differences in black and white asset and liability combinations, financial services review, volume 8. *Issue*, 3:129–147, 1999.

P. Bayer and K. K. Charles. Divergent paths: A new perspective on earnings differences between black and white men since 1940. *The Quarterly Journal of Economics*, 133(3): 1459–1501, 2018.

P. Bayer, F. Ferreira, and S. L. Ross. The vulnerability of minority homeowners in the housing boom and bust. *American Economic Journal: Economic Policy*, 8(1):1–27, 2016.

P. Bayer, F. Ferreira, and S. L. Ross. What drives racial and ethnic differences in high-cost mortgages? the role of high-risk lenders. *The Review of Financial Studies*, 31(1): 175–205, 2017.

D. G. Blanchflower, P. B. Levine, and D. J. Zimmerman. Discrimination in the small-business credit market. *Review of Economics and Statistics*, 85(4):930–943, 2003.

F. D. Blau and A. H. Beller. Black-white earnings over the 1970s and 1980s: Gender differences in trends. *The Review of Economics and Statistics*, pages 276–286, 1992.

F. D. Blau and J. W. Graham. Black-whites differences in wealth and asset composition. *The Quarterly Journal of Economics*, 105(2):321–339, 1990.

R. Blundell. Income dynamics and life-cycle inequality: Mechanisms and controversies. *The Economic Journal*, 124(576):289–318, 2014.

R. Blundell and I. Preston. Consumption inequality and income uncertainty. *The Quarterly Journal of Economics*, 113(2):603–640, 1998.

R. Blundell, L. Pistaferri, and I. Preston. Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921, 2008.

R. Blundell, L. Pistaferri, and I. Saporta-Eksten. Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435, 2016.

S. Bonhomme and J. M. Robin. Assessing the equalizing force of mobility using short panels: France, 1990–2000. *The Review of Economic Studies*, 76(1):63–92, 2009.

J. Campbell and A. Deaton. Why is consumption so smooth? *The Review of Economic Studies*, 56(3):357–373, 1989.

D. Card and A. B. Krueger. School quality and blacks-whites relative earnings: A direct assessment. *The Quarterly Journal of Economics*, 107(1):151–200, 1992.

D. Card and A. B. Krueger. Trends in relative blacks-whites earnings revisited. *The American Economic Review*, 83(2):85–91, 1993.

K. K. Charles, E. Hurst, and N. Roussanov. Conspicuous consumption and race. *The Quarterly Journal of Economics*, 124(2):425–467, 2009.

K. Y. Chay and D. S. Lee. Changes in relative wages in the 1980s returns to observed and unobserved skills and blacks–whites wage differentials. *Journal of Econometrics*, 99(1): 1–38, 2000.

P. Cheng, Z. Lin, and Y. Liu. Racial discrepancy in mortgage interest rates. *The Journal of Real Estate Finance and Economics*, 51(1):101–120, 2015.

R. Chetty, N. Hendren, P. Kline, and E. Saez. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623, 2014a.

R. Chetty, N. Hendren, P. Kline, E. Saez, and N. Turner. Is the united states still a land of opportunity? recent trends in intergenerational mobility. *American Economic Review*, 104(5):141–47, 2014b.

R. Chetty, N. Hendren, M. R. Jones, and S. R. Porter. Race and economic opportunity in the united states: An intergenerational perspective. *The Quarterly Journal of Economics*, 135(2):711–783, 2020.

A. Deaton. *Understanding consumption*. Oxford University Press, 1992.

G. Dymski and L. Mohanty. Credit and banking structure: Asian and african-american experience in los angeles. *American Economic Review*, 89(2):362–366, 1999.

P. Ganong, D. Jones, P. Noel, D. Farrell, F. Greig, and C. Wheat. Wealth, race, and consumption smoothing of typical income shocks. Technical report, University of Chicago, Becker Friedman Institute for Economics Working Paper, (2020-49), 2020.

R. Hall and F. Mishkin. The permanent income hypothesis: Evidence from panel data. *Econometrica*, 50:461–482, 1982.

J. Heathcote, F. Perri, and G. L. Violante. Unequal we stand: An empirical analysis of economic inequality in the united states, 1967-2006. *Review of Economic dynamics*, 13 (1):15–51, 2010.

J. J. Heckman, T. M. Lyons, and P. E. Todd. Understanding blacks-whites wage differentials, 1960-1990. *American Economic Review*, 90(2):344–349, 2000.

J. S. Heywood and D. Parent. Performance pay and the whites-blacks wage gap. *Journal of Labor Economics*, 30(2):249–290, 2012.

R. A. Hummer and E. M. Hernandez. The effect of educational attainment on adult mortality in the united states. *Population bulletin*, 68(1):1, 2013.

Ò. Jordà, M. Schularick, and A. M. Taylor. Macrofinancial history and the new business cycle facts. *NBER macroeconomics annual*, 1(31):213–263, 2017.

D. Krueger and F. Perri. Does income inequality lead to consumption inequality? evidence and theory. *The Review of Economic Studies*, 73(1):163–193, 2006.

P. Menchik and N. Jianakoplos. Black white wealth inequality: Is inheritance the reason? *Economic Inquiry*, 35(2):428–442, 1997.

B. D. Meyer and J. X. Sullivan. Measuring the well-being of the poor using income and consumption. *Journal of Human Resources*, 38:4, 2003.

R. L. Oaxaca and M. R. Ransom. On discrimination and the decomposition of wage differentials. *Journal of econometrics*, 61(1):5–21, 1994.

J. Peoples and W. K. Talley. Black-whites earnings differentials: privatization versus deregulation. *American Economic Review*, 91(2):164–168, 2001.

T. Piketty and E. Saez. Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, 118(1):1–41, 2003.

G. E. Primiceri and T. Van Rens. Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of monetary Economics*, 56(1):20–39, 2009.

Appendices

Appendix A Data

A.1 Sample Selection

As explained in Section 2, to create our dataset, we append together all waves from 1968-2017. The full PSID dataset contains 1,856,953 individual-years. We limit our sample to the SEO and SRC samples, eliminating households from the Immigrant and Latino surveys. We also include only current heads, since they are the households with the richest and most consistent set of observables overtime. We also create a consistent race indicator for all households. The PSID asked heads to identify their race in every wave. For all heads, we assign race as the mode value of race from all reported years. Due to the limited sample size of some reported races, we only keep households identifying themselves as Black or as White. Our full sample, using all waves of data, includes 457,286 individual-year observations. In our main regression analysis, however, we define a base year of 1981 - so only households present in the 1981 wave and beyond are included. This brings us to 342,679 individual-year observations (from waves 1981-2017).

A.2 Variable Definition

A.2.1 Demographics

Over time, the PSID has altered how it collects educational data. From 1968 until 1990, households reported educational buckets; afterwards education was given in yearly denominations. To create a consistent education status, we assign households to the four categories used by Attanasio and Pistaferri [2014]: 1) 0-11 grades completed, 2) High school degree or 12 grades plus nonacademic training, 3) College dropout (some college), and 4) BA degree or college and advanced/professional degree. Heads report educational information every year in the period 1968-2017. To account for missing data, we generate a new variable for each individual that contains their maximum education status attained.

The PSID allows for five different classifications for marital status: married, single, widowed, divorced, separated, and married with spouse absent. We use these same definitions in our analysis.

Another demographic variable of interest is disability status. The PSID asks households if they have "any physical or nervous condition that limits the type or amount of work" they can do. We use an affirmative answer to this question as an indicator for presence of a disability.

To define employment status, we create a binary variable. households who report that they are working or temporarily laid off are considered employed. The PSID also includes a variable on self-employment. We define as self-employed only those households who report being exclusively self-employed, not those who indicate they are employed by themselves and by someone else.

Information is also present on total hours worked, defined as the total annual hours worked for money on all jobs, including overtime. We replace two wild codes in the data with missing values.

A.2.2 Consumption

For all waves of the survey except 1973, 1988, and 1989, the PSID consistently collects information on food consumption. Starting in 1968, interviewees are asked to provide their *annual* expenditures on food used at home. This value includes the cost of food delivered to the home, but excludes alcohol and cigarette consumption and excludes expenditures from food stamps. Then in 1994, the question switches to a *varying time* unit form. The interviewees themselves choose the time frequency to report at, whether it be weekly, monthly, or yearly. Therefore, we convert expenditures to annual values by multiplying the reported values by the appropriate constant based off the given time unit (i.e. by 12 if the time unit is monthly, by 52 if the time unit is weekly, etc...). Post 1994, if an individual reports $0 spent on food at home, we set their home food expenditures to 0 regardless of the time unit. In addition, food delivery expenditures become a separate variable beginning in 1994, so we have to manually add these values to our measure of food-at-home consumption.

The PSID follows a similar format to collect information on food away from home. With the exception of 1973, 1988, and 1989, households provide the dollar value of annual expenditures spent on food away from home between 1968-1993. Money spent on meals at school or work is excluded. Then in 1994, the question switches to a varying time unit format. We use an identical procedure as described above to convert expenditures to annual values.

Though the PSID asks respondents questions about food stamps in every wave except 1973, they change the wording on the questionnaire. Between 1968-1979 respondents are asked about the amount they saved by using food stamps in the previous year, calculated as the dollar value of food bought with stamps minus the amount spent to purchase the food stamps. Then from 1980 -1993 they are asked about the dollar value of stamps they received in the previous year. In 1993 and in subsequent waves, the PSID also asks about the value of food stamps received, but with a varying time unit. For our own measure, we use the annual values up until 1992, and the time varying values from 1993 on. If an individual reports \$0 received in food stamps, we set their food stamp expenditures to 0. For the year 1993, if the *time-varying* value is missing, we fill it in with the *annual* value. Since the time frame of collection for food stamps does not align with the time frame for collection of other food expenditures (i.e. most food expenditure questions are asked about current food consumption, while food stamps are reported for the prior year) we assign food stamp values to the year of the wave they were collected in.

To create a total food consumption measure, we add together the expenditures for food at home (and food delivery when this is separate), food away from home, and food stamps for each wave.

We noticed the presence of dramatic outliers in total food consumption. These come from later waves of the survey, and we suspect were due to errors in the time unit reporting. For instance, if the correct time unit for food expenditures is monthly, but it is coded as weekly, we would multiply the value by 52 instead of the correct 12 to achieve the annual amount. To correct for extreme outliers, we drop the top 0.1 percentile of food consumption each year.

Our rent equivalent measure combines values for both renters and homeowners. We define homeowners as households who report a non-zero positive house-value. We create a rent equivalent by taking 6% of this house-value. For those who do not report a positive house-value, the PSID provides annual rent payments from 1968-1993. Then for 1993-2017, rent is given in varying time units, defined by the interviewee. To convert these payments to an annual rent, we multiply the reported rent by the appropriate constant based off the given time unit (i.e. by 12 if the time unit is monthly, by 52 if the time unit is weekly). This procedure applies to interview years 1993-2011. In all waves, rent values can be either positive or 0.

This leaves us with some missing values. In 1993, if an individual does not report a positive house-value, but is also missing *time-unit* rent information, we fill in the value of rent given in the *annual* variable when applicable. We are only able to do this in 1993 because this is the only year that includes both the annual and the time-varying rent variables. Furthermore, the PSID includes another variable that indicates the interviewee's self-reported house-status. For households with missing house-values and missing rent information but who self-report that they are not renters or homeowners, we set their rent equivalent to 0. In 1978, some households claim an annual rent of \$768 but also report they are not homeowners or renters. Communication with the PSID indicated that 768 was a wild code in that year. The rent equivalents for these households are therefore also set to 0.

In summary, our analysis includes one measure of rent equivalent. For people with positive house-values, we take 6% of this value. For people without positive house-values, we generate an annual version of their reported rent payments (whether the payments are positive or 0). For households missing information on both house-value and rent payments and who self-report being neither homeowners nor renters, we set their annual rent equivalent to 0.

The PSID asks about amounts paid for utilities such as electricity, water and sewage, gas and other heating fuel, and miscellaneous utilities. We convert all quantities to an annual measure by multiplying the reported value by the appropriate time constant.

PSID transportation variables are all reported at the monthly level. For the month of the interview, respondents are asked how much they paid for parking expenses, gas, bus and train, cab, and other transportation costs. We again annualize these values. households also provide their car insurance payments for all family vehicles per year.

Annual school-related expenses (such as tuition, books, computers, tutors, room/board, uniforms, and other school related expenses) are asked of households regarding the previous year. Families are also asked how much they paid for childcare in the previous year. This question is one of the few consumption measures asked beginning in 1970, but in earlier years it is only asked to families with working female heads or wives. In the waves relevant to our purposes (1999 and on), all families are asked about childcare costs.

We also use various healthcare expenditures in our analysis. For instance, the PSID asks households how much they pay for health insurance premiums for all health insurance coverage in their family. This includes amounts both paid directly and automatically deducted from pay. Furthermore, information is also collected regarding out-of-pocket costs paid for nursing homes, hospital bills, doctors' visits, outpatient surgery, dental bills, prescriptions, in-home medical care, and specialty facilities. Healthcare costs correspond to the prior two-year period, so we divide reported values in half to get annual values.

The final consumption variable we use in our analysis is home insurance. Interviewees provide their total yearly homeowner's insurance premium.

A.3 Missing Data

For all variables, whenever an individual gives an answer of "Don't Know" or "Not Available" (indicated with specific codes in the PSID data), we set this value to missing at first.

Unfortunately, the systematic presence of missing values would eliminate a large number of observations from our consumption imputation. To make sure these observations are still included, for each categorical demographic variable we create a new group to identify households with missing information. For example, all households missing marital status information are assigned the code 999 for the marital status variable, so they are grouped together in the imputation. This procedure applies to marital status, maximum education, state, number of children, employment status, self-employment status, disability status, and homeowner status. For continuous variables (such as age, expenditures, and income), missing values remain missing.

A.4 Other Considerations

One peculiarity about the PSID is the discrepancy that sometimes arises between the year of the survey wave and the year that a variable is collected for. For example, in each interview the PSID asks respondents about their current house value and rent, so these values correspond to the year of the survey wave. The same pattern also arises for food consumption at home - interviewees are asked about their current expenditures on food consumption, so the value corresponds to the year of the survey. However, for some variables the PSID asks respondents about values for the prior year. For example, households report their family income for the year prior to the survey. Food stamp value is also collected for the year prior to the survey. Beginning in 1999 when the PSID includes more consumption measures, this inconsistency continues. Utilities, transportation, and car insurance costs are reported currently, and therefore apply to the year of the survey. Other consumption expenditures, such as education and childcare expenses, are reported for the prior year. In addition, healthcare costs - including drug and hospital costs - are reported for the prior *two* years. Since the time frame that the PSID uses to collect data varies for different variables, we standardize our measures of consumption and income by assigning all values in a particular interview to the year of that survey wave. For instance, all information collected in the 1995 survey wave is assigned as pertaining to the year 1995. This becomes relevant when we adjust our values by the CPI - we use the CPI of the year of the survey wave.

In our regressions, we cluster the standard errors at the family level, using our own definition of family. We consider families to be households where the identity of the head and the wife remain the same (though in the actual regressions only the heads are present). If at any point and time the identity of the head and/or wife changes (i.e. if a couple splits, if a head or wife dies, or if a previously single individual gets married), we consider this to be a new family.

We would also like to note that for all types of analyses involving consumption expenditures, we do not use data from years 1973, 1988, and 1989 since food information was not collected in those surveys. Family income, however, was collected in those years. For all analyses pertaining to income, therefore, we keep years 1973, 1988, and 1989 in order to increase the sample size. One more final consideration is that when values in the PSID are topcoded, we keep the topcoded values. This applies to very few observations.

A.5 Attrition

From 1968 until 1991, the PSID only interviewed households if they had been interviewed in the previous wave. People who could not be found or refused to participate in one year were lost to the survey. However, in 1992 the PSID began an effort to recontact some of these nonresponse households from previous years. Furthermore, starting in 1993, households who were nonresponsive in a particular wave were still followed for the subsequent wave. If an individual remained missing for two waves, they were then dropped. In a similar effort, 1993 marked the year when the PSID began to follow sample children who left their family units before the age of 18 to join a non-sample family. This meant that for the first time, both the head and the wife of an interviewed family could be non-sample. The family just needed one sample member in order to be interviewed, regardless of this member's relational status. Due to budgetary constraints, in 1997 the PSID dropped approximately 25% of its sample families, with reductions made mainly to the SEO subsample.

Appendix B Insurance

B.1 Response to Income Shocks

The evidence provided confirms a large difference in the amount of savings and wealth of Whites and Blacks. This is suggestive of a potential important difference in the level of consumption insurance achieved. More specifically we conjecture that even for households in the same top quintile, Blacks have a lower degree of partial insurance than Whites: having then a more volatile consumption with potentially larger downfalls. In order to test this hypothesis, we adopt the framework developed by Blundell et al. [2008].

The question of how much do income shocks reverberate into consumption shocks has been widely discussed in the literature. According to the complete market hypothesis, consumption is fully insured against any idiosyncratic income shocks. This hypothesis has been usually rejected in micro data (Attanasio and Davis [1996]). On the contrary, the standard permanent income hypothesis assumes that self-insurance through savings is the only mechanism that can be used to smooth income shocks. According to this latter theory, intertemporal consumption is smoothed against transitory, but not against permanent income shocks (Deaton [1992]). However, both aggregate and micro data exhibit what is called "excess smoothness", i.e. consumption is found to react too little to permanent income shocks. Further, consumption data also exhibit excess sensitivity to transitory income shocks (Hall and Mishkin [1982], Campbell and Deaton [1989], Attanasio and Pavoni [2006]. In the light of these studies, Blundell et al. [2008] propose a model in which they start assuming that there is some degree of insurance, which is not necessarily full, and focus on the importance of distinguishing between transitory and permanent shocks (Blundell et al. [2008]).

In line with that work, we disentangle the permanent and transitory income component and we allow the variances of the permanent and transitory factors to vary over time. Further, we assume that the permanent component follows a random walk. Suppose log income, $\log Y_t$ can be decomposed into a permanent component P and a mean-reverting transitory component v. Then the income process for an household i is:

$$\log Y_t = Z'_{i,t}\varphi_{i,t} + P_{i,t} + v_{i,t} \tag{9}$$

where Z is a set of observable income characteristics such as demographic, education, race and other variables. We allow the effect of these characteristics to shift with calendar time and we also allow for cohort effect. The impact of the deterministic effects $Z_{i,t}$ on log income and (imputed) log consumption is removed by separate regressions of these variables on year and year-of-birth dummies, and on a set of observable family characteristics (dummies for education, race, family size, number of children, region, employment status, residence in a large city, outside dependent, and presence of income recipients other than husband and wife). As in Blundell et al. [2008], we then work with the residuals of these regressions. We assume that the permanent component follows the following process:

$$P_{i,t} = P_{i,t-1} + \zeta_{i,t} \tag{10}$$

where $\zeta_{i,t}$ is serially uncorrelated and the transitory component $v_{i,t}$ follows an MA(q) process, whose order is established empirically. We are interested in assessing how income shocks differently transmit to consumption for Blacks and Whites households. We write unexplained change in log consumption as:

$$\Delta c_{i,t} = \phi_{i,t}\zeta_{i,t} + \psi_{i,t}\varepsilon_{i,t} + \xi_{i,t} \tag{11}$$

where $c_{i,t}$ is the log of real consumption net of its predictable components. We allow permanent income shocks ($\zeta_{i,t}$) to have an impact on consumption with a loading factor of $\phi_{i,t}$. On the other hand, the impact of transitory income shocks $\varepsilon_{i,t}$ is measured via the factor loading $\psi_{i,t}$. The random term $\xi_{i,t}$ represents innovations in consumption that are independent of those in income (this may capture measurement error in consumption, preference shocks, etc.). Our aim is to estimate $\phi_{i,t}$ and $\psi_{i,t}$, which are our insurance parameters. In case of full insurance, they would be both equal to zero, whereas in case of no insurance they would be both equal to 1. These parameters are estimated by diagonally weighted minimum distance.

	Whites	Blacks
ϕ	0.7687***	0.7959***
	(0.0650)	(0.1182)
ψ	0.1026***	0.1699***
	(0.0322)	(0.0550)

Table B.1: Degree of partial insurance of Blacks and Whites towards permanent vs transitory income shocks. Bottom 0.5% of consumption has been trimmed.

The parameter ϕ represents the degree of insurance with respect to permanent income shocks, whereas the parameter ψ stands for the degree of insurance with respect to transitory income shocks. In both cases, the lower the value of the parameter, the higher the degree of partial insurance, the smoother the consumption profile and the smaller the consumption responses to both types of income movement. From Table 1, it emerges that

Blacks are less insured than the Whites, both with respect to transitory and to permanent shocks. This is likely to explain the remaining differences in consumption persistence between Blacks and Whites, in particular at the top of the distribution. Different racial degrees of partial insurance are likely to be at the root of racial differences in persistence across the consumption distribution. However, these differences in the estimated coefficients for partial insurance across race are not statistically different from each other, neither for the permanent nor for the transitory shock coefficient. This has been verified by performing 100 bootstrap replications of the estimation presented above. While the parameters are not statistically different from each others, those differences are economically quite substantial in particular for the transitory component. Indeed, a 1 USD temporary shock translates into a 17 cents consumption fall for Blacks, whereas it only translates into a 10 cents consumption fall for Whites. This finding is consistent with those by Ganong et al. [2020], who find that black household cut their consumption on average 50% more than white households in response to an unexpected temporary shock in income. A permanent shock has clearly a much larger impact on consumption, as predicted by the theory, on both Blacks and Whites, with 1 USD of permanent fall in income pushing down consumption by 80 and 77 cents for Blacks and Whites respectively. Notice that the variance of the two components are very similar for Blacks and Whites, confirming the validity of our original assumptions on similar income processes between Blacks and Whites after controlling for a few demographic and labor market characteristics.

	Whites	Blacks
Var of permanent component	0.0395	0.0509
Var of transitory component	0.0432	0.0586

Table B.2: Variance of the permanent and of the transitory income component, by race. Bottom 0.5% of consumption has been trimmed.

This suggests that Blacks and Whites are subject to similar shocks during their lifecycle, but what determines their different degree of positional persistence in the consumption distribution is how they react to these shocks (i.e. drawing from their savings or reducing consumption permanently). Note that the similarity in income variances between Blacks and Whites is not merely a consequence of the modelization adopted, but is instead a feature present in the data. A simple descriptive statistics shows that the overall cross-sectional standard deviation of log wages, which can be considered as a rough measure of income volatility, is equal to 1.28 for the Blacks and to 1.59 for the Whites, i.e. these standard deviations are rather close (households with zero wage have also been included, as having the value of 1, in this computation).

B.2 Sources of Insurance for Whites and Blacks

When hit by a shock, an individual may resort to one or more of three main sources of insurance, i.e. social or government insurance, family insurance and self-insurance.

While the previous section provide some insights on how much insurance is achieved by Blacks and Whites, it doesn't directly decompose the contribution to the final nexus income-consumption mediated by all the possible sources of insurance. We here investigate those different sources. As far as family insurance is concerned, we are not able to precisely estimate how much this channel accounts for in case of a shock for Blacks and for Whites. However, based on some descriptive evidence in our data, we can deduce that Blacks in general have a lower access to this insurance channel. Indeed, Blacks usually have more out-of-wedlock children and they also on average get married more times during their lifetime than Whites. It appears that with multiple and changing family ties the fundamental for informal insurance aren't particularly solid. Just to provide an example, in the top consumption quintile, 20% of Blacks are divorced, whereas only 10% of Whites are. Further, Blacks are less likely than Whites to receive an inheritance and, when they do, the average amount is substantially lower.

Finally, as far as social or government insurance is concerned, it is not straightforward that Blacks have more access to it than Whites. It is plausible that the poorer Black households somehow lack knowledge of the administrative procedures which are necessary to obtain social security transfers, and hence are less likely than the (equally poor) White households to obtain them.

Panel (a) Panel (b)

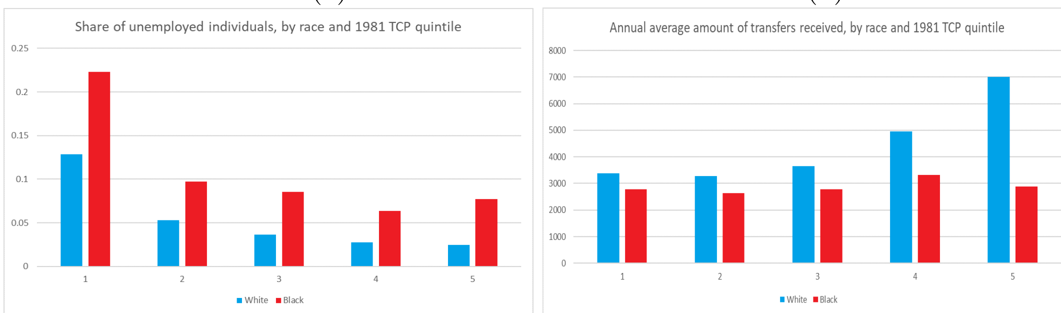

Figure B.1: Panel (a): percentage of people who are currently unemployed and looking for a job, by race and by 1981 TCP quintile. Data for 1981-2017. Panel (b) average amount of transfers that households have received in a year, either from public or private sources, by race and 1981 TCP quintile. Data for 1981-2017.

In Panel (a) of Figure B.1, we report the share of unemployed households, by race and 1981 TCP quintile, in order to assess whether there are relevant racial differences in the probability of being hit by an income shock. From this Figure, we deduce that, in general, the share of unemployed households is higher for the Blacks than for the Whites. In order to obtain an overview of how much support can Blacks and Whites households receive in case they are hit by a shock, in Panel (b) of Figure B.1, we report the average amount of transfers (this time including both private and public sources) received in a year. From this panel we notice that this average amount of transfers is higher for Whites than for Blacks in any TCP quintile, and in particular at the top, where the average amount is more than double for the Whites than for the Blacks.

In order to dig further into the issue of the different degrees of insurance for Black and White households, we analyze whether Blacks and Whites have different degree of health insurance and whether they are differently exposed to health shocks.

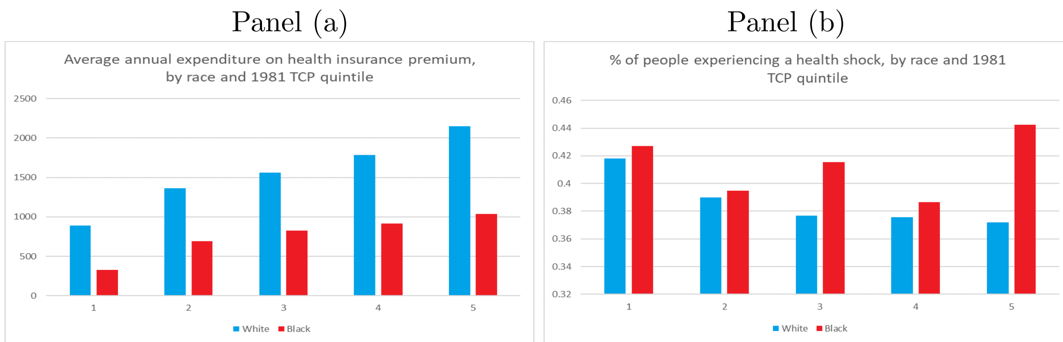

Figure B.2: Panel (a): average annual amount (in US dollars) of expenditure for health insurance at the family level, by race and by 1981 TCP quintile. PSID data for the period 1999-2017. Panel (b): share of households being affected by an health shock. An health shock is defined as insurgence of any of the nine major health problems recorded in the PSID in the period 1999-2017 (asthma, arthritis, cancer, diabetes, lung disease, heart attack, heart disease, stroke, high blood pressure).

As a first exploratory analysis, in panel (a) of Figure B.2 we report the average expenditure on health insurance, by race and by 1981 TCP quintile. This information is available in the PSID from 1999 onward. The amount paid by Whites for health insurance each year is on average way higher than the amount paid by the Blacks and the difference widens in the upper part of the consumption distribution. This is a first supporting evidence to the claim that Blacks are less insured than Whites against health shocks, even when they are at the top of the consumption distribution. Further, in panel (b) of Figure B.2 we find evidence that the Blacks are notably more exposed to health shocks than the Whites, and this is especially true in the top consumption quintile.

As far as other external channels of insurance against income shocks are concerned, there is a large literature on racial differences in credit market access. Just to mention some examples, Arrow [1998] claims that credit market is one of the many aspects in which economic discrimination may manifest itself, and Blanchflower et al. [2003] find evidence that, all other relevant factors being equal, Black-owned small businesses are around twice as likely to be denied credit than White-owned ones. Dymski and Mohanty [1999] further suggests that one of the reasons why Blacks have lower access to the credit market may be that there are fewer bank branches in the urban areas which are mostly populated by Blacks. As an important point, on average, the interest rate paid by Whites on their first mortgage is 5.61%, whereas that paid by Blacks is 5.87%. This is consistent with the findings by Cheng et al. [2015], who claim that Black borrowers on average pay about 29 basis points more than comparable Whites borrowers, even after controlling for mortgages characteristics. Further, Cheng et al. [2015] report that the median mortgage amount for Blacks is 105'000 US dollars, while for Whites is 120'000 US dollars; this is consistent with our hypothesis that Blacks have a harder time in accessing the credit market than Whites. Moreover, Blacks seem to prefer long-term mortgages (30-year-loans) than Whites (71% vs 57.8%, Cheng et al. [2015]). Similarly, on the basis of more detailed data, Bayer et al. [2016] show that African-American and Hispanic borrowers were respectively 103 percent and 78 percent more likely to receive high-cost mortgages for home purchases before the Great Recession, even after controlling for individual credit scores and other risk factors. Moreover, Blacks have been more exposed to foreclosures than Whites during the crisis (Bayer et al. [2017]).

Appendix C Robustness checks

In this section, we discuss a series of robustness checks that we performed. All the graphs relative to the robustness checks are reported in Appendix B.

C.1 Actual Consumption

As far as persistence probabilities are concerned, we re-estimated them by using actual consumption data (for years from 1999 onward) instead of our measure of imputed consumption. In this case, too, the Blacks/Whites differences in consumption persistence at the top of the distribution do not disappear, even after the industry of employment has been included among the controls. Further, we estimate positional persistence in the top consumption quintile by using different sets of control variables. In particular, we exploit geographical information (four US macro-regions), explore non-linearities in the impact of the number of children and investigate the role of household wealth (i.e. estimated house value). However, none of these variables totally closes the gap between Blacks and Whites positional persistence at the top.

C.2 Equivalence Scale

As a further check, we estimate persistence by adopting a different equivalence scale than the one in the main body of the paper. This means that we compute Total Family Income or Consumption by dividing total family income/consumption by an alternative equivalence scale, i.e., the Square Root Scale. The formula applied here is the following:

$$SR = \sqrt{\text{Number of people in the household}} \tag{12}$$

The results of the previous sections are confirmed, in the sense that persistence differences between Blacks and Whites, both at the bottom and at the top of the income distribution, almost disappear once a standard set of explanatory variables is included in the quintile-quintile regression. The behavior of consumption is the same as in the previous Sections, with persistence differences not disappearing at the top of the consumption distribution.