Economic Brief

November 2021, No. 21-39

COVID Transfers Dampening Employment Growth, but Not Necessarily a Bad Thing

By Felipe F. Schwartzman

Overall employment levels have remained below their pre-pandemic level and are growing only slowly despite rising wages and vacancies. In this *Economic Brief*, we examine whether historically high government support may have empowered workers to pull back from labor markets. While that support presents a clear benefit to recipients, a simple calculation based on recent estimates indicates that transfers of close to $2 trillion to households approved over the course of 2020 and 2021 implies a reduction of 0.58 percentage points in the employment-to-population ratio.

As the U.S. recovers from the pandemic, the demand for goods and services has skyrocketed, as evidenced by widespread shortages. Still, overall employment levels have remained below their pre-pandemic levels and are growing slowly despite rising wages and vacancies. This apparent labor shortage has been framed two ways: (i) the consequence of high unemployment benefits, which disincentivize work, and (ii) the result of workers holding back from the labor market because of infection fears and unpredictable child care.

In either case, the implication is that low labor-market participation reflects a lingering distortion that will hopefully resolve itself as unemployment benefits are exhausted and COVID-related conditions improve.

In this *Economic Brief*, we advance a third option: Historically high government support may have empowered workers to pull back from labor markets, even if the support did not create a disincentive per se.

Means-Tested Support vs. Universal Support

The difference between this option and the disincentive effect of high unemployment benefits is subtle, but important. When government assistance comes in the form of unemployment benefits or other "means-tested" programs, it may disincentivize work because individuals eager to join the labor force will instead not join so they can keep their benefits.[1]

In contrast, support that is not dependent on income or employment status doesn't force this choice on workers, giving them more options. Yet, labor force participation may still decline as workers use the additional resources provided by the programs to scale back their labor hours in favor of more highly valued pursuits. This decrease in labor hours may persist for as long as those workers have higher-than-baseline cash balances.

To more concretely make the point, consider two large recent provisions: (i) the more than $1 trillion in direct payments that households received from various fiscal packages designed to mitigate economic fallout from the pandemic, and (ii) the expanded child care tax credit, made available to most families with children in 2021.

In both instances, workers do *not* lose transfers if they work. However, they may still work less than they would otherwise, as money becomes a less pressing issue. This is a clear net gain to those workers.

In particular, because they receive transfers regardless of their working status, any reduction in their labor earnings is more than made up by the benefits of additional time away from work. In other words, any change in the behavior of recipients follows from having greater choices and opportunities. The net effect is positive for recipients, even if it is not accounted for in GDP.

To what extent can such a channel explain the current shortfall in employment? As discussed in depth in my 2017 paper "Does Redistribution Increase Output? The Centrality of Labor Supply" (co-authored with Kartik Athreya and Andrew Owens), this depends on workers' marginal propensities to supply labor (MPLs).[2] This is the flip side of the more commonly evoked marginal propensity to consume (MPC).

Just as individuals may take advantage of a cash windfall to increase their consumption, they may also take advantage of a windfall to increase their time away from work. The paper emphasizes that this response can substantially dampen the positive effect of redistributive transfers on aggregate output.

The measurement of those MPLs is still relatively sparse compared to the large body of work on MPCs. More recently, though, there has been progress in the area, including via careful analysis of data from lottery winners — surely recipients of a windfall unconnected to employment — to obtain tight estimates of such propensities. We discuss next what those different methods imply for one's view of MPLs in the U.S. economy.

Measuring MPL with a Quantitative Model

Our paper bases estimates of MPLs on a richly specified model of labor decisions of U.S. households. A key element of the model is that households have dual earners: One is a primary earner who either works full time or not at all, and one is a secondary earner who may flexibly choose to work any number of hours.

This corresponds to the model estimated in the 2016 paper "Consumption Inequality and Family Labor Supply." In such an environment, for the most part, the primary earner has a job whenever available, and the secondary earner is most likely to react to changes in wealth.

Another key assumption is that disutility from work — which captures the full range of forces that make working outside the home difficult — differs across individuals, as noted in the 2015 paper "Wealth and Labor Supply Heterogeneity." This accounts for the fact that low labor force participation is concentrated in the lowest wealth quantile. This heterogeneity could capture many factors, including disability, long-term effects of incarceration on employment opportunities[3] or any other barriers to employment.

Finally, the model allows for strong means-tested consumption support — which is present in the U.S. welfare system — and thus a source of disincentives to work for those in the lowest income groups. All these elements are set to allow the model to match parameters estimated in the "Consumption Inequality and Family Labor Supply" paper, as well as key moments of the joint distribution of wealth and employment in the U.S.

Under those assumptions, our paper finds that MPLs are likely to vary substantially across the U.S. wealth distribution. In particular, MPLs are fairly low among households with low wealth, among which primary earners are more likely to not work. The quantitative model implies that households at the bottom of the wealth distribution would on average be willing to reduce earnings by 8 cents for each dollar that they receive as a transfer. In contrast, households at the middle of the wealth distribution would reduce their earnings by 5 cents.[4]

Labor Lessons from Lottery Winners

One might worry that the MPL estimates in our paper depend heavily on detailed modelling of labor supply decisions. Alternatively, recent work aims to obtain a "model-free" estimate of MPLs using data from lottery winners. Reactions to lottery winnings capture the effect of a mostly unexpected windfall on individual behavior without other incentive effects mixed in. As such, they provide a close-to-ideal laboratory to isolate the effects of increased wealth on labor-market decisions.

A very recent example is the 2021 working paper "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income." It finds that for each $100 in lottery winnings, households reduce per adult family earnings by $2 per year in the

subsequent five years. Also, the probability that an individual remains employed after receiving $100,000 in a lottery declines by close to 3.6 percent. Similar to our paper, it finds that lower income households are more likely to stop working.

Those estimates are based on the microeconomic behavior of households. To translate this to aggregate outcomes, one needs to take a stance on labor-market functioning and wage setting. We turn to this next.

MPLs and Sticky Wages

MPLs play an important role if workers can choose whether to work at a given wage. However, while workers can always look for a job, whether they get one or remain unemployed depends on firm hiring decisions. This means that, to use MPLs to assess the impact of a transfer policy, one needs to have a sense of how unemployment is determined by both sides of the labor market.

This issue is most likely to be pressing in environments where unemployment is high and employers do not perceive an advantage in hiring from a larger pool of eager-to-work people at lower-than-normal wages (that is, wages are downwardly sticky). In its most extreme form, this can make individual participation decisions completely irrelevant for aggregate labor market outcomes.

Downwardly sticky wages have received increased attention since the Great Recession and have been found to be operative in detailed studies using various forms of data, such as microeconomic data and regional data.[5]

Those studies also emphasize that there is asymmetry in the degree of stickiness. For example, my 2021 paper "Local Scars of the U.S. Housing Crisis" — co-authored with Saroj Bhattarai and Choongryul Yang — finds substantial wage gains in counties that experienced faster housing booms and employment growth before 2006, but do not find commensurate wage losses when those same counties experienced steeper housing and employment declines in subsequent years.[6]

Yet, the current situation is not one where *downward* stickiness in wages is likely to be particularly binding. In fact, wages have been rising, making the current situation one where labor supply considerations are more plausibly a constraining factor. With that in mind, we now turn to some simple calculations of whether the wealth mechanism is a plausible channel to explain low labor force participation in the recovery from the pandemic.

Application to the Pandemic Recovery

According to the Bureau of Labor Statistics, the employment-to-population ratio in the U.S. was 58.7 percent in August 2021, down from a peak value of 61.1 percent in February 2020 but up from the deep trough of March 2020. This corresponds to a reduction in the number

of employed workers from about 158 million in February 2020 to about 153 million in August 2021.

To gauge the wealth effect mechanism, we multiply the average MPL reported in the "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income" paper by the total transfers received by households. This assumes that those transfers were evenly spread across the population and that the average MPLs that we use are similarly representative. While the former is not true for all of the programs, it provides a useful ballpark value to assess whether the wealth channel could be relevant.

We include in our calculation not only the expanded child tax credit and direct payments to households, but also student loan deferrals, disaster and housing relief, earned income tax credit (EITC) expansion, and unemployment expansion. We include programs such as unemployment insurance that condition transfers on employment status or income since they also increase household wealth.

In particular, our calculation isolates the wealth effects of those measures and thus does not include the direct incentives or disincentives to work that they might generate.[7] We rely on estimates of the net cost of those policies from tabulations by the Committee for a Responsible Federal Budget encompassing all COVID assistance packages passed since 2020.

Together, those programs amounted to transfers of close to $2 trillion to households — or about $16,000 per household — approved over the course of 2020 and 2021. A direct application of estimates in the "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income" paper implies a reduction in the employment-to-population ratio of 0.58 percentage points relative to a baseline where those transfers had not happened but all other elements of the economic environment had remained the same (including pre-pandemic household wealth). This corresponds to one-fifth of the shortfall in employment between February 2020 and August 2021.

It is important to emphasize that this value could be underestimated, since discussions of persistent large expansions of U.S. federal transfers are ongoing. To the extent that individuals may expect those to be approved (even if only in part), this may lead them to similarly reduce current labor supply in anticipation of those additional transfers.

A final word of caution on those calculations is that while some of those expenditures might be absorbed by a persistent increase in U.S. debt, their large sizes suggest that they may also need to be absorbed by higher tax rates over time. Those higher tax rates would have the opposite effect on labor supply by reducing the net present wealth of taxpayers. To the extent, however, that those effects are deferred over a long period of time, they are likely to have a comparatively muted effect on current labor supply decisions.

Discussion

A general economic mechanism implies that transfers can lead households to restrict their labor supply, even when uncoupled from incentives to work. Some simple calculations suggest that such a channel can account for as much as one-fifth of the shortfall in labor market participation between February 2020 and August 2021.

Returning to a point we stressed at the outset, this shortfall cannot automatically be taken as an indication that those transfer policies have a negative effect on the economy. For one, those transfers might also be making up for lost income due to the pandemic. More generally, the aim of these programs was in fact to support workers as they elected care for family members or to retool workers for new jobs that fit better as we exit the pandemic. While none of these activities show up clearly in conventional measures of economic activity, it is clear that fiscal policymakers chose programs that support these activities for reasons that are widely valued by society.

The extent and the particular ways in which unconditional transfers have might have allowed individuals to make better use of their time is still an open area of research. As proposals for an expanded social safety net and even universal basic income-type programs become more widespread, it becomes increasingly important to have a good understanding of those effects.

Felipe Schwartzman is a senior economist in the Research Department at the Federal Reserve Bank of Richmond.

---

[1] *More on this topic can be found in the 2021 article "Benefits Cliffs."*

[2] *This research was also summarized in my 2017 article "Does Redistribution Increase Output?," co-authored with Athreya, Owens and Jessie Romero.*

[3] *For more on this topic, see the 2021 working paper "Incarceration, Earnings and Race."*

[4] *Those figures are found in Table 2 of my paper in the column that says MPL. However, they refer to changes in earnings, which is denoted MPE ("marginal propensity to earn") in subsequent literature. Here, we use MPL to denote the change in hours worked or employment rather than earnings.*

[5] *An example using microeconomic data would be the 2021 paper "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," and an example using regional data would be my 2021 paper "Local Scars of the U.S. Housing Crisis," co-authored with Saroj Bhattarai and Choongryul Yang.*

[6] *A summary of this finding is also available in my 2020 article "Will COVID-19 Leave Lasting Economic Scars?," co-authored with Tim Sablik.*

[7] *In particular, their wealth effects will persist even after unemployment benefits stop being paid for as long as household's cash balances remain higher, so the timing at which benefits stopped*

in different states will be unlikely to help identify those effects.

Subscribe to Economic Brief

Receive an email notification when Economic Brief is posted online:

Email address

Subscribe

Contact Us

RC Balaban
(804) 697-8144