Number 2022-16
December 19, 2022

ECONOMIC COMMENTARY

Student Debt Incidence: Recent Data and Conceptual Issues

Jordan Manes, Emily Moschini, and Thomas Phelan

In recent years, the rising level of student debt has led to calls for increased government assistance in the form of partial or full student debt cancellation. In this *Commentary*, we use the 2019 Survey of Consumer Finances (SCF) to study the incidence of student debt and cancellation benefits along quantiles of household income, net worth, and our estimate of lifetime wealth. We show that student debt is highly concentrated among households with low net worth, but much more evenly distributed across income and lifetime wealth. We then highlight several challenges in using such statistics to draw conclusions about whether cancellation will ultimately increase or decrease inequality in lifetime wealth, and we outline open questions for future research.

Education debt in the United States has risen sharply in the past decade and now stands at roughly $1.6 trillion, constituting the second-largest category of consumer debt behind mortgages.[1] This growth has led to calls for partial cancellation (which the current administration is aiming to enact) together with a subsequent debate about the distributional impact of such assistance.[2] In this *Commentary*, we estimate the distributions of both student debt and the benefits of cancellation schemes using the 2019 wave of the Survey of Consumer Finances (SCF), the primary source of information concerning household balance sheets in the United States. We explore how these estimates depend on assumptions regarding the value of future income, and we then discuss the extent to which debt cancellation is likely to increase or decrease inequality in either net worth or our estimate of lifetime wealth.

Much of the debate over student debt cancellation centers on its likely distributional implications, that is, whether it will primarily benefit the rich or the poor. However, since households differ along multiple dimensions, such as age, education, income, and net worth, the appropriate way to define "rich" and "poor" is not obvious, as there is no single quantity that uniquely measures a household's financial means. To illustrate this point more concretely, consider two hypothetical individuals: the first, a recent graduate from medical school with low net worth but expectations of high future earnings; the second, a middle-aged individual who has some college but no degree and some accumulated wealth but relatively poor prospects for future earnings. Which of these individuals is likely to have greater difficulty in paying down student loan debt? The answer to this question is not obvious and depends crucially on the growth of income over the lifecycle, a quantity that is likely unknown to both the household and to the economists studying their balance sheets.

For this reason, we first use the SCF to calculate the distribution of student debt under a variety of different assumptions attached to income, net worth, and expected future income.

Jordan Manes is a research analyst at the Federal Reserve Bank of Cleveland. Emily Moschini is an assistant professor of economics at William & Mary. Thomas Phelan is an economist at the Federal Reserve Bank of Cleveland. The authors thank Peter Hinrichs, Deborah Lucas, and Sergio Salgado for helpful discussions. Replication materials for this *Economic Commentary* may be found at https://github.com/tphelanECON/student_debt_SCF.

Economic Commentary is published by the Research Department of the Federal Reserve Bank of Cleveland. *Economic Commentary* is also available on the Cleveland Fed's website at www.clevelandfed.org/research. To receive an email when a new *Economic Commentary* is posted, subscribe at https://www.clevelandfed.org/subscriptions.

ISSN 2163-3738 DOI: 10.26509/frbc-ec-202216

FEDERAL RESERVE BANK of CLEVELAND

At one extreme, when we ignore income and focus solely on net worth to define our quintiles, we find that student debt is concentrated primarily in households within the lowest quintile. In contrast, if we define quintiles by income, the distribution of student debt is hump-shaped, first rising among the lower quintiles before falling at the top quintile. To address the question motivated by the above hypothetical example, we then provide estimates of lifetime wealth, that is, the sum of net worth and the present value of their expected future income. We find that even if households attach relatively low value to expected future income, the distribution of student debt by lifetime wealth looks qualitatively different than by net worth, and, in particular, student debt is no longer concentrated within the lowest quintile.

After presenting our estimates of both lifetime wealth and the average benefits of student debt cancellation by quintile, we then discuss two difficulties with using these estimates to draw conclusions about whether cancellation will ultimately increase or decrease inequality in net worth or lifetime wealth.[3] First, we believe that one ought to specify the source of funding for cancellation proposals when assessing redistributive impact. Because the partial cancellation of student debt reduces future government revenue, it may necessitate an increase in either current or future taxes, and these changes in fiscal policy will likely have distributional implications of their own. Although we do not analyze the impact of such hypothetical tax changes in this *Commentary*, we note that any redistributive effects of cancellation are contingent upon the mode of financing and that this ought to be the subject of future work. Second, focusing solely on average benefits by income or net worth quantiles obscures the fact that households with identical income and net worth in general differ in their outstanding amount of student debt. The average amount of student debt forgiven within any group of households, whether that grouping is by net worth or income or another means, may be a poor guide to the impact of a proposal if these benefits are experienced by a minority percentage of the group. This point is particularly relevant for the case of student debt; we show that within each quintile of income or lifetime wealth, the majority of households hold no debt and so would not benefit directly from cancellation proposals.

Student Debt in the Survey of Consumer Finances

The Federal Reserve Board's SCF is a triennial survey designed to provide an overview of the balance sheets of a representative sample of US households.[4] The survey contains variables that influence income and wealth. Households are asked about salaries, wages, interest income, government assistance received (such as the Earned Income Tax Credit), self-employment income, and debts and assets. Since we are interested in estimating how the burden of student debt is distributed among student debt borrowers, we focus on the variables that best represent the ability of households to pay this debt today and in the future. We first compare the income and net worth distributions of borrowers with those of the whole population and then estimate the distribution of student debt within each population. We then use income and net worth and the cross-sectional distribution of income by age to arrive at an estimate for lifetime wealth, and we then record the benefits of student debt cancellation broken down by this last quantity.

Income, net worth, and age

The simplest measures of a household's financial means are current income and net worth (assets minus liabilities). Figure 1 provides us with a first look at how the population of student loan borrowers differs from the general population, recording a selection of percentiles of per capita income and net worth across the two groups of households, those with student debt and those without.[5]

Households with student debt have a higher median income than the general population, while the mean and median net worth of borrowers is much lower than that of the general population. Further, the distribution of net worth is far more unequal than that

Figure 1: Percentiles of Per Capita Income and Net Worth

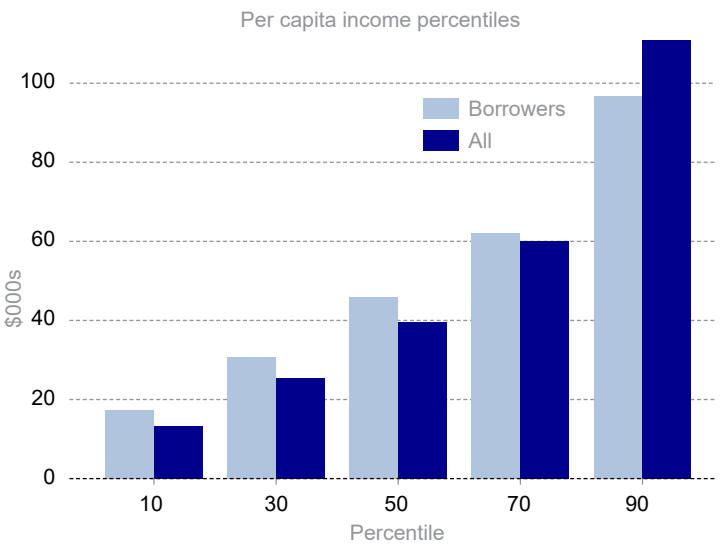

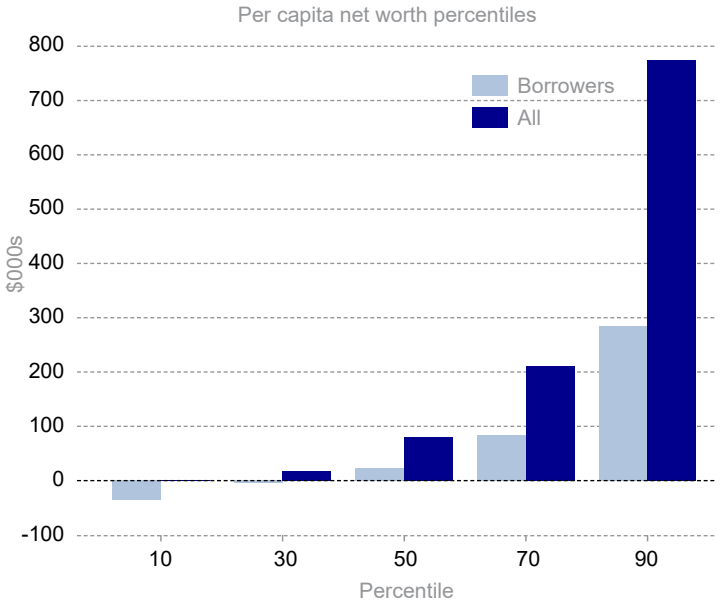

Sources: 2019 Survey of Consumer Finances and authors' calculations

Figure 2: Average Per Capita Student Debt by Income and Net Worth Quintile

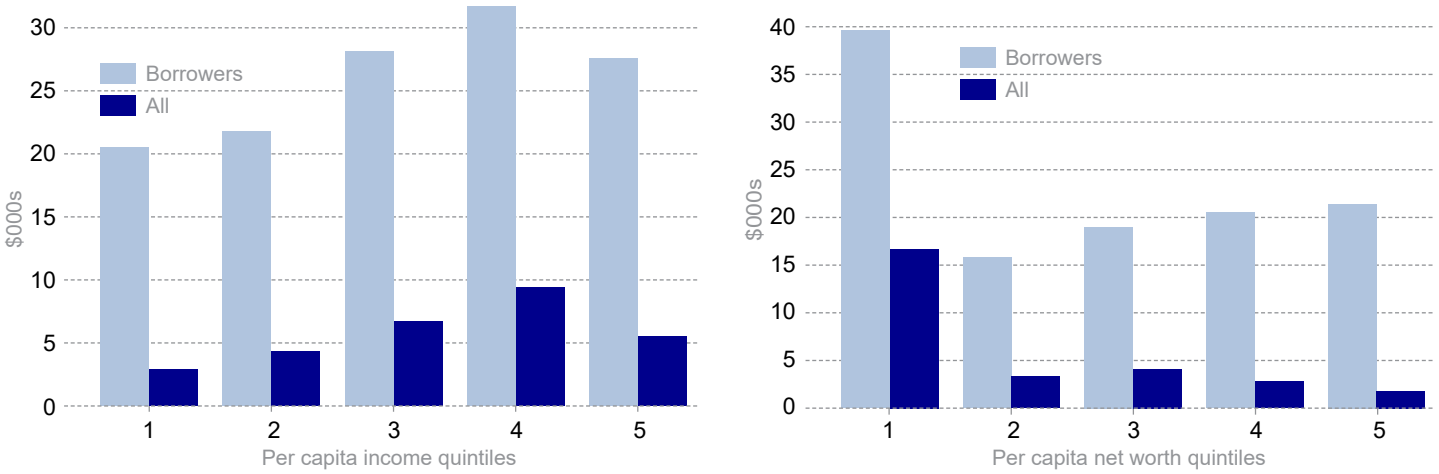

Sources: 2019 Survey of Consumer Finances and authors' calculations

of income for both populations. The difference between the two charts in Figure 1 illustrates the difficulty noted in the introduction in quantifying the burden of student debt. If we were to focus solely on income, we would conclude that borrowers are typically richer than the general population, while the reverse would be true if we were to focus solely on net worth.

To show how student debt is distributed among the population of borrowers, Figure 2 depicts average per capita student debt per quintile of the income and net worth distributions for both borrowers and the population as a whole.[6] For both groups, average per capita student debt exhibits a hump-shaped pattern as a function of income, rising through most of the distribution before declining for the highest quintile. However, among the population of student debt borrowers, average per capita debt does not vary greatly across quintiles, lying between $20,000 and $33,000 in every quintile. Further, the average per capita student debt among borrowers in the highest income quintile is less than 50 percent higher than that in the lowest quintile even though the associated average income is roughly five times higher. In contrast, average per capita student debt varies substantially by quintiles of net worth, and for both the whole population and the population of borrowers, a plurality of debt is held by households in the lowest quintile.

Figures 1 and 2 show that focusing exclusively on either net worth or income materially changes our assessment of the burden of student debt. So which of the two offers the most salient picture? Obviously both are relevant when quantifying the financial means of households, but what we really desire is some measure for lifetime wealth, that is, the sum of their current net worth and the value of their potential future income. However, in contrast with both income and net worth, the value that households attach to their future income cannot simply be directly observed in the data, so we must make further assumptions in order to produce our estimate.

To motivate our choice of estimate, we first explore how income and net worth differ by age. Students typically attend college when young, prior to having accumulated much savings, and so Figures 1 and 2 may simply reflect the dependence of net worth on the age of the household head. Indeed, the median age of a household head in the 2019 SCF with student debt is 37, much

Figure 3: Lifecycle of Net Worth and Income

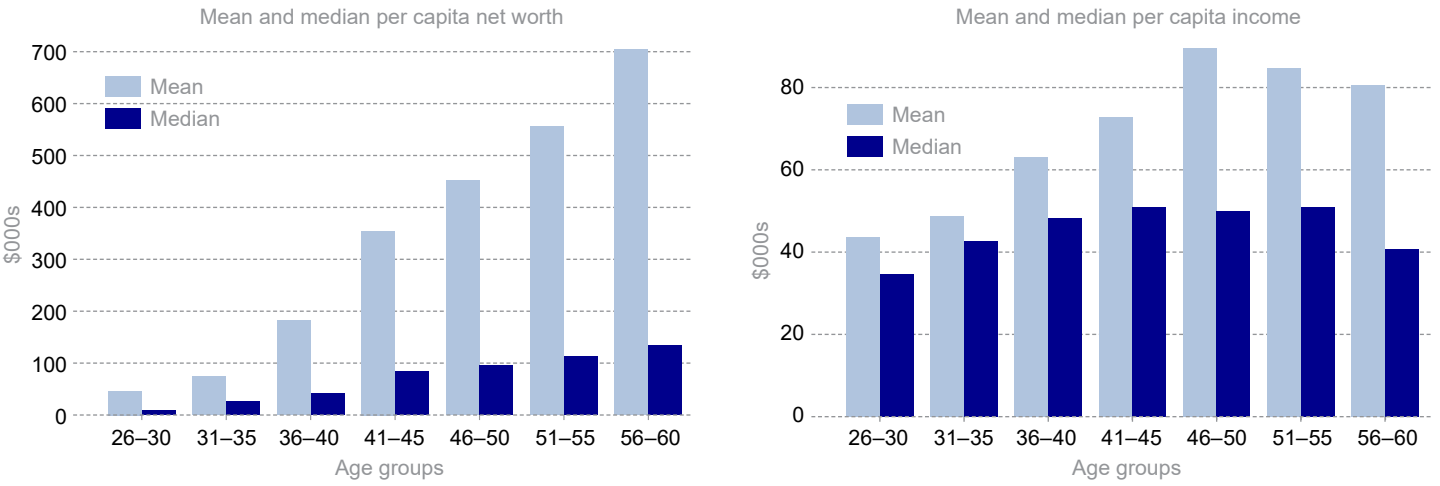

Sources: 2019 Survey of Consumer Finances and authors' calculations

Figure 4: Student Debt by Lifetime Wealth

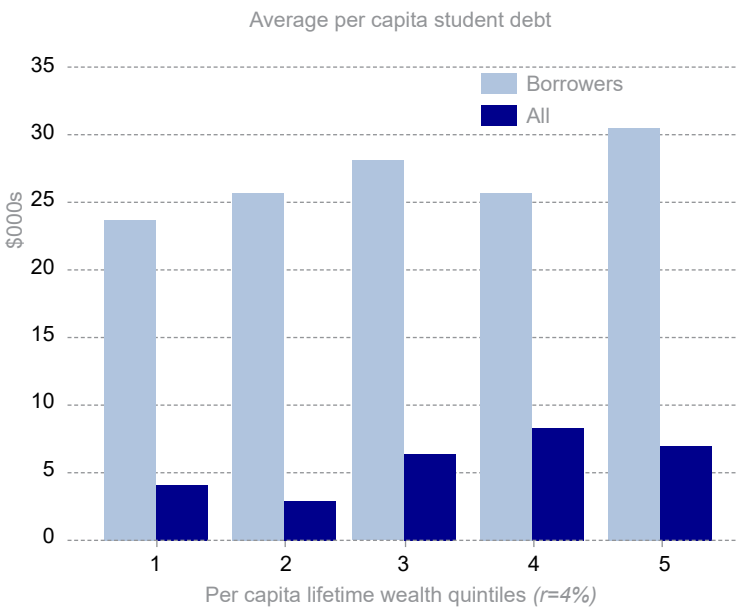

Sources: 2019 Survey of Consumer Finances and authors' calculations

lower than the median age of household heads in the general population, which is 52. Figure 3 explores this point by depicting the mean and median of net worth and income by age group. Both mean and median net worth increase by more than 10 times over the prime-age working life, while the corresponding mean and median for income increase by less than three times and exhibit a more hump-shaped pattern.

Figure 3 shows that the relationship between income and net worth varies systematically by age. These findings suggest that to estimate a household's ability to pay its debt, we ought to account for the fact that younger households can expect a different pattern of future income growth than older households. The precise manner in which we do this is described in the following section.

Estimation of lifetime wealth

There are many ways in which one could use survey data to approximate the lifetime wealth of households. For simplicity, we will suppose that the income growth experienced by households over the next few decades is consistent with the cross-sectional evidence in the 2019 SCF in the following sense.

We group households into eight different brackets based on the age of the household head. The bracket length is five years and begins at age 26 (thus the first bracket is 26 through 30) and ends at age 60 (56 through 60). We begin at age 26 to ensure that the majority of borrowers have left school and entered the workforce, and we end at age 60 in order to encompass the majority of borrowers working lives. We then estimate future income by assuming that the growth of household income is consistent with the cross-sectional distribution of income across ages. Specifically, we compute median income in each bracket and assume that the difference in income in consecutive brackets represents the growth that members of that bracket will experience in the next five years.[7] For a given interest rate, we then define lifetime wealth as the sum of the discounted value of future income and the net worth of the household. This approach has its limitations since the distribution of income across age groups may change in the future and because borrowing costs vary across both time and households. However, this method allows us to combine net worth and current income with the lifecycle aspect of income growth to produce an estimate of ability to pay.

For our main estimates, we follow the Congressional Budget Office (CBO) and use an interest rate of r = 4 percent.[8] Figure 4 plots average per capita student debt broken down by lifetime wealth quintiles. In contrast with the corresponding breakdown by net worth in Figure 2, the average per capita student debt by lifetime wealth is not concentrated in the lowest quintile for either the whole population or the population of student borrowers.

Repeating this exercise for the higher discount rates of 7 percent and 10 percent leaves the qualitative features of the plots unchanged. Even for such high discount rates (or, equivalently, low values attached to future income), it is not the case that average per capita debt levels are highest in the lowest quintile.

Distribution of cancellation benefits

We want to estimate how cancellation benefits are spread over the distributions of income, net worth, and lifetime wealth. To this end, we will consider proposals that cancel up to a fixed amount of student debt per borrower and document per capita benefits across households. We first focus on income and net worth before turning to our estimates of lifetime wealth. Figure 5 plots average per capita cancellation benefits by quintiles of income and net worth when up to $10,000 and $50,000 are cancelled per borrower.

Figure 5: Per Capita Cancellation Benefits

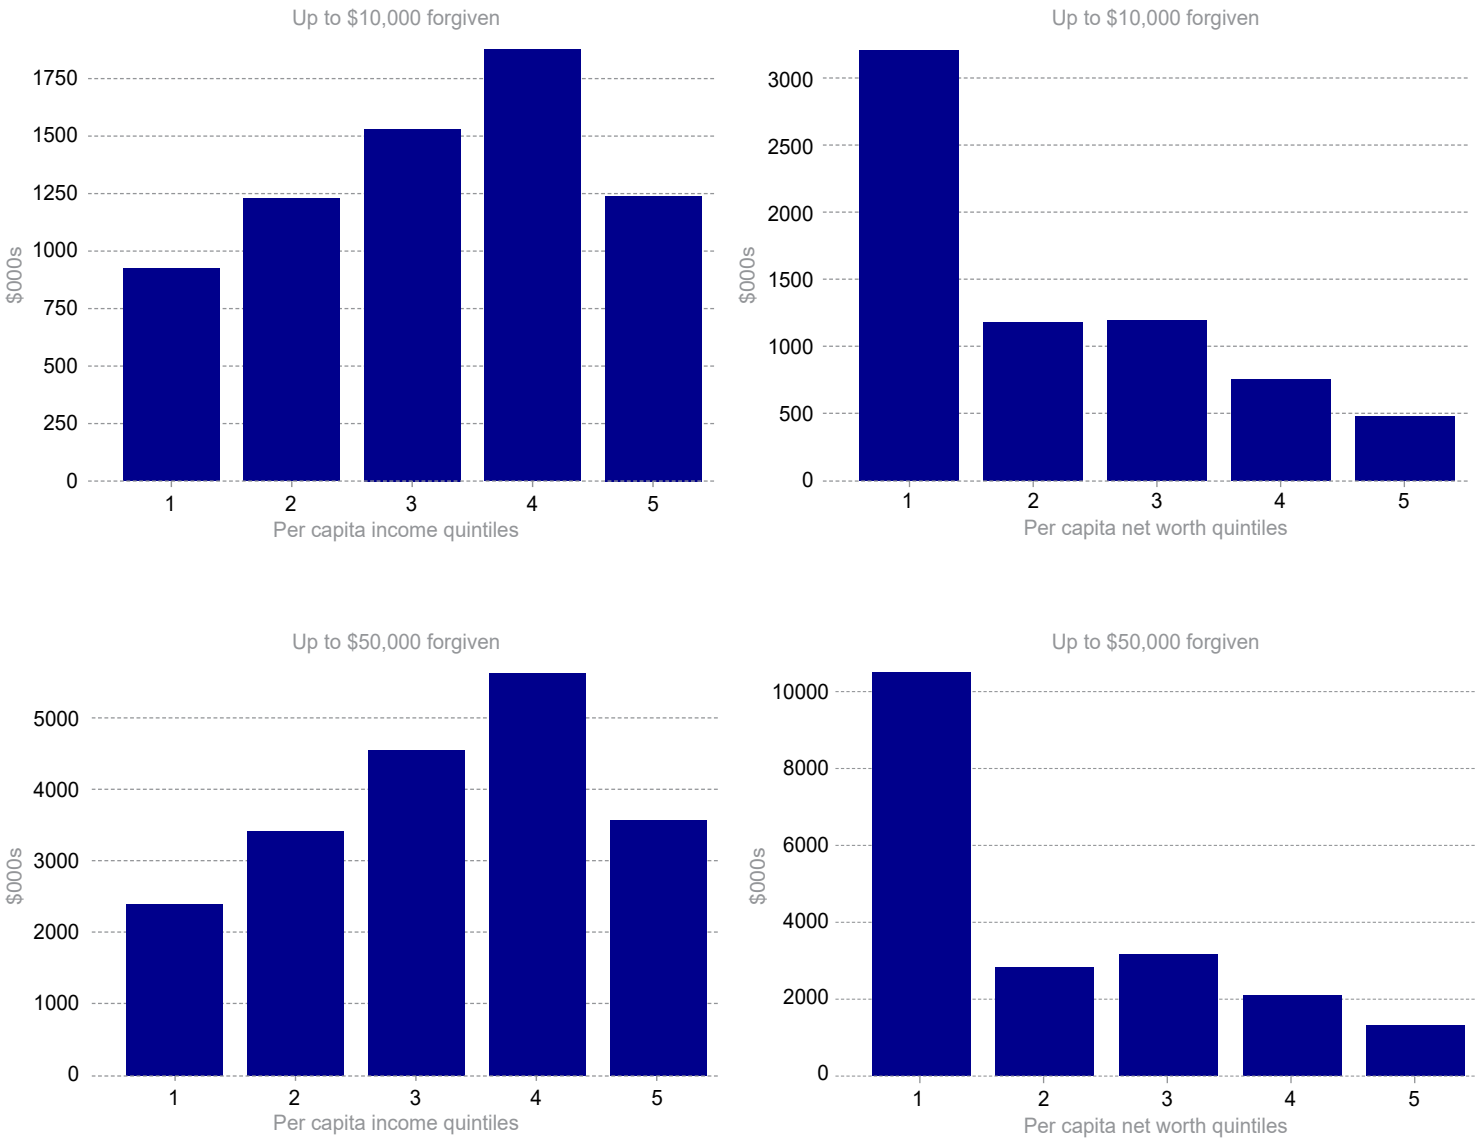

Sources: 2019 Survey of Consumer Finances and authors' calculations

Figure 5 exhibits similar qualitative features as Figure 2, which documents average per capita student debt by income and net worth. For both $10,000 and $50,000, the benefits of cancellation are hump-shaped in the distribution of income but highly concentrated among the lowest quintile in net worth. Figure 6 depicts the distribution of cancellation benefits by lifetime wealth under the assumptions that up to $10,000 and $50,000, respectively, are cancelled per borrower.

In contrast with the categorizations by net worth depicted in Figure 5, for both cancellation levels, the distribution of benefits is not concentrated in the lowest quintile. This pattern of average cancellation benefits across lifetime wealth quintiles is similar to that of average per capita debt depicted in Figure 4.

Progressive or Regressive?

There has been much recent debate as to whether or not student debt cancellation is a "progressive" or "regressive" policy. This debate reflects the natural desire to understand whether cancellation primarily benefits those with the fewest resources, and, relatedly, whether it is likely to increase or decrease inequality in net worth or lifetime wealth. However, the definition of progressive is not obvious in this context, and there are two difficulties in proceeding directly from the above estimates to assessments of the redistributive impact of cancellation.

The first difficulty is that the definition of a progressive transfer scheme is ambiguous when the source of funding is not specified. Economic progressivity is usually discussed in the context of a single tax that is deemed "progressive" if the ratio of taxes paid

Figure 6: Per Capita Benefits by Lifetime Wealth

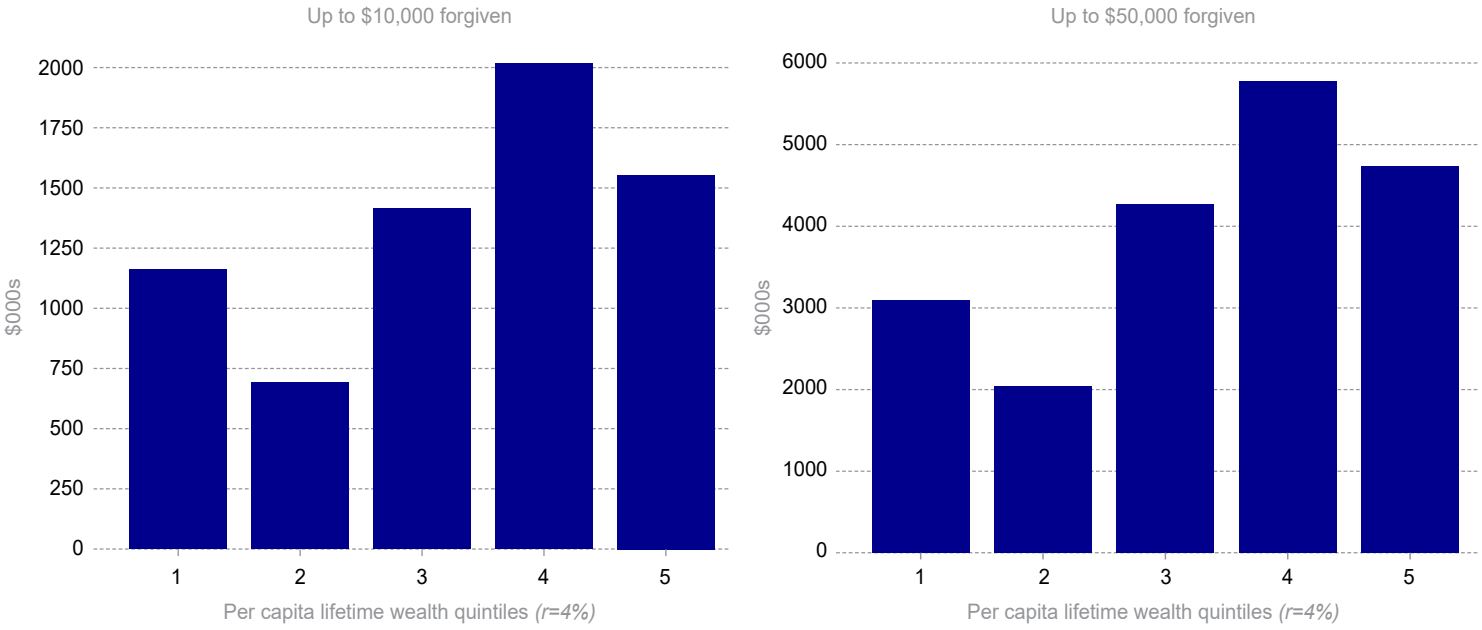

Sources: 2019 Survey of Consumer Finances and authors' calculations

to pretax income rises with income. Implicit in this definition is that the tax incorporates all the transfers between the government and the household (or individual). However, the tax system in place in the United States, one in which individuals are subject to multiple forms of taxation (sales tax, payroll tax, and so on) and are entitled to various separate transfers (such as social security benefits), complicates such designations, and the appropriate equivalent of the definition of progressivity is thus unclear.

This complexity is important when considering which households will benefit most from student loan cancellation. In a discussion of the meaning of progressivity, Slemrod (1993) highlights the importance of understanding the interdependence of different tax and transfer schemes when assessing their distributional impacts:

> The progressivity of the tax structure cannot be judged by looking at only one component of taxes. . . . In recent years the fastest-growing component of federal taxes has been the payroll tax, which is regressive (the opposite of progressive) in its impact, because it taxes at a flat rate only on wages below $63,400 (in 1991). The Social Security system, however, is progressive because it pays higher benefits—relative to taxes paid in—to lower-income workers.

If a policy institutes a tax to fund a spending program, one cannot assess its progressivity or regressivity without reference to the incidence of both the tax and transfers together. As the above passage illustrates, the Social Security system, reflecting both the payroll tax and the benefits paid, is progressive because the ratio of net benefits to taxes paid is higher for lower-income workers. This point illustrates that one difficulty in characterizing student debt cancellation schemes as progressive or regressive is that such an analysis is incomplete until the implications for the government budget constraint are specified. If a fraction of the population has its debts cancelled, then to satisfy the government budget constraint, debt must increase, spending must fall, or tax revenue must rise, and each of these possibilities will have distinct distributional effects. For instance, even if aggregate benefits per quintile increase with income, net worth, or lifetime wealth, so, too, may the associated tax receipts, a situation thus leading to an ambiguous effect on the net benefit by quintile.

The second difficulty concerns the importance of considering differences in student debt incidence among households with the same level of net worth and income. If student debt balances were solely a function of income or net worth, then the standard definition of progressivity could perhaps be applied by considering benefits minus expected future taxes. However, it is not true that households with identical lifetime wealth necessarily have the same debt, and so individuals with the same level of net worth and income will benefit differently from student debt cancellation. Student debt cancellation therefore violates the principle of horizontal equity, the idea that individuals with equal ability to pay ought to be subject to the same burden of taxation and receive the same net benefits from the government.[9] Although departures from horizontal equity are unavoidable in any transfer scheme, this problem

Figure 7: Distribution of Debt within Quintiles

Sources: 2019 Survey of Consumer Finances and authors' calculations

is particularly acute in the case of student debt because of the significant differences in student debt held by individuals even within the same quintile. Figure 7 documents the fraction of households within each quintile that lie within a particular range of per capita student debt and shows that for both income and lifetime wealth, there exists a great deal of difference within quintiles.

In particular, within each quintile of either income or lifetime wealth, the majority of households hold no student debt. Focusing on the average benefits of cancellation in a given quantile is, therefore, potentially misleading because the majority of households in the quantile will receive no direct benefit. The fact that student debt does not depend solely on income or net worth implies that the standard definition of progressivity is not applicable even if the source of the funding were fully specified and lifetime income could somehow be measured without error. This also makes it difficult to proceed directly from the distributions depicted previously to claims about the distributional impact of debt forgiveness proposals.

Conclusion and Areas for Future Work

In this *Commentary*, we have documented several facts regarding both the distribution of student debt and the benefits of cancellation proposals. We have shown that whether or not one considers households with student debt to be richer or poorer than the average household depends crucially on whether we assess financial means using income or wealth. Further, we have used the SCF to provide an estimate of lifetime wealth, or net worth plus the value of future income, and have shown that the distribution of debt by this latter quantity is closer to income than to net worth.

However, areas of future work remain. We have argued that knowledge of average benefits by quintiles of income, wealth, or our estimate of lifetime wealth is not sufficient to determine whether such a policy will ultimately increase or decrease inequality in net worth or lifetime wealth. In order to more fully assess whether a student loan cancellation policy will increase or decrease inequality, the estimates of the distribution of benefits ought to be complemented by two things: an indication of where funding for cancellation and the attendant tax burden will come from and a move beyond averages by quintile to a measure of overall welfare that considers differences in student debt burden and projected lifetime wealth within quintiles. The method by which one compares the merits of various proposals ought to be stated explicitly, as simply computing average benefits per quintile appears inadequate. When analyzing the effects of various economic proposals, economists typically assume that the goal is to maximize a weighted average of the welfare of individual citizens, taking into consideration that an additional dollar is worth more to a poor household than to a rich one. Such an objective will lead one to prefer to provide some form of social insurance via redistribution. With this in mind, the benefits of redistributing from rich to poor households must be weighed against the costs of treating individuals with the same lifetime wealth differently.

Endnotes

1. See the Federal Reserve Bank of New York's *Quarterly Report on Household Debt and Credit* for 2021:Q3 at https://www.newyorkfed.org/medialibrary/interactives/householdcredit/data/pdf/HHDC_2021Q3.pdf.
2. See, for instance, Catherine and Yannelis (2020) and articles published by the Roosevelt Institute at https://rooseveltinstitute.org/2021/06/08/student-debt-cancellation-is-progressive/ and the Brookings Institution at https://www.brookings.edu/blog/up-front/2019/04/24/how-progressive-is-senator-elizabeth-warrensloan-forgiveness-proposal/.
3. As we note below, these difficulties are also relevant for the question of whether cancellation ought to be viewed as a "progressive" or "regressive" policy, terms that we believe are ambiguous in meaning in this context.
4. For a more comprehensive analysis of the 2019 wave of the SCF, see Bhutta et al. (2020).
5. Throughout this *Commentary*, we define per capita quantities by dividing household quantities by two if the household head is married or lives with a partner.
6. In Figure 2, the quintiles are defined using the entire population.
7. For example, if the median income of households in the 31–35 age bracket is 10 percent higher than those in the 26–30 age bracket, then we assume that the income of a household headed by a 26-year-old grows by 10 percent over the next five years.
8. See Appendix A of Income-Driven Repayment Plans for Student Loans: Budgetary Costs and Policy Options, CBO, available at https://www.cbo.gov/publication/55968.
9. See, for example, Musgrave (1990) and the references therein for the history of this idea and further discussion.

References

Bhutta, Neil, Jesse Bricker, Andrew C. Chang, Lisa J. Dettling, Sarena Goodman, Kevin B. Moore, Sarah Reber, Alice Henriques Volz, and Richard A. Windle. 2020. "Changes in US Family Finances from 2016 to 2019: Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 106 (5): 1–42. https://doi.org/10.17016/bulletin.2020.106.

Catherine, Sylvain, and Constantine Yannelis. 2020. "The Distributional Effects of Student Loan Forgiveness." Working paper 28175. National Bureau of Economic Research. https://doi.org/10.3386/w28175.

Eaton, Charlie, Adam Goldstein, Laura Hamilton, and Frederick Wherry. 2021. "Student Debt Cancellation IS Progressive: Correcting Empirical and Conceptual Errors." Brief. Roosevelt Institute. https://rooseveltinstitute.org/publications/student-debt-cancellation-is-progressive/.

Karamcheva, Nadia, Jeffrey Perry, and Constantine Yannelis. 2020. "Income-Driven Repayment Plans for Student Loans: Budgetary Costs and Policy Options." CBO report 55968. Congressional Budget Office. https://www.cbo.gov/publication/55968.

Looney, Adam. 2019. "How Progressive Is Senator Elizabeth Warren's Loan Forgiveness Proposal?" *Brookings* (blog). April 24, 2019. https://www.brookings.edu/blog/up-front/2019/04/24/how-progressive-is-senator-elizabeth-warrens-loan-forgiveness-proposal/.

Musgrave, Richard A. 1990. "Horizontal Equity, Once More." *National Tax Journal* 43 (2): 113–22. https://doi.org/10.1086/NTJ41788830.

"Quarterly Report on Household Debt and Credit." 2021. Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/interactives/householdcredit/data/pdf/HHDC_2021Q3.pdf.

Slemrod, Joel B. n.d. "Progressive Taxes." Library of Economics and Liberty: The Concise Encyclopedia of Economics. Accessed April 28, 2022. https://www.econlib.org/library/Enc1/ProgressiveTaxes.html.