Note: Comments and discussion cover two papers presented at the Session 1 of Summer 2020 BPEA conference on labor markets and consumer spending.

Comments and Discussion

COMMENT BY

JONATHAN A. PARKER Following the identification of a novel form of coronavirus in China at the end of 2019, COVID-19 has spread rapidly around the world causing death and economic destruction. On March 13, 2020, with hundreds of new cases identified each day (soon to be thousands), the United States declared a national emergency. In response both to the spreading virus and to government-ordered partial shutdowns, significant swaths of the US economy simply stopped during much of April. As the spread of the virus was slowed in the late spring and early summer, the reductions in economic activity have been partially reversing. These two excellent papers present some of the first broad-based, quantitative measurement of the massive disruption and partial rebound of employment and consumer spending during the first few months of this pandemic recession. They should be a good guide for what is happening now at the end of July, as new cases currently number in the tens of thousands.

In this discussion, I will first briefly describe how aggregate consumption and income have reacted during these first few months that followed COVID-19 reaching the United States, and then compare these measures to the average consumption and income responses documented in each of the two papers. But the main contributions of these papers lie in the careful analysis of the heterogeneous impact of the pandemic using high-quality, large microeconomic data sets.[1] So, second, I will emphasize what I take as the two main lessons from the combination of the papers. During this initial period of the pandemic, the economic collapse is almost entirely due to the

1. Both papers treat the related literature well, and I choose not to use my space comparing these papers to other rapid-response analyses of the economic effects of COVID-19 using micro data.

Figure 1. Real Personal Disposable Income and Personal Consumption Expenditures

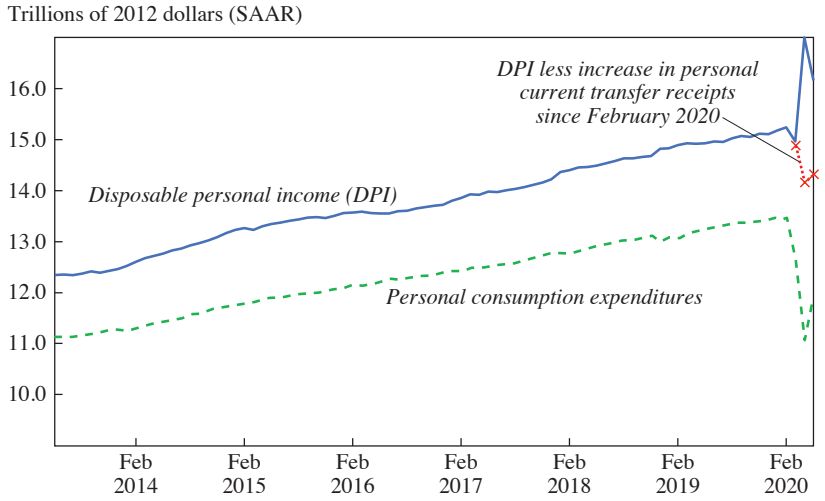

Sources: Bureau of Economic Analysis and authors' calculations.

pandemic directly. That is, whether one defines the shutdown of some sectors of the economy "demand" (e.g., people do not want to consume certain goods) or "supply" (e.g., certain firms cannot produce), the key point is that the papers show that the collapse in consumption is not due to, or amplified by, current income losses, and that the declines in employment are not due to, or amplified by, current low income or liquidity. The other lesson is that this lack of observable propagation through incomes and reduced demand is significantly due to the large policy response. These lessons apply only to these first few months. The pandemic is continuing, and I expect that this recession will slowly turn from a pandemic shutdown into a more typical recession and exhibit the usual economic propagation of recessions through demand channels.

Since these lessons lean more heavily on the consumption results, in the last part of my discussion, I will highlight a few other interesting findings about the decline in employment, and conclude with a few thoughts about real-time analysis by academic economists and economic policy going forward.

Figure 1 shows the disruption that this caused the economy as documented in aggregated statistics. Aggregate personal consumption expenditures (the dashed line) fall by about 6 percent from February to March and

then decline precipitously from March to April by 12.2 percent, reaching a level not seen since about seven years earlier. Both to emphasize how rapid and large this decline is, and because I hope never to have the opportunity to write a sentence like this again, let me note that this implies that real personal consumption expenditures declined by 79.1 percent at an annual rate between March and April 2020. Consumption subsequently rebounded in June, rising by 8.1 percent to roughly 11 percent below its February level.

These numbers are broadly consistent with the patterns shown in Cox and colleagues mapped into monthly averages, with two exceptions. First, the high-frequency nature of the paper's data—a significant contribution of the paper—shows some evidence for a spike in spending that occurs before the shutdown, possibly as people stock up ahead of expected increases in infection rates, store closures, and shelter-in-place orders.[2] Second, the paper documents a larger decline in consumption from February to March and a smaller decline from March to April than in the Bureau of Economic Analysis (BEA) data. In favor of the results in Cox and colleagues is that the month-to-month timing of consumption expenditures in official statistics is not particularly reliable. A major source of the data, for example, is a survey of retail establishments about sales volumes in which establishments can choose different time horizons over which to report sales amounts, overlapping horizons that must then be unpacked by the BEA to create monthly data. On the other hand, the Chase data capture only Chase customers and omit certain types of consumer spending. Using account-level data requires that one infer whether an outflow is consumption, saving, or debt payment from the observed counterparty. Paper checks do not have readily observable counterparties, nor do electronic funds transfers (EFTs) in this paper. One might also be concerned that the pandemic changed the means of payment for different types of consumption.[3]

Figure 1 further shows that disposable personal income (solid line) falls slightly in March and then *rises* by 13.6 percent in April before falling back down slightly in June. The personal savings rate—the difference between the two series as a percent of disposable income—rose to more than 30 percent in April, a number consistent with the unprecedented use of the word *unprecedented* during this pandemic.

2. See also Baker and others (2020).

3. That said, cash withdrawals are measured in consumption both before and after the pandemic, so that switches in the composition of spending between cash and cards should only affect the allocation of spending to categories (Cox and colleagues, figure 4).

Figure 2. Consumption and Income by Ex Ante Income Level

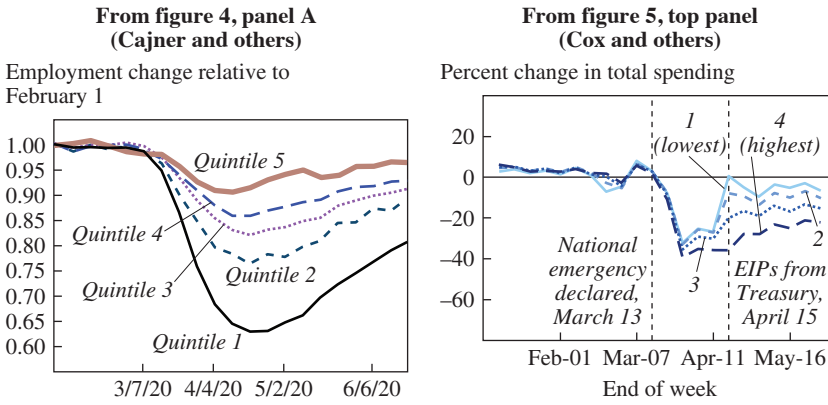

Sources: ADP anonymized payroll records and authors' calculations (left); JPMorgan Chase Institute (right)

However, the increase in disposable personal income looks like good news, and nothing like the labor market disruptions documented by Cajner and colleagues. The reason for this discrepancy is that government transfers increased by $231 billion from March to April ($2.8 trillion at an annual rate!) mostly due to the disbursement of economic relief payments. The final line (solid with *x*'s) on figure 1 removes the increase in current transfer receipts since February from the disposable personal income series and shows that income less these transfers declined by 2.3 percent in March and then 4.9 percent in April before slightly rebounding by 1 percent in May.

These declines in income are large for one-month movements, but are still slightly lower than one might infer from Cajner and colleagues. The most likely cause of this discrepancy is the difference in populations. The BEA data include government workers and retirees for example, whose regular incomes, as best we know, experienced less of an impact early in the pandemic. This highlights a difference not just between the aggregate data and the population studied by Cajner and colleagues, but also between the populations studied by the two papers. That said, I will now pretend the papers cover the same people and measure what we want them to measure, two assumptions that appear reasonable given the size of the pandemic shock and the point I want to emphasize.

Figure 2 simply reproduces the two figures in the two papers that show changes over time by quintiles or quartiles of the ex ante income

distribution. The figure on the left, from Cajner and colleagues, shows that there are dramatically larger declines in employment for ex ante lower income workers. The figure on the right, from Cox and colleagues, shows that there are not larger declines in spending for ex ante lower income workers. The implication: *the initial aggregate collapse of consumption during these months was driven by the unwillingness or inability of people to consume rather than by declines in income.*

The second main point: the most important reason that low-income (harder hit) households have not on average had to cut consumption on average by more than high-income households appears to be the increase in government transfers (economic relief payments, automatic stabilizers, and extended UI benefits). While this point is suggested by figure 1, Cox and colleagues show two important pieces of evidence in favor of this conclusion. First, consumer spending jumps up significantly at exactly the same time that most of the Economic Relief Payments—a major part of the policy response—were disbursed. Second, the paper shows that low-income households maintained substantial liquidity during these first few months despite significant income losses.

Let me note one caveat about the evidence for these conclusions. As shown in the left panel of figure 2, Cajner and colleagues show that employment has rebounded more strongly for ex ante low-income workers (although it remains below its prepandemic level). Further, the paper also shows that wage cuts are more common among high-income workers. Thus, my reading may be exaggerating somewhat the differences in income losses by ex ante income level. If there were only small differences in income, then we would expect little difference in spending responses across income levels and be less confident in the conclusion that income losses in general were not substantially holding back the economy.

My juxtaposition of the results of the two papers in figure 2 can be complemented by a similar comparison of the set of results in the two papers by industry. The pandemic shut down certain industries, and both papers nicely document how this has caused quite different income losses across workers. Yet again, we see little differences in consumption of households who work in different industries.

This conclusion implies that for the first few months of the pandemic in the United States, the goal of policy should have been insurance rather than stimulus. And policy largely met this insurance need. Policy surely also contributed to the lack of economic damage from a collapse in spending, but it did not stimulate the economy beyond this point, which I think is appropriate. When from a public health perspective (and so a welfare

perspective), it is optimal to shut down some sectors of the economy, then there is a reduced multiplier from government transfers and income support. In the typical recession, stimulus tends to raise purchases the most for the goods for which demand fell the most. So stimulus tends to generate spending that leads to hiring or maintaining employment for the workers most affected by the recession, who then tend to turn around and maintain consumption instead of cutting it. This is the Keynesian multiplier. However, when some employers are shuttered for health reasons, no stimulus is spent there, so any increase in demand and in resultant incomes go to those workers and industries who are already the least affected.

Further, unlike in a typical recession, when a sector of the economy is temporarily shut down, everyone in that sector with the same skills is out of work at the same time. While some workers can gain employment by moving across industries, if the shutdown is temporary, there is little benefit to having people searching for work which requires employers to on-board and train people whose skills are a poor match for the job at hand. Instead, as emphasized by Guerrieri and others (2020), fiscal targeted transfers are an important part of optimal policy as pandemic insurance. Only once the pandemic is past and as the economy reopens, to the extent that we are in a recession or a slow recovery, then more traditional demand stimulus may be beneficial.

There are also many more fascinating details in each paper, but I only have space to discuss two and will focus on Cajner and colleagues, which I have focused slightly less on up to now.

First, data on average wages show that they have risen substantially in the crisis so far, a point that has received a fair bit of attention. One can only measure wages for employed workers. Cajner and colleagues show that the average wage rises precisely because, as I focused on above, low-wage workers disproportionately lost their jobs. The paper shows that in fact, wages for continuing workers are on average unchanged through these first few months of the recession. The paper also shows that this constant average wage masks lots of different wage changes, and indeed a substantial share of workers have experienced wage reductions.[4] This finding sheds light on the theoretical factors that may constrain wage reductions in typical recessions. In particular, menu costs models predict that wages are more flexible in response to large shocks. Further, in such models, aggregate wage adjustment is slowed by strategic complementarities

4. For example, most senior faculty at MIT experienced a (temporary, we hope) wage cut for 2020, a first as far as I know.

and non-simultaneous adjustments. Both modelling ingredients thus predict that wages are more flexible in response to large and simultaneous economic shocks to firms, which is what this early evidence appears to show.

Second, the pandemic has had significantly different impacts on workers by gender. Cajner and colleagues show a much larger decline in the employment of women than men. Further, this difference is largely unrelated to the fact that women and men tend to work in different industries and at different sized firms. As such, in addition to the unequal burden of childcare and housework as the pandemic has shuttered schools, women may experience longer-lasting and more serious consequences from the labor market shutdown (Alon and others 2020).

Before concluding, I wanted to both praise and make a few suggestions for the conduct of research in this new world in which academic economists work with private-sector companies and conduct nearly real-time analysis, sometimes now directly for high-quality journals like the BPEA. Real-time analysis used to be purely the purview of newspapers, Gallup-type survey firms, and economists in bank research departments that had access to data and produced analyses for clients. And these organizations do still tend to produce analyses of important economic events first.[5] These early analyses partly lay out key questions and partly set narratives that persist in our understanding of events. The involvement of academics in this process is a boon. We can expand the resources available for these analyses, and also add a set of skills and knowledge—about theory, causation, and economic behavior—that can improve these analyses.

But these benefits should also involve some changes in how we operate. First, our usual strengths—the added value of academics—is about getting things right, at the cost of being slow. We are often trusted because we are correct, which involves being diligent, careful, and taking our time. We have to protect that trust, which means being clear that rapid analysis of firm data is not the same product as established research based on painstaking analysis. To be clear, I praise both papers in this regard. Each is extremely careful to delineate its strengths and to clearly state caveats.

Second, we have to be careful ourselves not to take early narratives (like my main conclusions in this discussion) as final truths. As an example, the latest Commerce Department estimate is that the homeownership rate

5. Bank of America, for example, produced an analysis of the spending caused by the economic relief payments on April 22, within days of the first payments being distributed; Michelle Meyer and Anna Zhou, "COVID-19 and the Consumer: Data through April 16," Bank of America Data Analytics.

rose by more than 3 percent between the first and second quarters of 2020. The pandemic has played havoc with the collection of lots of economic data, and I will go out on a limb and say that this large an increase in homeownership is very, very unlikely. Another example is that it has taken academics many years and many papers to overturn the early media consensus that the subprime crisis caused the financial crisis. And we are still—more than a decade later—parsing the relative roles of lending standards, low interest rates, and optimistic beliefs in the housing bubble, which is great. We, and many of the first contributors, have a dogged persistence to refine early findings and get to the truth. But early findings are more persistent the farther one looks from the core researchers. So, to again try to be clear, I have no reason to doubt the results in these papers, but we as a field need to avoid first-impressions bias, and I look forward to updating and refining my understanding from future analyses of the data from more companies and from traditional representative surveys.

To conclude let me return to interpreting the main lessons of these papers. The dramatic aggregate declines in employment and consumption appear to be due to choices rather than responses to low incomes or liquidity. The income declines represent a combination of responses to the pandemic: government policies, motivated by a desire to stop or at least slow the spread of the disease, and human behaviors that incorporate the additional motivation of self-preservation.[6] The effects of the income declines on consumption appear, in the data so far, to have been largely mitigated by fiscal insurance policies.

REFERENCES FOR THE PARKER COMMENT

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt. 2020. "This Time It's Different: The Role of Women's Employment in a Pandemic Recession." Working Paper 13562. Bonn: Institute of Labor Economics. http://ftp.iza.org/dp13562.pdf.

Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan. 2020. "Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to COVID-19." ArXiv:2005.04630v1. https://arxiv.org/pdf/2005.04630.pdf.

6. My reading of the literature so far is that, given the individual responses that already occur in response to the disease, there are few medium-term economic costs of government policies that shut down economic activity, at least on average outcomes, and there are substantial benefits in terms of health and lives saved from the disease (Andersen and others 2020; Correia, Luck, and Verner 2020).

Baker, Scott R., Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. 2020. "How Does Household Spending Respond to an Epidemic? Consumption during the 2020 COVID-19 Pandemic." *Review of Asset Pricing Studies* 10, no. 4: 834–62.

Correia, Sergio, Stephan Luck, and Emil Verner. 2020. "Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu." Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561560.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" Working Paper 26918. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w26918.

GENERAL DISCUSSION Adriana Kugler appreciated Cox and coauthors' simulation to estimate changes in income from stimulus policy, and she considered the possibility that Cajner and coauthors could employ a similar simulation to estimate the effects of Paycheck Protection Program (PPP) payments for firms. She postulated that PPP payments could help explain the trajectory of firm size and the nature of employment rebounds being primarily recall-driven (as opposed to being driven by new hires). She proposed that researchers explore the timing differentials that occurred in the implementation of the stimulus to verify this idea.

Kugler also commented on the breakdown of employment changes by gender in Cajner and coauthors' presentation, wondering if perhaps the idiosyncratic effects of the pandemic within individual sectors have been responsible for job losses being concentrated against women. She raised the possibility that using three-digit industry codes, as opposed to two-digit or one-digit codes, could be useful in answering this question.

Lastly, Kugler turned back to Cox and coauthors' paper, suggesting the potential for this study to focus more strongly on differences in pandemic unemployment assistance and relief transfers by state. Kugler proposed a more in-depth exploration of how individual states differed in their program implementation timing, worker composition, welfare infrastructure, and welfare qualification criteria and how those differences drove the results of the paper. Kugler concluded by urging both groups of authors to detail more explicitly the effects of the Coronavirus Aid, Relief, and Economic Security (CARES) Act and general government policy on their respective analyses.

Olivier Blanchard reiterated the finding from Cox and coauthors that saving has increased at the top of the income distribution during the pandemic, noting that he found it particularly striking. He asked the authors

and discussant if they had any general predictions for the consumption behavior of the rich in the future.

Hilary Hoynes commended the authors of both papers for examining how their respective analyses found different results for different groups and for articulating important levels of heterogeneity in the trends they observed. She raised the possibility that certain individuals were left out of the samples in each study, a problem she believed was particularly concerning for Cox and coauthors. She questioned if economically disadvantaged Americans, who are disproportionately likely to lack the bank accounts necessary to be included in the JPMorgan Chase data, were properly represented in the study. Similarly, Hoynes pointed out that the lowest income quartile in the study was limited to those with $12,000 or more, again raising the possibility that the sample did not accurately represent the most vulnerable Americans. She echoed discussant Jonathan Parker's comment that cautioned against drawing strong conclusions based on this issue and promoted the papers to be presented later in another conference session (Han and coauthors and Bitler and coauthors) for their results that focused on the poorest Americans.

David Wilcox continued along a similar line of thought. He asked if Cox and coauthors had any information regarding indicators of financial distress for the households in their sample, such as potential delinquencies on rent or a mortgage. He cautioned against concluding that relief measures had succeeded in staving off financial distress in the absence of measures of these indicators to confirm such a conclusion. Wilcox also noted that communities of color have been disproportionately harmed by COVID-19 and asked if it would be possible to examine the data along racial and ethnic lines to draw out additional insights.

Daron Acemoglu suggested to Cox and coauthors that they could use a shift-share composition analysis to explain the differences in their results across income groups. He noted that certain types of consumption would decline more than others because of social distancing (among other pandemic-related factors) and that those types of consumption are not homogeneously distributed across income brackets. Acemoglu proposed that understanding how the composition of different income groups' consumption was affected by the pandemic will be important for understanding changes in savings and consumption as the economy turns toward recovery.

Claudia Sahm was unable to comment directly due to technical difficulties, but moderator James Stock summarized from a comment she posted via the teleconference platform. According to Stock, Sahm said she strongly

disagreed with a sentiment expressed by Jonathan Parker urging moderation in the responses of economists, and she exhorted the conference to take action with respect to solving the pressing crises of the pandemic.

Wendy Edelberg said she agreed with Sahm that economists had an imperative to act and also agreed with Parker in cautioning against drawing conclusions too quickly. Turning to her central point, she argued that if one looked at how much consumption by low-income workers changed relative to high-income worker consumption, one might be led to believe that consumption in low-income households was not particularly affected by changes in income. She noted that such a conclusion would upend standard conceptions of marginal propensities to consume and how they differ among the rich and the poor. However, she pointed out that such an understanding would largely require ignoring the actual levels to which low-income consumption fell. Edelberg asked whether, by looking at how much consumption fell for low-income individuals, economists could gain new understandings of marginal propensities to consume among the poor. She concluded by noting that such lessons could have implications for the design of future stimulus policies and for determining whether or not stimulus payments to low-income individuals would largely be spent or saved.

Alessandro Rebucci asked if it was possible to determine the extent to which the pandemic differently affected the markets for goods and for services. He noted that services have been affected more than goods. He continued by commenting on the importance of understanding what has driven saving behavior in the pandemic, and he asked if increases in saving have been more due to precautionary saving in response to increased uncertainty or to declines in nonessential and conspicuous consumption, like vacation spending, due to the lockdowns and travel restrictions. He argued that understanding these drivers is important to form expectations about the recovery and also for policy design.

Ryan Decker, a coauthor on the Cajner paper, represented his colleagues in answering questions. He first noted, in relation to Sahm's comment and the related discussion, that while there has been pressure to release results quickly, he felt confident in his team's ability to work with ADP data. To illustrate this point, he noted that they have released papers using those data going back to 2018 and have other forthcoming papers that have been subject to rigorous academic scrutiny. He credited Wilcox for helping guide the research team to using the ADP data set.

In response to questions about the PPP, Decker said that his team's paper did not specifically examine that initiative. He pointed toward work that David Autor, Crane, and colleagues were presenting the same day at an

Automatic Data Processing, Inc., conference, examining PPP and analyzing how small and large businesses had different experiences with the program.

Lastly, Decker responded to an earlier comment about industry coverage, affirming that his data set was comprehensive across industries, and said that his team's findings about employment differences by gender held true even within detailed industries.

Peter Ganong, a coauthor on the Cox paper, fielded questions for his research team. He directed attention to a figure from his presentation, a bar plot showing changes in debit card spending, income without transfers, and income with transfers ("Estimated changes in income and spending," on figure 13 of his team's presentation).

First, Ganong answered questions regarding the representativeness of the sample in his study, noting that it was comprised of bank account data, and that roughly 95 percent of Americans have bank accounts.[1] He commented that this obviously left out some Americans, particularly low-income ones. Additionally, he affirmed the point raised by Hilary Hoynes, that individuals had to have at least $12,000 in labor income to be included in the sample. As a result, if, prior to the COVID-19 pandemic, an individual earned less than $1,000 per month, they would not be represented in the study sample. Ganong elaborated, stating that the lowest quartile of individuals in the study had annual post-tax labor incomes between $12,000 and $24,000, so while many low-income individuals were represented in the study, those with the lowest incomes were not.

Second, Ganong turned to questions regarding how his study calculated income changes, and what assumptions he and his team made regarding the receipt of unemployment insurance and stimulus. Briefly, he noted that while not all stimulus checks had gone out at that time (referring to economic impact payments), enough had for that fact not to be a large source of uncertainty. What was more important, Ganong said, was the fact that some states have been slower than others in processing unemployment insurance claims and in issuing unemployment insurance payments. He noted that their study does not assume that everyone left unemployed in the pandemic has received unemployment insurance but instead uses information from the Department of Labor to infer the fraction of unemployed Americans receiving unemployment insurance. He stated this was roughly 50 percent in April and 75 percent in May.

1. Economic Inclusion, "FDIC Survey of Household Use of Banking and Financial Services," 2019, https://economicinclusion.gov/surveys/.

Finally, Ganong turned to questions regarding heterogeneity in the team's sample. Noting that the average consumption in the lowest quartile of the sample remained *approximately* constant, this did not mean that every individual in this quartile had constant consumption but rather that some people increased their consumption and some people decreased their consumption. He then concluded briefly in response to Wilcox's earlier point, saying that the JPMorgan Chase Institute is engaged in studies regarding mortgage delinquency and that while those considerations were not in his team's line of research, they would be addressed by others soon.