Washington Center for Equitable Growth
Evidence for a stronger economy

equitablegrowth.org

Comments from the Washington Center for Equitable Growth on BEA's "The Feasibility of a Quarterly Distribution of Personal Income"

I'm writing to provide comments on BEA's report, "The Feasibility of a Quarterly Distribution of Personal Income." This feasibility study is an important first step toward more timely data on the distribution of income growth in the U.S. economy. While I appreciate BEA's efforts to place some error bands on one possible technique for improving the timeliness of the product, I believe there are methodologies that could make more timely reporting sufficiently accurate for publication as a federal statistical series. The conclusions of the feasibility report should not preclude studying these options for suitability in a BEA product.

The Distributing Personal Income data series is an incredibly useful new product, and BEA has done laudable work both in launching this entirely new data series with limited resources and in integrating improvements into the series with each new release. There is no other data series produced by a federal statistical agency that provides a comprehensive look at how the U.S. economy is working for Americans up and down the income ladder.

Yet the current compromises made on timeliness are holding this product back from wide adoption. With the U.S. Congress poised to provide financial resources specifically for the development of this product, BEA should research further improvements to the timeliness of the data.

The feasibility report should be a starting place, not an end point, for thinking about the challenges that may confront BEA in attempting to increase the frequency and decrease the lag of this important data series. This comment to BEA examines, in three sections:

- Why more timely reporting is valuable and should be BEA's goal
- A promising method that has now been applied by two teams to achieve more timely reporting of distributional growth
- Some suggestions for how BEA can proceed

Briefly, here is a summary of my comments: BEA should continue to investigate methods for reducing lag and increasing frequency in the Distributing Personal Income dataset. I make four arguments in favor of investigating more timely reporting:

1. The Distributing Personal Income series is a disaggregation of an existing national income data series, Personal Income. There is no compelling reason to publish these two datasets on vastly different timelines.
2. There is considerable demand for a more timely product from policymakers, the media, and businesses. Members of each of these groups have specifically indicated that it is inadequate to receive data on income inequality in the U.S. economy that is 2 years out of date.
3. The feasibility report did not look at whether integrating new data sources into the prototype series could reduce error in nowcasting. Contrary to claims in the feasibility report, doing so does not require panel data nor a data source with a joint distribution of transfers and wages.
4. Teams from Realtime Inequality and Statistics Canada are using microsimulation to reallocate wages and transfers on a timely and subannual basis. Modeling and microsimulation methods are increasingly common in the production of federal economic statistics.

Based on these arguments, I believe BEA should pursue both lower-lag and subannual reporting. I suggest the following:

1. Use the QCEW and other data sources, such as the monthly CPS and CES, to reduce lag for annual estimates as much as possible. Full annual QCEW data is published in early June of the subsequent year. Annual distributional data could therefore be published on a two-quarter lag, with 2022 estimates released in June 2023 (or earlier, if BLS is able to furnish estimates based on QCEW to BEA sooner). Only one revision would be necessary, when SOI tables are released or become available to BEA.
2. Investigate the possibility of quarterly reporting using data sources such as the QCEW, CES, and monthly CPS to update distributions of income subannually. Only BEA can evaluate whether these methods will produce valid and informative estimates when applied to the existing dataset.

BEA should pursue more timely reporting

The current Distributing Personal Income prototype could be improved by increasing frequency and by reducing lag. The issues sometimes overlap, but I will try my best to clearly distinguish them in this comment. Although they are distinct issues, there is good reason to care about both. A general point is that BEA is undertaking the disaggregation of an existing product—Personal Income—that is released both on a low lag and at quarterly frequency. Many of the same objections that apply to lower-lag and higher-frequency distributional national accounts can equally be levied against national accounts themselves.

Gross Domestic Product calculations, for example, rely on significant imputation in the first estimate. It's difficult to put an exact number on the amount of imputation, interpolation, and modeling that goes into the calculation of GDP, but it is almost certainly greater than half of the data that are used as inputs into a first estimate (see Chart 1 on page 5 here). Some inputs into GDP, such as the rental value of owner-occupied housing, must be imputed in any vintage. Other inputs, among them corporate

profits, start out as extrapolations based on other data series before eventually being replaced with IRS data when they are available.

GDP estimates are also revised many times. Aside from the second and third estimates, there are annual, 5-year, and further revisions. According to the Federal Reserve Bank of Philadelphia's Real-Time Data Set for Macroeconomists, nominal GDP for the first quarter of 2000 has been revised nine times as of this writing, with the most recent revision in 2021. These revisions sometimes alter the substantive interpretation of a release. As the quarterly feasibility report notes, these errors are more common in recessions. In the 60 quarters that have received a third GDP estimate since the beginning of 2007, the year of the financial crisis, the sign of the advance estimate and the third estimate have been different on three occasions, a 5 percent chance (my calculations from the Real-Time Data Set for Macroeconomists).

Despite these challenges, NIPA estimates are valuable intelligence for many economic actors, including households, businesses, and economic policymakers, elected or not. No one would argue today that we should reduce the frequency of these releases or accept longer lags to have more stable, less-revised estimates. The challenges of disaggregating NIPA estimates, as the Distributing Personal Income prototype does, are similar to those faced by BEA in other areas, and parity, or near-parity, with existing data series is an important and achievable goal.

Lower-lag reporting

Reducing the amount of lag between the close of a year and the first estimate of income shares in that year is the more straightforward methods problem and should be prioritized over quarterly frequency. Under BEA's current release schedule, estimates of the income distribution in 2020 will arrive in December 2022, approximately a 2-year lag from the close of the period that is being reported. Those estimates will give us some idea of how the income distribution was affected by one of the most consequential economic events in modern history—the swift and sudden onset of the COVID-19 pandemic—but the estimates will document just the first year of a pandemic that, by then, will be nearly 3 years old.

This is a long time to go with very little knowledge of how the income distribution has been altered by a massive disruption of the U.S. economy and the passage of some of the most significant economic stimulus policies ever implemented in U.S. history. Reducing this lag is critical to the usefulness of the product.

A lower-lag product will attract a much larger user base. Journalists, business analysts, and interested U.S. residents are all more likely to find value in a timely dataset. A larger user base would make it easier for BEA to raise resources to improve the product. Federal statistical products that do not strongly appeal to potential audiences risk losing funding and stagnating or being eliminated altogether.

In fact, this was a significant topic at the June 2021 Federal Economic Statistics Advisory Committee meeting. Catherine Rampell of *The Washington Post* repeatedly cited timeliness as an important attribute for federal datasets to meet journalistic demand. She further indicated that at the time, she was especially interested in better understanding the possibility of a "K-shaped recovery," which directly implicates BEA's Distributing Personal Income dataset.

Similarly, at the same meeting, Dick Rippe of Evercore ISI noted that as a business analyst, he was interested in learning more about what kinds of households were suffering income losses due to the pandemic. In his remarks (at 4 hours and 15 minutes in the recording of the meeting), he mentions that he considered BEA's prototype data series, but that it was not useful to him because of the long lag. According to Rippe, "it would have been nice to have more useful breakdowns by income distributions."

Policymakers also express interest in a more frequent product. The House Committee on Economic Disparity and Fairness in Growth sent a letter to U.S. Secretary of Commerce Gina Raimondo, signed by eight members of the committee, urging greater timeliness in the product. The Measuring Real Income Growth Act, sponsored by Rep. Carolyn Maloney (D-NY), likewise indicates a preference for quarterly frequency.

BEA may be able to reduce this lag if SOI or another agency can share aggregate tables for high earners sooner. This is clearly desirable since it would require no modifications to the current methodology, but there are limits to how much lag can be reduced under the current methodology. BEA should additionally research the possibility of integrating lower-lag data products, such as the QCEW, into advance annual estimates. These advance annual estimates would only require one revision, when the annual SOI tables become available.

Quarterly reporting

Quarterly estimation offers a steeper challenge, but the quarterly feasibility report does not provide compelling reasons to abandon this path. Per the section on annual reporting above, subannual reporting is valuable to data users. Economic realities can shift quickly in recessions, and subannual reporting is critical to understanding economic conditions in real time and responding effectively to them. The U.S. Census Bureau acknowledged this when, during the early months of the pandemic, it released its Census Household Pulse survey. Significant policy decisions were made in 2020 that could have benefited from better subannual estimates of household incomes.

There are multiple claims made in the quarterly feasibility report about higher-frequency reporting that do not convincingly rebut the idea that quarterly statistics can be valid, informative, and transparent. Several sections of the quarterly feasibility report make arguments that are not entirely relevant to this debate. Sections 2.3 and 2.4 serve mostly as rebuttals to the idea that BEA should naively impute quarterly income without regard for subannual changes in the income distribution, but this is not a method for which proponents of subannual reporting have advocated. Rather, I and others propose that BEA should use datasets that are gathered subannually to update distributions throughout the year.

BEA further argues that to provide subannual updating, the agency requires either panel data following U.S. households over time or data that capture the joint distribution of wages and transfers. Neither is necessary, as the Realtime Inequality team has demonstrated. I consider these two data requirements in the next section.

Methods for current year and subannual reporting have now been demonstrated

Reporting about inequality in the current year and on a subannual schedule is challenging because many of the relevant data sources that go into the creation of the current prototype product are not available subannually. Likewise, some of the data that are used to adjust the dataset to account for very-high-income households in the United States, such as tables from the IRS's Statistics of Income program, are only available on significant lags.

In the quarterly report, BEA investigates one possible method for reducing lag by taking known distributions of categories of income in a year for which it does have data and simply applying those distributions to known aggregates in more recent years. For instance, if the agency knows that in 2019, the lowest-income 10 percent of households earned 1 percent of compensation of employees, then this method would simply apply that percentage to the aggregate for compensation of employees that is reported on a timely basis in 2021. This method implicitly assumes that the division of income aggregates between household percentiles changes only slowly.

For some income concepts, this assumption is largely accurate. Interest income, for example, is relatively stable over time. In 2007, BEA reports that the highest-income 10 percent of households earned 55.9 percent of all household interest income. More than a decade later, in 2019, this figure had dipped only slightly, to 55.2 percent. In addition to the slow rate of change, errors in this category are somewhat attenuated because it is a relatively small share of total income. By contrast, compensation of employees can shift more significantly, and those shifts are important to capture because it is the single largest income category.

We can visualize these distinctions by decomposing annual changes in household income in the existing BEA product into change that is a result of changes in aggregate income categories and those that are a result of changes in how income is divided between households at different points in the income distribution. In Figure 1 below, I show this decomposition, which is based on calculating what income would be for each group in a counterfactual where distributions don't change at all. The difference between the actual BEA estimates and these counterfactuals is the change in household incomes due to a changing distribution of aggregate income concepts. I have also highlighted what part of these changes is due to changes in transfers and wages.

Figure 1 shows that for the bottom 90 percent of households, most income comes from transfers and wages, and changes in the level of aggregate income dominate those caused by changes in the distribution. In 2001, for the bottom 50 percent, the growth in the general economy would suggest that aggregate income for this group increased by $83.5 billion (all numbers in 2012 dollars), which is shown in the left side of the top panel in Figure 1. Essentially all of this increase was due to wages and transfers.

The distribution of wages and transfers, however, actually shifted slightly away from this group in 2001. The righthand panel shows that when the distributional changes are accounted for, the total change is $44 billion less than what the aggregates alone suggest, of which 74 percent is due to a shift in the distribution of wages and transfers.

FIGURE 1

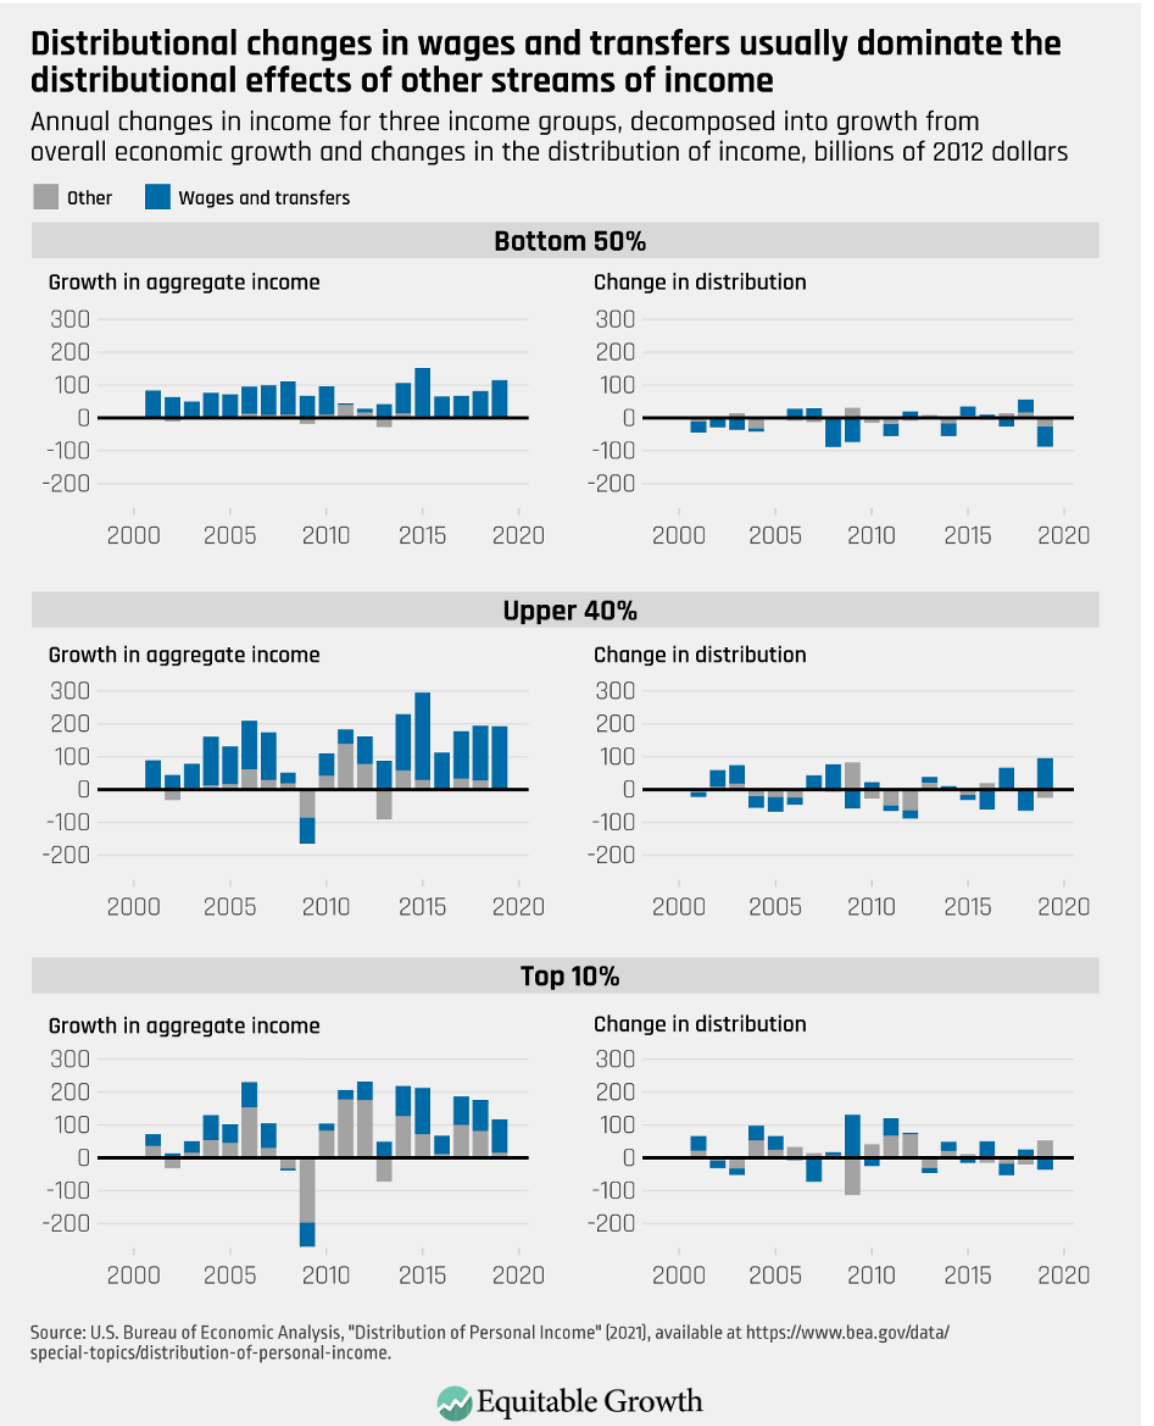

Indeed, much of the change in income is a result of changes in the distribution of wage and transfer income, even though other sources of income clearly do matter in some years. This is important because methods do exist to redistribute these categories of income both with little lag and on a subannual basis.

The challenge, then, is to find low-lag and subannual data sources that can be used to redistribute some of the larger categories of income between income groups. Section 2.3 of the quarterly

feasibility report makes two arguments about the data that would be necessary to increase reporting timeliness. The first is that subannual estimates cannot be made without panel data. The second argument is that subannual (and presumably low-lag) estimates cannot be made without a data source that contains information on both wages and transfers, so their distribution can be jointly observed. Neither argument is convincing.

The prototype data series is a snapshot of the distribution of income at a given time. BEA states that without panel data, "we will not be able to understand clearly how changes in households' quarterly income 'add up' to an annual distribution of income." I interpret this as concern that quarterly estimates based on nonpanel data will not agree with the annual estimates based on CPS and SOI data. Some form of reconciliation between annual and quarterly estimates, however, are inevitable regardless of whether panel data are employed or not.

Any form of panel data from a dataset that is not currently used to create the annual estimates would require reconciliation with the annual estimates because the distribution of wages and other income concepts in that dataset will not match those derived annually from the CPS. In fact, even if the CPS ASEC had quarterly panels, allowing BEA to create subannual estimates using only datasets that are already employed in the annual estimates, mobility of households during the year means quarterly estimates of inequality would not necessarily average out to annual estimates.

A joint distribution of wages and transfers, while potentially useful, is not strictly necessary either. Two teams of researchers have started producing nowcasts of the distribution of income that use microsimulation techniques to update the distribution of wages and transfers. Realtime Inequality, a project from University of California, Berkeley economists Thomas Blanchet, Emmanuel Saez, and Gabriel Zucman, and Statistics Canada, in its Distributions of Household Income Economic Accounts dataset, both use similar methods to make subannual and low-lag estimates of how much households at different levels of income are benefitting from economic growth.

Both teams use microsimulation to model changes to wages, transfers, and employment in their household- or individual-level microdata. Statistics Canada is additionally able to redistribute other categories of income thanks to subannual survey questions that do not have close analogues in the United States. In the absence of a dataset that can provide a joint distribution of wages and transfers, the operative question is whether microsimulation performs well enough in this particular data series. This is an empirical question that BEA must test.

For U.S. estimates, the Realtime Inequality team uses the QCEW and the CPS to recalculate the distribution of wage income on a quarterly basis and the CPS and ACS to estimate employment changes. Byoungchan Lee at the HKUST Center for Economic Policy first showed that the QCEW, due to its immense level of geographic and industry-type disaggregation, can be used to approximate shifts in the composition of earnings inequality on a quarterly basis. This approach uses the QCEW like a survey, with employment counts in each cell as a weight and average wage as an observed income.

Realtime Inequality uses rules-based microsimulation to distribute transfers. In other words, it uses its underlying microdata and the rules of new transfer programs to simulate which households receive transfers in its dataset. The first COVID-19 stimulus payments, for example, were distributed to single

tax-filers with up to $99,000 of income, with the amount reduced for those with incomes of more than $75,000. Not every program is so easy. It is not immediately clear, for example, how to distribute the Paycheck Protection Program, and the Realtime Inequality team uses a 2022 NBER working paper by MIT economist David Autor and his co-authors to decide on incidence, which would not have been available in the first 2 years of the pandemic.

The Realtime Inequality approach is relatively accurate. The project looks at four income categories: the bottom 50 percent of individuals by income, the next 40 percent (the 50th to the 90th percentile), the next 9 percent (90th to 99th percentile), and, finally, the top 1 percent. For each income category, Realtime Inequality predicts an income growth rate for that group for the next year and for the next 2 years—that is, the project predicts the overall growth rate for the next 2 years, not for the year that is t+2—based on its lagged annual source data, the same sort of calculation BEA would have to perform.

Realtime Inequality's 1-year-ahead estimates predict the correct sign close to 90 percent of the time for its top three income groups. For the bottom 50 percent, the project record is not as strong, with about 73 percent accuracy in its disposable income measure. I believe BEA has distinguished itself by putting significant work into producing accurate estimates of income in the bottom deciles, which may prove to be an advantage over Realtime Inequality for this kind of forecasting. In its 2-year-ahead forecasts, Realtime Inequality correctly predicts the sign of change in growth 88 percent of the time for the bottom 50 percent and 95 percent of the time for its next 40 percent and its next 9 percent groups.

These prediction rates compare favorably to the 95 percent accuracy BEA has recorded for the sign of GDP since 2007 that I discussed above. GDP estimation is, of course, a much more mature measurement methodology. If we look at the same window of time that Realtime Inequality does, starting from 1975, quarterly GDP forecasts have had the right sign 96 percent of the time when comparing advance estimates to revisions 1 year past the advance estimate and 94 percent of the time when comparing to revisions 2 years past the advance estimate.

The Realtime Inequality approach differs from the BEA approach in that its annual source data is a synthetic file based on federal tax information, whereas BEA starts with the CPS. Whether the Realtime Inequality approach will perform well when combined with BEA's CPS-based microdata is not currently known but can be tested empirically.

Evaluating the accuracy of microsimulation should be a priority

Microsimulation techniques, such as those described above, are not new or experimental. Microsimulation is the workhorse methodology of the analysis of proposed new taxes and new spending legislation. It is important to CBO's distribution of income report and to the Distributing Personal Income data series itself, which uses NBER's TAXSIM to impute both taxes and transfers. Techniques that broadly fall under the umbrella of modeling and microsimulation are becoming more common in federal data products as demand for more granular data grows and methodologies for

delivering that data improve. The U.S. Census Bureau is increasingly employing small area estimation techniques, for example.

The BEA dataset is significantly different from the Realtime Inequality dataset, which starts with IRS data, targets a different income concept (Realtime Inequality targets National Income), and incorporates different source data and assumptions. Accordingly, it is impossible to know whether similar methods will lead to useful nowcasting in the BEA dataset until it is attempted. Statistics Canada has not yet published an analysis of the validity of its new quarterly distributional data, so the reliability of its prototype is unknown.

The utility of more timely data is considerable, so BEA should consider appropriate trade-offs that might make the production of more timely data feasible. It may be difficult, for example, to provide reliable estimates for all 10 deciles of income. If nowcast estimates are more reliable for quintiles or for coarser categories, such as those used by the Realtime Inequality team, then there is no need for BEA to retain the existing decile structure of the dataset in more timely estimates.

Ultimately, only BEA can assess whether the Realtime Inequality methodology will produce good estimates of current-year and quarterly income growth distributions in its microdata. Assessing the validity of these methods is difficult because there is no one metric that provides a sufficient picture. Getting the sign of growth right is important, but if growth is near zero for a group, it might be less important to get the correct sign. Likewise, it might be acceptable to obtain the wrong sign on occasions where the general trend of inequality is correct, as measured by ratios of high-income group-to-low-income group average income.

If BEA were to add nowcasting to the prototype data series, then it would require resources to make periodic revisions to the product. Investigating this issue in the prototype phase will help BEA decide what resources are necessary and request them in annual budgeting. Creating a timely, useful product for policymakers will create the constituency to pass these budget requests and ensure that BEA has the resources it needs to continue to expand and improve the data series.

Conclusion

In the quarterly feasibility report, BEA states that, "an important test of any method for estimating the quarterly distribution of income will be whether it is informative during these quarters [during recessions and recoveries] of rapid change." This is indeed critical but should be viewed in the context of NIPA calculations generally, which often have larger errors at these same moments but are nonetheless incredibly valuable to businesses, policymakers, media, and households in the United States.

Measurement often requires finding the right balance of speed and accuracy. I believe that the demonstrated demand for these estimates from business analysts, the media, and policymakers is a strong argument for accepting some measurement error to increase timeliness. This trade-off is endemic to federal statistics and, as I showed above, GDP itself is not exempt.

The Distributing Personal Income dataset is an important addition to federal statistics. It is unique among all data published by federal statistical agencies in tracking inequality in a comprehensive income concept. The Census Bureau's annual Income and Poverty Report uses the money income concept, which does not include noncash benefits and does not account for taxes and other deductions from households' income. CBO's distribution of household income report captures some of these concepts, but it is not part of the federal statistical system, is not released on a set schedule, and is not compatible with any NIPA concept.

Although there are many extensions BEA could make to this prototype data series, I believe improving timeliness is most likely to lead to widespread adoption. If current appropriations bills for FY2023, now before the U.S. Congress, are passed, then BEA will receive $2.7 million to improve this project. At least some of this money should be spent on investigating whether the microsimulation techniques described above can be adapted to produce useful estimates of current year and subannual distributions of income.