

Student Debt and Racial Wealth Inequality

Marshall Steinbaum*

August 7, 2019

Abstract

This paper analyzes the effect of cancelling student debt on racial wealth inequality using the 2016 Survey of Consumer Finances. It concludes that cancelling student debt reduces racial wealth gaps as measured by both the ratio of white wealth to black wealth at a given wealth quantile, across the wealth distribution, as well as the absolute difference in wealth between quantiles of each distribution. It then discusses why cancelling student debt disproportionately increases the net wealth of black households and thereby reduces racial wealth gaps.

Jain
Family
Institute

*Assistant Professor, University of Utah Economics Department, and Senior Fellow in Higher Education Finance, Jain Family Institute. Gerald Daniels and Louise Seamster gave thoughtful feedback on the previous draft. Thanks to Will McGrew for outstanding research assistance.

1 Introduction

Among borrowers with outstanding student loans, those with the largest balances tend to earn higher incomes than borrowers with lower balances. But minority student loan borrowers tend to have more debt, and take longer to pay it off, than comparable white households. For those reasons, the effect of cancelling student debt on various measures of individual and group-level inequality has been a matter of controversy, especially given presidential candidates’ recent, high-profile proposals to eliminate either all, or a significant part of, outstanding student debt. This paper investigates the effect of cancelling student debt on racial wealth inequality.

It does so by analyzing the Survey of Consumer Finances for 2016, the most recent nationally-representative dataset that gives a picture of the demographics of student debt. For the purpose of these exercises, two candidates’ student debt cancellation proposals (those of Senators Elizabeth Warren and Bernie Sanders) are implemented within the SCF data and compared to the dataset’s baseline wealth distribution, which includes student debt.

Throughout this paper, the measure of racial wealth inequality used is the ratio of wealth at a given quantile of the distribution of white households to the wealth at the same quantile in the distribution of black households.¹ This extends the common racial wealth inequality metric of comparing medians of the two distributions. By examining racial wealth gaps across the wealth distribution rather than solely at the median, we gain a clearer picture of those gaps and improve on the ability to interpret why they arise. We do not consider the ratio of the total amount of “white wealth” to the total amount of “black wealth,” as that type of measurement is invariant to the within-race-group wealth distribution. Relatedly, we do not attempt to divide the “cost” of a given student debt cancellation plan between effective payments made to households grouped by race, since such a measurement is contaminated by the relative sizes of each race’s population of student debtors.

¹In fact, we use a transformation of wealth that is designed to mitigate the influence of very large outliers, as discussed in Subsection 3.2.

Finally, we confine attention to white and black households, ignoring those of other races, because the SCF is not large enough to construct a full wealth distribution for other race groups, especially when we narrow the sample to younger cohorts and to only those who have outstanding student debt.

Our conclusion is that *racial wealth inequality is reduced more the more student debt is cancelled*. With respect to the two presidential candidates' plans, this means that the Sanders plan, completely eliminating outstanding student debt, reduces racial wealth inequality more than does the Warren plan, which only forgives \$50,000 of debt, and phases that out for high earners.² But the difference between the two plans as measured by the reduction in the racial wealth gap is not large. It would be fair to say that the Warren plan achieves the vast majority of the racial wealth equity gains that the Sanders plan achieves, while leaving the student debt held by the highest-income borrowers intact. The Warren plan would achieve substantially all of the gains in racial wealth inequality as the Sanders plan if the forgiveness amount were uncapped (rather than \$50,000), but retained the high-income phaseout that is its other distinguishing characteristic.

The reason why student debt cancellation is, in general, a racially egalitarian policy is that black people have much more student debt for a given level of income or wealth than do white people, and therefore black people hold a share of outstanding student debt that is disproportionate to their share of the population. This holds true no matter how far up the wealth distribution you go, because among wealthier people, white people are less likely to have any student debt at all than black people. But the dynamics of racial wealth inequality at the top of the wealth distribution are not driven by student debt, as the analysis in Section 4 shows. Student debt matters much more for the non-wealthy. At the bottom and middle of the wealth distribution, black borrowers are likely to have more student debt than white borrowers conditional on income or wealth, which means that Warren's plan, which leaves intact student debt over \$50,000 (or more, for households that earn more) leaves more

²See Section 4 for a full description.

debt outstanding for black households than for white households, holding constant household income. However, the most important direct effect of cancelling student debt, vis-a-vis racial wealth inequality, is to increase the share of households with a positive amount of wealth, and of those households who begin with a small amount of wealth, to increase it by a large relative proportion. Black households are over-represented among such households.

There are a variety of reasons for the disproportionate representation of black borrowers among all student debt holders, of the importance of student debt to the overall balance sheets of black households, and for the importance of student debt to middle class black families relative to white families. Exploring and evaluating all of them are beyond the scope of this paper, but important elements of the story are that white households have much greater wealth in their family background and so do not need to go into as much debt to acquire higher education or to obtain secure financial status, that white workers are more likely to enter higher-paying jobs more quickly after obtaining a given level of education, and as a result, black workers acquire more credentials, and therefore more debt, to obtain a given job and a given level of income. In some ways, the central question at issue is not “who has student debt?” but rather “who is able to live their economic lives without it?” The answer to this last question is, increasingly and for younger cohorts for which contact with higher education is becoming a near-universal experience, white households who can obtain higher education credentials without going into debt.

This paper proceeds as follows. Section 2 reviews what literature exists about student debt cancellation and the racial wealth gap. Section 3 explains the dataset, including the construction of three different samples for which the analysis is conducted, as well as the methodology used to evaluate racial wealth gaps. Section 4 presents the counterfactual results of cancelling student debt according to the Sanders and Warren plans in each of the three samples, as against a status quo baseline. Section 5 explains why student debt cancellation is racially egalitarian, in light of the results and other studies. And section 6 sums up.

2 Literature Review

Any consideration of the racial distribution of student debt indicates that black households suffer disproportionately from their larger amount of student debt for a given level of income or wealth than do white households, who tend to have lower student debt for a given level of income or wealth. On its face, then, eliminating student debt would close the gap between the net wealth position of black households and white households. Notwithstanding that, the notion that student debt cancellation might be inegalitarian, including along racial lines, has obtained wide currency in the popular debate over student debt. This paper shows that that supposition is incorrect.

The closest analog to the results reported here appears in Huelsman et al. (2015), which models full and partial student debt forgiveness scenarios in the 2013 Survey of Consumer Finances. It reports that partial forgiveness, and forgiveness directed at lower-wealth households, would reduce the racial wealth gap, but that total forgiveness for all households would increase it.

In fact, the results reported in that paper are not fundamentally different than the ones reported here. Three things account for those authors' different presentation of their results, however: **first**, their measure of the racial wealth gap and changes in racial wealth inequality is different. Those authors measure the racial wealth gap as the difference in the number of dollars between the wealth of the median white household and the wealth of the median black household, while this paper measures the racial wealth gap as the ratio between wealth at a given quantile of the white household wealth distribution and wealth at that same quantile of the black household wealth distribution. Appendix A utilizes the absolute wealth gap measure along the lines of Huelsman et al. (2015).

For those authors, if median white wealth goes up by more dollars than median black wealth as a result of a cancellation counterfactual, the racial wealth gap has widened. Whereas in the ratio measure, since black households have (much) less wealth across the board, an increase in wealth for black households tends to narrow gaps. This contrast be-

tween absolute and relative measures of inequality is, of course, not unique to this paper, and there is no one right answer on how to measure inequality, including racial wealth inequality. We adopt the relative approach (i.e., we report wealth ratios) because that is the more common way that racial wealth gaps are presented in the popular debate, e.g. “white families have ten times as much wealth as black families,” as well as in the academic literature. In general, relative measures of inequality highlight the social implications of that inequality more so than absolute ones. The fact that the imbalance between white and black wealth is so large in relative terms is, fundamentally, why racial wealth inequality signifies historical disadvantage and exploitation. In fact, Huelsman et al. (2015) do discuss wealth ratios as measures of economic and social inequality, though their quantitative analysis does not utilize them.

The **second** difference between this paper and Huelsman et al. (2015) is that we report racial wealth gaps for every quantile of the wealth distribution (for black and white households), while they report gaps measured only at the median.

The reason why it matters to report the full distribution of wealth effects, and not just the median, is not simply that it gives a more complete picture of racial wealth inequality, though it does, and that is important for understanding *why* student debt cancellation is egalitarian. In addition, studying the entire wealth distribution highlights the importance of wealth dynamics around the threshold of zero net wealth, which matters a lot for measures of racial wealth inequality. Specifically, cancelling student debt puts more households (of all races) above the threshold of zero wealth. Depending on the sample used (discussed in Subsection 3.1), that threshold can occur around the median wealth of either white or black households, which means that the behavior of a wealth ratio at the median is “contaminated” by a spurious outlier. This is explained further in Subsection 3.2, and is evident in the results reported in Section 4.

The richer picture of net wealth we can paint from comparing all wealth quantiles, not just the median, gives rise to richer interpretations of the role that student debt plays in

the economy and in racial wealth inequality. We will show that student debt matters less to household net wealth positions the higher you get up the wealth distribution (although, as stated above, it always matters more for black households than white households). The significance of this is that student debt matters a great deal for households in the broad vicinity of the middle class, and it matters especially to black households because it is constitutive of their ability to attain and retain that status. That interpretation has the same general flavor as the findings in Steinbaum and Vaghul (2016)—that controlling for income, the strongest positive correlation between black population share and student debt delinquency is found in middle-class zip codes, presumably because middle class black people are so much more burdened by their student debt than are middle class white people.

Finally, the **third** difference between the findings in this paper and Huelsman et al. (2015) is simply that they study the 2013 SCF and we study the 2016 SCF. As Morgan and Steinbaum (2018) show, student debt is moving down the income (and wealth) distribution over time, and is already much more equitably distributed than income or wealth. The more recent dataset will include more debt held by lower-income people, both because lower-income people are increasingly taking on debt, but *especially* because they are having more trouble paying it off (Looney and Yannelis, 2015). This dynamic means debt-to-income and debt-to-wealth ratios are dramatically higher among lower-income people, in contrast with the rhetoric in the popular debate about how student debt is a problem affecting the rich or upper middle class. The trickling down of student debt as a social phenomenon to the broad middle and working class, and the consequent non-repayment of student debt, is more apparent over twenty years than over the three years between 2013 and 2016, but it matters nonetheless to the distributional impact of these debt cancellation counterfactuals, including by race.

To conclude this review, a large literature focuses on black students' and black borrowers' experience with student debt and its implications for their financial status. We will not review the entirety of it. But the paper from that literature closest to this one, in terms of both

methodology and data, is Kakar, Daniels and Petrovska (2019). Those authors analyze the combined 2013 and 2016 SCFs in order to investigate the effect of student debt on wealth, by race. They conclude that student debt contributes to the racial wealth gap between white and black households (which they measure as a ratio, in effect). The flip side, of course, is that cancelling student debt would narrow the black-white wealth gap.

3 Data

We use the Survey of Consumer Finances (SCF) for 2016, a nationally-representative survey of households for their financial status, including their student debt holding and repayment behavior, as well as demographic and economic characteristics.

The SCF has shortcomings as a dataset for analyzing the prevalence and distribution of student debt. It is designed to survey households, and a growing share of individuals do not live in what the SCF considers to be economically independent households. For that reason, the gap between the total amount of student debt in the SCF and in the most comprehensive alternative data source, the Consumer Credit Panel constructed from credit reports by the Federal Reserve Bank of New York, is getting wider over time (see Figure 1). Researchers involved with constructing the SCF have noted the non-comprehensiveness of the survey in comparison with the CCP and speculated about its causes. One representative discussion includes

“The published SCF statistics refer only to the debts and assets of this economic core. Thus, student loan information is not collected for members of the household that are outside of the household economic core. It is likely that most of the student loans of these non-core household members are included in G.19 and CCP statistics [...] A sampling frame is a list of all eligible housing units for a survey, and the frame for the SCF excludes institutional residences, including college dorms [...] Thus, the differences between the SCF and CCP and G.19 aggregate levels suggests that a considerable portion of aggregate student loan debt is

held by individuals outside of the economic core of a household” (Bricker et al., 2015).

It is likely that the student debt “missing” from the SCF is overwhelmingly held by low-income, low-wealth individuals, and therefore that if anything, counterfactual analyses of the SCF understate the degree to which debt cancellation would be racially egalitarian or egalitarian in any dimension.³

Moreover, there is reason to believe that the decline in household formation that makes the SCF a less-than-comprehensive source of data is itself caused by the rising prevalence of student debt among younger cohorts, which makes it more difficult for them to achieve economic independence (Dettling and Hsu (2013); Morgan and Steinbaum (2018); Bruenig (2019a)).

Finally, it should be noted that it’s not just student debt that’s missing from the SCF, but poorer and non-white people altogether, who are less likely to live in households that are included in the SCF sample frame. For example, housing instability, incarceration, cohabitation without economic dependence, and any number of other “non-traditional” living situations would mean someone is not included in an SCF household and therefore does not appear in our household wealth distributions. The further down the wealth distribution you

³Looney (2019b) claims that the student debt not reported in SCF is likely both in high-income and low-income households, so the SCF overstates the degree to which student debt is held by the middle class versus rich or poor households. But this analysis is flawed in and of itself and otherwise not replicable, given his reference to a privately-held administrative dataset “adjusted” to replicate the demographics of student debtors. Looney (2019b) presents the Looney and Yannelis (2015) estimate that shows a debt distribution of 9%, 15%, 23%, 26%, 27% for each income quintile for 2013. But in Table 4 of Looney and Yannelis (2018), he presents what appears to be the same analysis and same data set, but for 2012 data. He reports debt shares by quintile of 15%, 15%, 19%, 23%, 27%. The median income reported for each of these quintiles is very close in each of the two studies. It seems highly implausible that the “true” distribution of student debt shifted dramatically up from the bottom quintile to the third and fourth quintiles between 2012 and 2013, but that is what his data suggest. Moreover, in comparing the administrative data to student debt by quantile in SCF in Looney (2019b), the median income for each quintile (in 2013) is \$394; \$18,719; \$32,874; \$56,529; and \$113,603. These income quintiles are adjusted to correspond to the “age and gender of borrowers.” But the comparison appears to be with unadjusted SCF data. In that case, the median income for each quintile is \$15,200; \$30,400; \$48,700; \$77,900; \$121,700 for 80th to 89.9th percentile, and \$223,200 for the 90th to 100th percentile). The two distributions might have the same general shape, but the distribution from the administrative data is less dispersed and centered way to the left of the unadjusted SCF data. In fact, the first two quintiles of the administrative data nearly fit entirely into the first SCF quintile, meaning the unadjusted bottom quintile from the administrative data might hold about a quarter of the debt. If so, that would confirm that what the SCF is missing is debt held by low-income individuals. But, again, it’s hard to interpret the administrative check that Looney (2019b) conducts.

go, the more missing data there is, and this can give an incomplete picture of the wealth distribution.




Figure 1: This plots the total amount of student debt outstanding in both successive Surveys of Consumer Finances and the New York Fed’s Consumer Credit Panel (drawn from credit reporting data.) The latter indicates a larger amount of student debt outstanding, and the gap is getting wider, probably because declining household formation rates mean that a smaller and smaller share of outstanding debt is detected in independent households. Declining household formation may itself be due to the rising prevalence of student debt. This chart first appeared in Morgan and Steinbaum (2018).

Notwithstanding these shortcomings, there is no comparable publicly-available dataset that enables analysis of the demographic breakdown and consequences of student debt across the wealth distribution, including counterfactuals. So, in effect, the SCF is our only option.

3.1 Three Samples

Various analyses of the distributional consequences of student debt utilize different subsamples of the SCF. For example, Looney (2019a) confines attention to households with student debt, and Morgan and Steinbaum (2018) focus on “young” households—those whose heads are between the ages of 25 and 40—who may or may not have student debt. Here, we construct three samples in order to ensure the findings on racial wealth inequality are not an artifact of focusing on the sub-populations among whom the student debt crisis is most acute.

Sample 1 is those households earning over \$1000 in income annually. This effectively rules out full-time students, as well as others who may be out of the labor market, whether voluntarily or involuntarily. The practical effect of this restriction is to remove what would otherwise be the lowest-earning quantile from the analysis, and student debt is, not surprisingly, very high as a share of current income in that group. That makes non-earning-but-indebted sub-sample an unrepresentative group within which to analyze the impact of debt cancellation. Those households remaining in Sample 1 constitute the significant majority of the population (of households, keeping in mind the SCF’s non-representativeness of the population as a whole). Thus, Sample 1 includes many households with relatively little experience of student debt, and for whom student debt cancellation would matter little for household wealth.

Sample 2 contains households earning more than \$1000, whose heads are between the ages of 25 and 40. This sample is thus comparable to the one analyzed in Morgan and Steinbaum (2018). Student debt disproportionately weighs on younger households, both because the structure of the policy is that it is (or “should be”) repaid within ten years of graduation, and because the enormous increase in the overall stock of student debt in the economy is a relatively recent phenomenon, as is the slack labor market that gave rise to the student debt crisis. Older households were better able to enter the labor market without incurring student debt (since tuition used to be much lower, and many more jobs could be

had without any kind of higher education credential), and having entered, they’ve had longer to pay it off, and also not lived their entire working lives in the slack conditions that have obtained since the end of the labor market boom of the late 1990s.

Thus, those households remaining in Sample 2 would benefit proportionately more from debt cancellation than the larger set of households in Sample 1. It should also be said that younger cohorts are less white.

Sample 3 consists of households earning more than \$1000, whose heads are between the ages of 25 and 40, and who have a positive amount of student debt outstanding. This last restriction, obviously, further selects a sample based on exposure to student debt and therefore one that benefits more from its cancellation. Since student debt is disproportionately held by black households, Sample 3 is also less white than Sample 2 or Sample 1.

3.2 Transforming the Wealth Data

Wealth distributions are difficult to analyze because they are so skewed. In addition, wealth *ratios* can be complicated by the reality of negative net wealth, or such low net wealth that ratios register as very large (negatively or positively) even where the real differences in wealth they represent are in fact small. In counterfactual analysis, changes in wealth across the threshold of zero net wealth become meaningless.

For all of those reasons, we transform the SCF’s measure of net household wealth using the Inverse Hyperbolic Sine function (“asinh”), which is commonly used in the literature on wealth inequality (Thompson and Suarez (2015) Kakar, Daniels and Petrovska (2019)). Inverse Hyperbolic Sine operates like a logarithm (an order-preserving transformation that mitigates extremely high values), but unlike logarithms, it is defined on negative numbers and equals zero at the value zero. This permits us to reasonably compare positive and negative values of wealth.

Figure 2 plots the value of the Inverse Hyperbolic Sine function on $[-20, 20]$. While transforming wealth data in this way makes changes in wealth (and wealth ratios) legible

when they take place at small amounts of wealth, including negative wealth, the transformation still exaggerates the significance of the zero wealth threshold. In particular, for our purposes, a wealth ratio between a positive amount of wealth and negative amount of wealth, or between a positive amount of wealth and a very small amount of wealth, either positive or negative, is still going to appear disproportionate to a wealth ratio between two larger wealth observations. For that reason, in all of our wealth ratio plots in Section 4, there is a spike around the zero wealth threshold. As a result of student debt cancellation (i.e., an across-the-board increase in wealth), that spike will move down the wealth distribution because a larger share of households are in a positive wealth position.




Figure 2: Plot of the Inverse Hyperbolic Sine function on the domain $[-20, 20]$. For any two arguments, the difference between the function of each argument is large relative to the difference between the arguments when they are near zero and when one is positive and the other negative. This is why plots of wealth gaps show humps for quantiles in which white households have a relatively small, positive amount of wealth and black households have a small negative amount of wealth, or where both are positive but black household wealth is very small.

4 Results

The analysis consists of comparing the wealth distributions for white and black households in the status quo versus following the implementation of each of the Warren student debt cancellation plan and the Sanders student debt cancellation plan.

The Warren plan cancels a maximum of \$50,000 of outstanding student debt for households with outstanding student debt. A household with less than that amount outstanding will have all of its student debt eliminated, and a household with more than \$50,000 outstanding will retain the balance. That cancellation amount of \$50,000 phases out for households earning more than \$100,000, at a rate of \$1 of reduced cancellation amount per \$3 of income. Therefore, a household earning more than \$250,000 will have none of its student debt cancelled. To implement the Warren plan, we add the amount of student debt cancelled, a household-specific number, to household net worth to arrive at counterfactual household wealth under the Warren plan.

The Sanders plan cancels all outstanding student debt. Therefore, to arrive at the counterfactual wealth distribution following implementation of the Sanders plan, we simply add outstanding student debt to status-quo household net worth.

All household wealth observations are transformed using the Inverse Hyperbolic Sine function, as described in Subsection 3.2.

We present the results graphically, in three types of distributional chart per sample.

4.1 Sample 1

Recall that Sample 1 is the largest sample, including all households earning more than \$1000 of income annually. That means it covers the whole age distribution and includes households who may have no experience of student debt or of higher education.

Figure 3 depicts racial wealth inequality in the form of a traditional empirical Cumulative Distribution Function (CDF). What that means is that the horizontal axis plots household

wealth, and the vertical axis plots population share. A point on the graph signifies that (for households in that race group), the share with less than or equal to a given amount of wealth is plotted on the vertical axis.

The fact that the CDF for black household wealth is uniformly to the left and above the CDF for white household wealth signifies that a racial wealth gap exists at all centiles of the wealth distribution.⁴ For example, a white household with \$100,000 of wealth is less wealthy relative to other white households (a lower centile in the white household wealth distribution) than a black household with \$100,000 would be. For this kind of chart, it is important to remember that a higher value, and a higher plotted line, is “worse” (less wealthy) than a lower value.

We label dollar values of wealth on the horizontal axes, even though what is actually plotted in the graph is the inverse hyperbolic sine of a given level of wealth. Thus, the values on the horizontal axis are not to scale. That is necessary to make the distance between the two wealth distributions legible.

What we are interested in in order to evaluate the effect of student debt cancellation on racial wealth inequality is the horizontal distance between the white and black wealth distributions: does it get smaller or larger when student debt cancellation is implemented? Is the impact different for different parts of the wealth distribution, or for different versions of student debt cancellation? It can be hard to tell that just by eyeballing these empirical CDFs, so we present two more charts.

Figure 4 plots inverse CDFs, which means that now, the dollar values of wealth are on the vertical axis and the empirical population share with less than or equal to that amount of wealth is on the horizontal axis. In this depiction, higher lines are “better”, and when student debt is cancelled, everyone (who has student debt) is wealthier, so the distributions shift up.

In the chart, the status quo, Warren, and Sanders scenarios have different color schemes,

⁴In fact, the poorest few centiles of white households are in a larger negative position than the poorest few centiles of black households.




Figure 3: CDFs of the wealth distribution in the status quo and under the two candidates' student debt cancellation plans, by race.




Figure 4: Inverse CDFs of the household wealth distribution, by race, for the status quo, Warren, and Sanders debt cancellation plans.

so the idea for measuring racial wealth inequality is to examine whether the distributions for white and black households get closer together when student debt is cancelled. What the chart shows is that as all the distributions shift upward with successively larger cancellation policies, the two distributions for a given scenario also get closer together. The wealth distributions for white and black households are closer together in the Warren scenario than in the status quo, and closer together in the Sanders scenario than in the Warren one. But if we think of the Sanders scenario as the largest possible cancellation policy, then we can see that the Warren scenario achieves most of the gains in terms of racial wealth inequality that the Sanders plan achieves. The difference between the household wealth distributions in the bottom ventile (centiles 1-20) in the Sanders versus the Warren scenarios indicates why it is that the Warren proposal doesn't quite match the reductions in racial wealth inequality in the Sanders proposal: because it leaves debt over \$50,000 intact, and there are relatively poor households with more than that much student debt. The fact that that gap is larger for black households than for white households (again, comparing the distribution lines for the Sanders and Warren plans) indicates that there are relatively more such highly-indebted, poor black households than there are such highly-indebted, poor white households. There's almost no difference between the Sanders and Warren plans above the 20th centile.

Overall, the gains in racial wealth inequality relative to the status quo overwhelmingly occur in the bottom half of the wealth distribution, for households in the vicinity of zero wealth in the status quo. Wealthier households are less affected by student debt, and so their wealth increases relatively less when student debt is cancelled.

Figure 5 shows this in a different way: the wealth ratio, or the ratio of white household wealth to black household wealth, at a given wealth centile. Now, instead of plotting (a transformation of) dollar values of wealth, we plot the ratio, which is the difference between the inverse hyperbolic sines of wealth for each race group at each centile. The vertical axis labels (as opposed to the plots) exponentiate the difference between the inverse hyperbolic sines of the wealth levels for each race group. For example, the chart shows that the wealth




Figure 5: The ratio of white to black wealth across the wealth distribution in the status quo, Warren, and Sanders cancellation plans.

ratio is about 20 at the 20th centile for both the Warren and Sanders plans. This means that the white household at the 20th centile of the white household wealth distribution has approximately 20 times the wealth that the household at the 20th centile of the black household wealth distribution has under the two cancellation scenarios. Under the status quo, that ratio is much higher.

The major thing to notice about this chart is that the plot of the wealth ratio shifts down when cancellation is implemented, and it shifts down more under the Sanders plan than it does under the Warren plan. That is how we can conclude that student debt cancellation reduces racial wealth inequality.

The second thing to notice is the spike in wealth ratio that occurs around the threshold of zero net wealth. As discussed above, that threshold is significant for measuring wealth inequality, and canceling student debt pushes that threshold down the wealth distribution, since a larger share of households are pushed into positive territory when their student debt is cancelled. Thus, in figure 5, the spike both shifts left, to lower wealth quantiles, and narrows (in the sense of occurring over fewer wealth centiles), because a larger share of households are further away from zero.

In addition to narrowing gaps between black wealth and white wealth, the major effect of cancellation is to push disproportionately many black households into positive net wealth territory, which significantly reduces racial wealth inequality. This is why the “width” of the spike narrows following debt cancellation.

Examining the whole of the wealth distribution at once is important to gain an accurate picture of the effect of debt cancellation for exactly this reason. Looking at Figure 5, it might be tempting to say that debt cancellation worsens racial wealth inequality, at least for some households, because, for example, households near the bottom of the distribution go from a small-but-positive wealth ratio to an astronomical one following debt cancellation, given the spike near the zero wealth threshold. But seeing that spike move from a higher quantile to a lower one, as well as narrowing, tells us what’s really going on: a large share of

households have been moved from either negative wealth or near-zero wealth into positive net wealth territory, and this dynamic matters disproportionately for black households, of whom there are many around the threshold of zero wealth. So to conclude that racial wealth inequality is worse for the poorest households when cancellation is implemented is to ignore, or worse, misinterpret, the overall dynamic that is exactly what benefits black households about student debt cancellation.

Finally, and as is also evident from figure 4, cancelling student debt has little effect on racial wealth inequality in the top half of the wealth distribution where, again, student debt matters less to the overall household wealth portfolio. For this sample, that top half of the wealth distribution consists of older households who probably had little experience of student debt, since (if they attended college at all) it was in the era when tuition was far lower, and they've had longer to pay it off.

4.2 Sample 2

As we move from Sample 1 to Samples 2 and 3, we are confining attention more and more to households who are likelier to be affected by student debt cancellation, first younger households, then younger households who actually have student debt.

The dynamics analyzed in Section 4.1 are essentially replicated for the more constricted samples, but in exaggerated form: larger changes in relative household wealth (because the households remaining in the sample tend to be less wealthy), narrowing gaps between the distribution of white household wealth and black household wealth, the shifting of the zero wealth threshold down the wealth distribution, and declining wealth ratios, the larger the amount of student debt cancelled.

In Figure 6, the wealth distributions are on the whole shifted to the left relative to the plots in Figure 3, because Sample 2 is less wealthy than Sample 1. The gains in wealth depicted in Figure 7 are relatively larger than they were in Figure 4, and the wealth gaps consequently narrow more than they did for Sample 1 when student debt is cancelled. There-




Figure 6: CDFs of the wealth distribution in the status quo and under the two candidates' student debt cancellation plans, by race, for Sample 2.

fore the wealth ratios plotted in Figure 8 change more: the zero-wealth threshold starts at a higher wealth centile and shifts a larger distance to the left, and the range at the higher wealth centiles where student debt cancellation is relatively insignificant to wealth ratios is attenuated, since fewer households in that category are retained in Sample 2 relative to Sample 1. With those households gone, a gap appears between the status quo wealth ratio and debt cancellation scenarios extending to higher wealth centiles (up to around the 80th centile, rather than the 50th), meaning that middle-wealth households from Sample 1 are now relatively high-wealth households in Sample 2, and just as before, student debt cancellation narrows wealth gaps for those households.




Figure 7: Inverse CDFs of the household wealth distribution, by race, for the status quo, Warren, and Sanders debt cancellation plans, for Sample 2.




Figure 8: The ratio of white to black wealth across the wealth distribution in the status quo, Warren, and Sanders cancellation plans, in Sample 2. The gap between the status quo and debt cancellation scenarios now extends to approximately the 80th percentile, signifying that this sample is less well-off than the earlier one, because it has excluded older, and therefore wealthier and less indebted households.

4.3 Sample 3

Moving from Sample 2 to Sample 3 further focuses the analysis on households likely to be affected by student debt cancellation: those who are both young and have outstanding student debt. And what we see in analyzing that sample is, once again, relatively wealthy households have been removed by the sample restriction and so the effect of debt cancellation on racial wealth inequality moves “up” the remaining wealth distribution.

Figure 9 shows the CDFs of the wealth distribution in the status quo and cancellation scenarios for Sample 3. They are shifted to the left relative to Figures 6 and 3. Figure 10 shows that cancellation shifts the distributions relatively further in this narrower sample, and that it also closes the gaps between white and black wealth proportionally more. That is confirmed by Figure 11, which shows the zero wealth threshold shifting down the wealth distribution a larger amount. It also shows the narrowing of racial wealth gaps extending to nearly the very top of the wealth distribution. This last finding is not surprising given that all of the households in this sample are affected by cancellation, since they all have outstanding student debt.

An important component of Figure 11 is its relevance to only focusing on comparing medians, as opposed to all centiles of the wealth distribution. If we considered only the median, we would report an enormous decline in the ratio of white household wealth to black household wealth. But that is due to the fact that the zero wealth threshold is in the vicinity of median wealth in this narrower (and poorer) sample. In fact, once again, we see that the pattern is that when focusing on households most affected by student debt (i.e., those who have student debt outstanding), the effect of debt cancellation is proportionally greater, in part because the zero wealth threshold shifts further down the wealth distribution.

In fact, that is the **only** way the Sanders and the Warren plans differ from one another according to this empirical measure (wealth ratios by quantile). That is because the Warren plan is capped at \$50,000 of debt cancellation, and there are some households with more than that amount of debt outstanding who have a household net worth in negative territory or near




Figure 9: CDFs of the wealth distribution in the status quo and under the two candidates' student debt cancellation plans, by race.




Figure 10: Inverse CDFs of the household wealth distribution, by race, for the status quo, Warren, and Sanders debt cancellation plans, for Sample 3.




Figure 11: The ratio of white to black wealth across the wealth distribution in the status quo, Warren, and Sanders cancellation plans, in Sample 3. The gap between the status quo and debt cancellation scenarios now extends to nearly the top of the wealth distribution, which is not surprising since this sample is constructed so that all of its households would be affected by debt cancellation to some extent.

zero wealth. The Sanders plan helps them more, because it forgives all their debt. And that reduces racial wealth inequality more than does the Warren plan, because those households are disproportionately black. All of this showed up in the lowest ventile in Subsection 4.1; now those gaps between the Warren and Sanders scenarios appear wider, and they extend further up the wealth distribution, to approximately the 40th centile. Since the gaps are larger, it's more apparent now that the cap matters more for black household wealth than white household wealth. But fundamentally, what's going on here is the same as what was going on in 4.1: poor households with more than \$50,000 in outstanding student debt get more out of the Sanders plan than they do from the Warren plan, and those households are disproportionately black.

There are two ways that the Warren plan differs from the Sanders plan on paper: that the cancellation amount is capped, and that cancellation phases out for higher earners. It is likely that the Warren plan would achieve **all** of the racial wealth inequality reduction that the Sanders plan achieves, were it to uncapped the debt cancellation amount, but retain or even tighten the phase-out.

5 Discussion

The analysis in this paper contrasts with the received wisdom on the impact of student debt cancellation on racial wealth inequality, and more generally with the conventional wisdom on the distributional impact of student debt cancellation along dimensions other than race.

Much of the criticism of such plans has contended that since higher-income people have more student debt, debt cancellation is a regressive policy. This narrative is false. As we move from Sample 1 to Sample 2 to Sample 3, we are confining attention more to households likelier to be affected by student debt cancellation, by construction. Figures 3, 6, and 9 show that that entails taking wealthier households out of the analysis, because the wealth distributions shift to the left between the three samples. The set of households with student debt is a

relatively disadvantaged, rather than a relatively advantaged, subset of the population. This paper does not analyze the time series of student debt distributions, but in combination with prior work, we can say that the distribution of student debt has been shifting down the distribution of income and wealth: those with lower economic status account for a larger and larger share of outstanding student debt, over time, because 1. they have to take on more of it than they once did in order to enter the labor force (i.e., more than zero), and 2. because they have more trouble paying it off, as wages have stagnated in the labor market and repayment periods have lengthened. Whereas once student debt was a marker of relative privilege in the labor market, when it was confined to high-status graduate degree holders, it is increasingly a mark of relative lack of privilege. This analysis confirms that is the case for the 2016 SCF.

Furthermore, Figures 5, 8, and 11 show that the share of each distribution in which white-black wealth ratios don't change as a result of cancellation starts at the top half of it, then is confined to the top ventile, and finally disappears almost altogether. The idea that debt cancellation worsens racial wealth inequality because it is a giveaway to rich white people could not be further from the truth. *The more debt cancellation there is, the more racial wealth inequality is reduced*, at least among the plans that have actually been proposed. It might be possible to construct a student debt cancellation plan that worsened racial wealth inequality, by explicitly or implicitly barring black people from it, but such a plan is not actually on the table.

Student debt matters less to the wealth portfolios of richer people, in contrast with the rhetoric of regressivity and college being “worth it” due to the lifetime college earnings premium. The fact is that high-wealth households get to be where they are not because they took on student debt in order to obtain valuable degrees, but because they **didn't** take on student debt to obtain valuable degrees and high social and economic status, or if they did, that debt has already been paid off. This is the sense in which Bruenig (2019b)'s analysis of the distributional impact of student debt cancellation is flawed: the set of households with

outstanding student debt is not an arbitrary delineation of households, relative to those households who didn't go to college or who paid off their debt. In fact, the set of people with outstanding student debt well approximates the set of people most harmed by the policy failure of the vast expansion of student borrowing against a credentializing labor market with stagnant wages and a last in-first out hiring dynamic that disadvantages people of color. This, fundamentally, is why student debt cancellation is racially egalitarian.

Confronting the policy failure of the expansion of student debt, and the shift that it embodies in the cost of job training from employers, institutions, and states and onto individuals, requires interrogating the skills gap ideology that motivated the formation, and subsequent expansion, of the federal student loan program in the first place. Starting in the late 1960s, the federal government began to adopt the thinking behind the economic concept of “human capital” into public policy: that the role of education, especially higher education, is for individuals to “invest” in their future earning power by obtaining skills in their youth that raise their value to employers over the course of their careers. (Berman and Stivers (2016); Mitchell (2019); Steinbaum (2017)). According to such thinking, the role of public policy in higher education is, if anything, to complete the capital market: to ensure that individuals can finance their upfront investment in human capital, at a time in their lives when they lack the ability to collateralize such an investment. Thus, the government has a role, albeit a limited one: to guarantee private loans (and much more recently, to act as the direct lender), through which individual students can obtain access to the skill-formation that is the core function of higher education.

It's important to recognize this conception of what higher education policy aims to accomplish post-dates the expansion of public higher education that occurred following the passage of the GI Bill, as well as state-led efforts to build out university systems, such as the California Master Plan of the early 1960s. Those policies radically expanded access to higher education, and in doing so, they opened it up to a student population that looked very different from state tax bases. This dynamic of racial division came to a head in the

public mind with the campus protests for civil rights and against the Vietnam War that characterized the late 1960s.

As a result of that, the expansion of institutional support for higher education out of public budgets appeared to have introduced societal dislocation and conflict onto campuses, along with the influx of new students. Shifting the focus from supporting institutions to supporting individual students, through student debt, and shifting the policy aim from providing a public good to completing the capital market for investment in needed labor market skills, thus had an attraction to federal policy-makers, as evidenced by the Rivlin report of 1969 and the subsequent re-authorization of the Higher Education Act in 1972, which created Pell Grants and Stafford Loans.

More recently, economists have stressed that wage stagnation and labor market inequality are caused by a “skills gap,” that workers lack the skills necessary for today’s jobs. The existence and supposed growth of the college wage premium, the gap between average labor market earnings of high school graduates and college graduates, has been interpreted as evidence of this skills gap. Thus, the policy solution is to expand the federal student loan program so that more students can have access to increased educational attainment. The availability of federal loans has operated in tandem with shrinking state support for higher education institutions, as policy-makers have used the college wage premium to justify shifting those costs onto students in the form of higher tuition. Supposedly, the increase in debt necessary to finance that tuition would ultimately be sustainable because the degrees obtained would pay for themselves by causing students’ earnings in the labor market to increase.

The final piece of the policy and economic background for the student debt crisis and its sharp racial dimension is labor market credentialization and underlying racial wealth inequality, as well as racial segregation in higher education institutions. When the costs of higher education were individualized, those with more family wealth in their backgrounds were better able to shoulder them, and that meant white people. Moreover, labor market

discrimination means that black students have to obtain more degrees, and go into more debt, to obtain the same jobs with the same salaries that white people can obtain with fewer degrees and less debt. Finally, white students have access to institutions with larger sources of external support, so racial wealth gaps are essentially replicated at the institutional level. In the worst case, budget cuts to black- and Hispanic-serving institutions have pushed minority students into the predatory for-profit sector, or induced such institutions to become more predatory on their own student bodies themselves.

The idea that higher education would equalize outcomes by race, regardless of family background, turned out to be a myth: the disposition of policy to de-fund institutions and increase tuition and debt operated against that racially inequalitarian background to widen racial wealth disparities through the “new” instrument of student debt, just as the expansion of housing credit to non-traditional borrowers in the 1990s and 2000s was presumed to be racially egalitarian, but in fact widened racial wealth gaps when the housing bubble burst. In the case of student debt, the bubble isn’t going to burst in the same spectacular fashion that it did in housing because the debt is unsecured, so a debt deflation scenario following a firesale cannot take place. Instead, there is a long, slow diminution in the value of the degrees the student debt financed, in the form of labor market credentialization that worsens the earnings distribution conditional on any one degree and drives people to therefore seek more of them, and take on more student debt to do so, without ever being able to pay it off.

Worsening student debt repayment schedules for sequential cohorts bespeaks this policy failure, and it is why student debt cancellation is on the table as a policy solution at all. In fact, it is already implied by existing Income-Based Repayment programs, which cap payments as a percentage of current income and so push repayment schedules into the indefinite future. This is de facto debt cancellation through the back door, since it reduces the present value of debt repayments, and since IBR programs come with cancellation provisions following a maximum number of years in repayment (10, 20, or 25, depending on the program), there is already a great deal of cancellation baked into the system. The worse the labor

market performs, the larger is the implied cancellation already entailed by current policy. And it should be emphasized that the expansion of student debt and, by that means, higher education attainment, was supposed to be the mechanism by which the labor market would be made to work better, and earnings to grow for better-educated workers.

Summing up, the racial inequality dimension of the student debt crisis requires us to question whether the skills gap and human capital ideology that motivates recent and current higher education and student debt policy is correct as a matter of social science and whether it achieves its normative policy aims. If one such aim is to mitigate wealth disparities by race, the current system is failing dramatically, and cancellation would, as this paper shows, rectify the disparities in racial wealth that the expansion of student debt has caused.

6 Conclusion

This analysis has tested the implications for the racial wealth gap of cancelling student debt, and concluded that the racial wealth gap narrows more the more student debt is cancelled. It's worth emphasizing that point: *across all samples, across all quantiles, the racial wealth gap narrows when student debt is cancelled, and it narrows more the more student debt is cancelled.* Even where it appears that wealth ratios increase, that is in fact due to the effect of lifting more households above the threshold of zero net wealth, and disproportionately many black households benefit from lifting household wealth into positive territory, or proportionally farther away from zero.

Furthermore, and again contrary to conventional wisdom, it's evident from the successively narrower samples used in this analysis that student debt matters more for the net worth of relatively poorer households. The samples are designed to focus more on populations directly affected by student debt as they narrow. That turns out to entail that they are more and more focused on relatively less-well-off, nonwhite households. The idea that student debt cancellation is a giveaway to the rich just isn't sustainable in the data.

References

- Berman, Elizabeth Popp, and Abby Stivers.** 2016. “Student Loans as a Pressure on U.S. Higher Education.” In *The University Under Pressure*. Vol. 46 of *Research in the Sociology of Organizations*, 129–160. Bingley, UK:Emerald Group Publishing Limited.
- Bricker, Jesse, Karen M. Pence, Simona Hannon, and Meta Brown.** 2015. “How Much Student Debt is Out There?”
- Bruenig, Matt.** 2019*a*. “Low Income People Have More Student Debt Than Realized.”
- Bruenig, Matt.** 2019*b*. “The Student Debt Forgiveness Muddle Continues.”
- Dettling, Lisa J., and Joanne W. Hsu.** 2013. “The State of Young Adults’ Balance Sheets: Evidence from the Survey of Consumer Finances.” *Federal Reserve Bank of St. Louis Quarterly Review*, 96(4): 305–330.
- Huelsman, Mark, Tamara Draut, Tatjana Meschede, Lars Dietrich, Thomas Shapiro, and Laura Sullivan.** 2015. “Less Debt, More Equity: Lowering Student Debt While Closing the Black-White Wealth Gap.” Demos-IASP.
- Kakar, Venoo, Gerald Eric Daniels, and Olga Petrovska.** 2019. “Does Student Loan Debt Contribute to Racial Wealth Gaps? A Decomposition Analysis.” *Journal of Consumer Affairs*.
- Looney, Adam.** 2019*a*. “How progressive is Senator Elizabeth Warrens loan forgiveness proposal?”
- Looney, Adam.** 2019*b*. “Who owes the most student debt?”
- Looney, Adam, and Constantine Yannelis.** 2015. “A Crisis in Student Loans? How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults.” *Brookings Papers on Economic Activity*, Fall.
- Looney, Adam, and Constantine Yannelis.** 2018. “Borrowers with Large Balances: Rising Student Debt and Falling Repayment Rates.” Brookings Institution.
- Mitchell, Joshua.** 2019. “The Long Road to the Student Debt Crisis.” *Wall Steet Journal*.
- Morgan, Julie, and Marshall Steinbaum.** 2018. “The Student Debt Crisis, Labor Market Credentialization, and Racial Inequality.” Roosevelt Institute.
- Steinbaum, Marshall.** 2017. “A Brown v. Board for Higher Ed.” In *The President’s House is Empty*. *Boston Review*.
- Steinbaum, Marshall, and Kavya Vaghul.** 2016. “How the Student Debt Crisis Affects African Americans and Latinos.”
- Thompson, Jeffrey P., and Gustavo Suarez.** 2015. “Updating the Racial Wealth Gap.” Federal Reserve Board FEDS Working Paper 2015-76.

A Absolute Wealth Gaps

Huelsman et al. (2015) calculate racial wealth gaps in absolute terms: the difference between the amount of wealth held by the median white household and the median black household. When those authors report percentage changes in the size of the racial wealth gap under alternative debt cancellation scenarios, they are referring to measurements of the type “dollar change in the absolute wealth gap over the ex-ante absolute wealth gap,” always measured at the median.

In this appendix, I report changes in the absolute dollar wealth gap at every centile of the wealth distribution, for each of the three SCF 2016 samples previously described. The charts here are analogous to figures 5, 8, and 11: the net worth percentile is plotted on the horizontal axis, and for each centile, the vertical axis plots the dollar difference in wealth between white households occupying that centile in the wealth distribution of white households and black households occupying that centile in the wealth distribution of black households, for each of three scenarios: status quo, the Warren debt cancellation plan, and the Sanders debt cancellation plan.

As with the previous charts, absolute wealth differences are transformed using the inverse hyperbolic sine. It is that function of the data that is plotted, even though the level of the wealth gap, in dollars, is what is labeled on the vertical axis. Hence, that axis is not to scale.

What these charts show is, once again, that student debt cancellation reduces racial wealth gaps, measured in absolute terms, just as it does when we measure wealth gaps as ratios. These charts show that absolute dollar wealth gaps are increasing the higher up the wealth distribution you go, such that student debt becomes a relatively insignificant component of them for wealthier households. This is also evident in the wealth ratio charts, but wealth ratios among wealthier people change more in response to eliminating student debt than do absolute wealth gaps.

The fact that absolute racial wealth gaps are so large among wealthier households tells us something important about student debt’s role in overall racial wealth inequality, namely, that it is not all that important to the phenomenon of racial wealth inequality. What is important is how much more wealth wealthy people have than everyone else, and the fact that there are almost no black people among society’s wealthiest households. Student debt cancellation does quite a bit to reduce wealth inequality in the middle and working class, but it does little to reduce it among the wealthy, and for the overall phenomenon, the wealthy are the ones who matter. If racial inequality is to be reduced altogether, then, that entails that society’s richest people must become a whole lot less rich in comparison to everyone else than they currently are.




Figure 12: Absolute wealth gaps for Sample 1, by quantile, under alternative cancellation scenarios.




Figure 13: Absolute wealth gaps for Sample 2, by quantile, under alternative cancellation scenarios.




Figure 14: Absolute wealth gaps for Sample 3, by quantile, under alternative cancellation scenarios.