Excess Savings and Inflation Dynamics After the Covid-19 Pandemic

Giovanni Sciacovelli

July 29, 2024

Abstract

One of the most important legacies of the Covid-19 pandemic is a massive increase in private assets, the so-called excess savings. Starting from this empirical observation, this paper studies the implications of excess savings for US inflation over the medium run. Relying on a calibrated HANK model, this paper finds that the expected drawdown of aggregate assets that follows the re-opening of the US economy will result in mild, albeit prolonged, inflationary pressures. The reason lies in the distribution of excess savings across households. The fiscal stimulus put in place by the US government has particularly benefited the accumulation of savings by low-income agents. As the economy re-opens, these households quickly decrease their assets due to their high MPCs, contributing to a timely recovery of consumption and inflation after their initial drop. Afterwards, the bulk of savings is left in the hands of higher-income households, who drawdown their assets smoothly over time, leading to contained inflationary pressures. Nonetheless, the paper shows that alternative fiscal policies may lead to different results: the longer the government decides to keep its increased debt levels unchanged, the larger the inflationary pressures out of excess savings.

1 Introduction

One of the most important legacies of the Covid-19 pandemic appears to be a massive increase in private savings. Figure 1 shows that personal savings followed a fairly stable path up until the beginning of 2020, and skyrocketed thereafter. To appreciate the magnitude of such increase, figure 1 displays the counterfactual path that personal savings would have followed if the growth rate of disposable and income and the personal savings rate had remained constant at their 5-year average before 2020[1]. The integral between the 'Actual' and 'Counterfactual' lines, amounting to approximately $2.5 trillions, represents *excess savings*: savings whose presence is due to the pandemic and that wouldn't have existed without it. This figure for excess savings is very large. As the chart makes clear, US households used to accumulate roughly $1.2 trillions in savings per year before the pandemic. This means that, in slightly more than one year since the beginning of the pandemic, households had already accumulated, on top of their normal savings, an amount of assets that would have been accumulated in over two years in standard circumstances.

Figure 1: Monthly Personal Savings

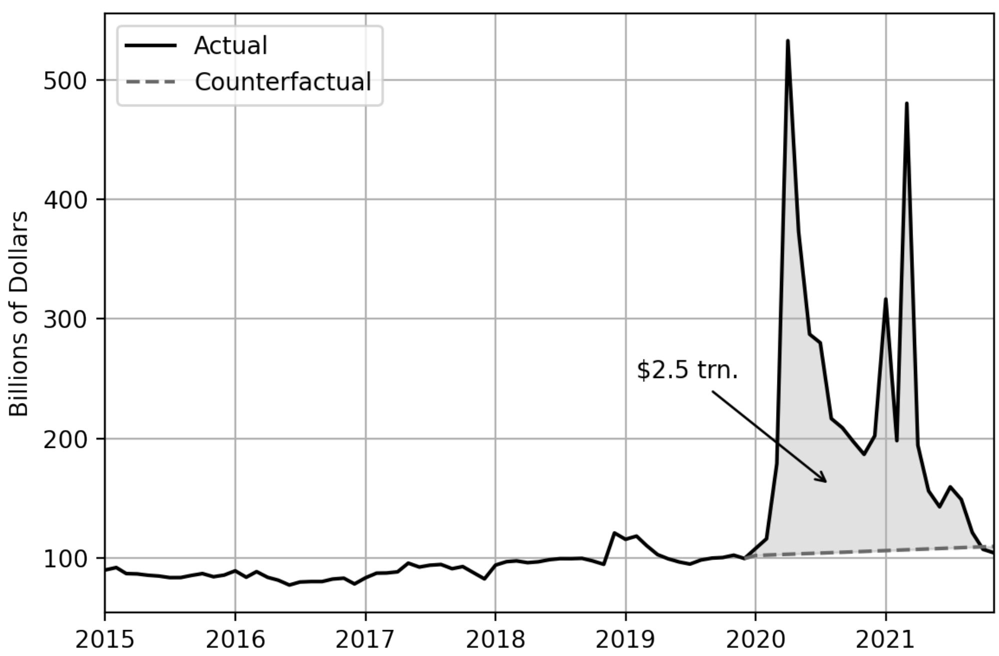

Notes: Data from FRED and BEA. The counterfactual path assumes that the personal savings rate and disposable income growth are constant to their 5-year average before 2020. As defined by BEA, personal savings amount to what is left once households have consumed and paid taxes. X-axis: monthly values from January 2015 to November 2021.

[1] Figure 1 extends the chart in Bilbiie et al. (2021).

What does this increase in private savings imply for the US economy over the medium run? In particular, one of the macroeconomic variables that has been most affected during the pandemic is inflation. As shown in figure 2, after stabilizing around 2% in the years before 2020, inflation first declined and then bounced to 6% after the beginning of the pandemic. Such high levels of inflation hadn't been experienced since the early 1990's, and have led to concerns about the possibility of having inflation expectations de-anchor (e.g. Blanchard, 2021; Summers, 2021).

Figure 2: CPI Inflation

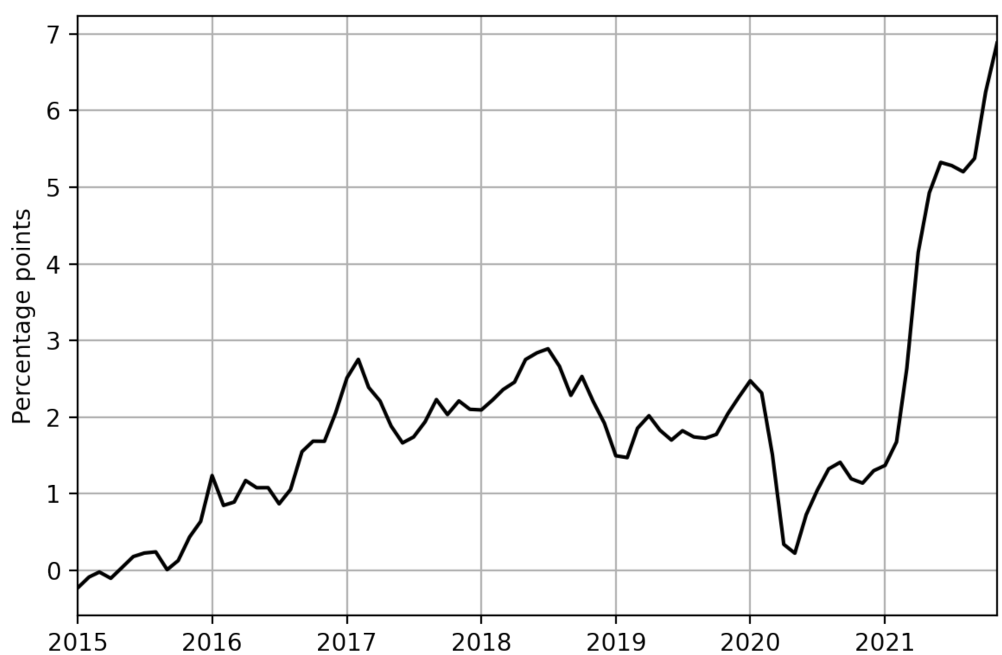

Notes: Data from FRED. Inflation is computed as year-on-year percentage changes of monthly CPI. X-axis: monthly values from January 2015 to November 2021.

In light of these empirical observations, the goal of this paper is to analyze how excess savings will impact the evolution of inflation, and to investigate how government policies may affect this impact. The rationale behind this research question is that there is a clear link between excess savings and inflation, which depends on the way in which US households will drawdown their savings. If assets that households have accumulated during the pandemic are quickly spent as soon as the economy re-opens, before the supply-side of the economy has fully recovered, then excess savings are likely to be a source of inflationary pressures. Differently, if households decide to only slowly spend their savings or to keep them, then they will represent less of a concern for inflation.

Crucially, the way in which households will make use of their savings relies both on the agents holding these assets and on government policy. If savings concentrate in the hands of high MPC households, then the first scenario is more likely: these agents will use their assets to increase consumption as soon as pandemic restrictions are lifted, creating inflationary pressures. However, government policy can act to affect households' savings choices: by changing its supply of government bonds, the government can affect the return on savings, altering the speed with which households will drawdown their assets.

Based on these considerations, which imply that capturing the distribution of assets across households with different MPCs is of utmost importance, I build a Heterogeneous Agent New Keynesian (HANK) model to answer my research question. The model is calibrated so that its steady state corresponds to the US economy in 2019, the year before the beginning of the pandemic. Afterwards, four shocks are sent to the model, reproducing both Covid-19 and the fiscal response implemented by the federal government.

Covid is captured by two shocks. The first one, following Bayer et al. (2020), is an exogenous increase in the probability of households being quarantined (a state in which they can consume but are not allowed to work). This shock is calibrated to reproduce the observed increase in unemployment starting from March 2020. The second one, following Carroll et al. (2020), is an exogenous drop in the utility of consumption, capturing the reduced desire to consume in the months following the beginning of the pandemic due to the fear of contagion. This shock is calibrated to match the observed drop in consumption.

Similarly, two shocks reproduce the fiscal response of the US federal government. On the one hand, the government stepped in to protect unemployed individuals by providing a top-up on the unemployment insurance administered by individual states. This policy is modelled as increased unemployment benefits up until September 2021, when such policy was dismissed. Secondly, the federal government has provided US households with unconditional cash transfers, which are included as unanticipated income shocks in the model.

The combination of these shocks leads to a drop in consumption, an increase in aggregate savings and, importantly, to an initial drop in inflation, quickly followed by a re-bound. Nonetheless, in the baseline scenario, the model does not predict a spike in inflation. The reason is that the combination of Covid and fiscal shocks, in line with the empirical evidence provided by Cox et al. (2020) and the JPMorgan Chase Institute[2], leads to an increase in savings that is particularly large for poorer households. As a consequence, as soon as the economy starts to re-open, savings are predicted to rapidly decrease, since poorer households display high MPCs. This leads to the quick bounce back of prices, but it is not sufficient to sustain a prolonged increase in inflation: after lower-income households have quickly decreased their savings and consumed, pushing inflation back to its pre-

[2] Charts available here.

pandemic levels, the remaining savings are in the hands of middle- and high-income households, whose lower MPCs imply mild inflationary pressures. Accordingly, the most important result of the model is to predict that excess savings will be a source of mild, albeit persistent, inflationary pressures. This result, however, does depend on two important assumptions.

The first one relates to the length of the pandemic, captured by the exogenous reduction in the utility of consumption. In the baseline scenario, this shock leads to the observed drop in consumption on impact and then reverts back, so that the shock is assumed to be fully absorbed in 2 and a half years. Admittedly, there is still much uncertainty about when the pandemic and its associated fear of contagion, will be over. It is then important to assess the sensitivity of the results to this assumption. Alternative calibrations show that excess savings would lead to higher inflationary pressures when assuming quicker recoveries of the desire to consume. The reason is that when households' fear of contagion decreases fast, consumption increases earlier on, during periods when unemployment is still high due to the Covid shock. Differently, when the drop in the utility of consumption is assumed to last longer, excess savings become less of a source of inflation: in this case households increase their consumption slowly over time, and only once unemployment has mostly recovered. Since this second scenario, in which the pandemic is fairly long-lasting, seems to be the most likely empirically, this analysis does not overturn the main result.

The second important assumption relates to future government policy. Importantly, the fiscal shocks that are sent to the model imply an increase in government debt. In the baseline calibration, it is assumed that the government reduces its debt increase in approximately 12 years. Again, this assumption is subject to debate, so that it is important to evaluate how results would be affected under alternative scenarios. I show that by keeping its debt high for longer, the model predicts much larger inflationary pressures. The reason for this result is twofold. On the one hand, when debt is left unchanged for longer, the government raises taxes slowly over time, allowing agents to spend more. On the other hand, high debt implies that the government needs to incentivize households to hold government bonds through higher interest rates, providing additional resources to be spent to households. The implications of this analysis are relevant, highlighting that government policy has the power to significantly affect how excess savings will affect inflation.

The remainder of this paper is organized as follows. Section 2 provides a review of related literature. Section 3 introduces the households' block of the model, the exogenous shocks, and shows how consumption and savings would be affected in a partial equilibrium setting. Partially out general equilibrium feedback is important to clearly identify households' choices. Section 4 moves on to the general equilibrium analysis, showing the main results of this paper, while section 5 concludes.

2 Related Literature

A rapidly growing literature is investigating the consequences of the Covid-19 pandemic on the macroeconomy. Instead of providing a comprehensive overview of this literature, I will only discuss the most relevant contributions for this paper.

Among early contributions, Guerrieri et al. (2020) provide a crucial analysis of the characteristics of the Covid shock: despite starting out as a typical supply shock, it can lead to demand shortages in a multi-sector economy. This paper does try to capture the demand shortages stemming out of the pandemic but, differently from the framework in Guerrieri et al. (2020), this is captured through a more parsimonious, reduced-form preference shock, since the focus of the analysis are households' savings decisions. Bayer et al. (2020) build a DSGE model that combines heterogeneous agents with the multi-sector environment introduced by Guerrieri et al. (2020), and make use of it to quantify the importance that government transfers had on reducing the drop in output in the immediate months after the beginning of the pandemic. A similar analysis is implemented by Carroll et al. (2020), where the impact of the fiscal stimulus is analyzed through a partial equilibrium exercise, and by Faria-e-Castro (2021), who implements the analysis in a TANK model. This paper shares with these contributions the goal of building a DSGE model to analyze the implications of the pandemic but, rather than focusing on quantifying the impact of fiscal policy, it assesses how fiscal policy, by contributing to the build-up of private savings, may affect inflation dynamics going forward.

Within empirical contributions, particularly important for these papers are the evidences provided in Chetty et al. (2020), Mongey et al. (2021), Ganong et al. (2021), Cox et al. (2020) and by the JPMorgan Chase Institute[3]. In particular, Chetty et al. (2020) and Mongey et al. (2021) show the heterogeneous impact that the Covid-19 shock has had on agents depending on their income level, with lower income households having a significantly larger probability of becoming unemployed after the beginning of the pandemic (this evidence is also reported in Eichenbaum et al., 2021). Ganong et al. (2021), Cox et al. (2020) and research by the JPMorgna Chase Insisute, on the other hand, investigate the heterogeneous response in savings behaviour across households in the initial months of 2020. The shocks sent to the model will be calibrated to match the facts emphasized by these empirical studies.

In spirit, this paper relates to work by Fornaro and Wolf (2021), which aims to understand the long-run consequences of the pandemic. Similarly, this paper investigates the legacies of the pandemic and what they imply for medium-run inflation dynamics. From a methodological standpoint, the model used in this paper is solved using the Sequence-Space Jacobian approached introduced in the

[3] Charts available here.

work by Auclert et al. (2021a).

3 Households & the Pandemic

This section introduces the household block of the model. Given that the goal of this paper is to understand the impact of excess savings, this block is of particular importance, since households are the agents that ultimately decide how to make use of their assets. To identify how they react after being hit by the combination of pandemic and fiscal shocks, this section shows households' impulse responses within a partial equilibrium framework. This simpler framework allows to better understand the mechanisms behind households' choices before moving to the general equilibrium setting of section 4. I will firstly introduce the households' problem, then detail the characteristics of the shocks that are sent to the model, and finally show the partial equilibrium impulse responses.

3.1 Households

The household block follows the workhorse one in Bewley (1977), with the main deviation being that households face idiosyncratic income uncertainty not only though exogenous labor productivity, but also through exogenous shocks to their quarantine status. In particular, the preferences of household i are represented by the utility function

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\frac{c_{i,t}^{1-\sigma}}{1-\sigma}\right] \tag{1}$$

where c_{it} is a generic consumption good, σ is the coefficient of relative risk aversion and $\beta \in (0,1)$ is the discount factor.

Given aggregate income Y_t, which assumed to be constant in this section, idiosyncratic productivity level e_{it} and quarantine status q_{it}, pre-tax labor income of household i is

$$y_{i,t} = q_{i,t} \cdot e_{i,t} \cdot Y_t. \tag{2}$$

It is important to mention two features of this equation. First, the quarantine shock q_{it} follows a Markov chain on the space $\{0,1\}$. Accordingly, households can either be quarantined ($q_{i,t} = 0$) and not perceive any labor income or be out of quarantine ($q_{i,t} = 1$) and perceive labor income. Second, when households are out of quarantine, their income is proportional to their exogenous productivity level $e_{i,t}$, which follows a Markov chain on the space $\{e^1, ..., e^S\}$, where S is the total number of productivity states.

Households are subject to the budget constraint

$$c_{i,t} + a_{i,t+1} = (1+r)a_{i,t} + y_{i,t} + \tau_{i,t} \tag{3}$$

where a_{it} are bonds, the only asset that can be traded in this economy, r is the real interest rate (assumed constant in this section) and $\tau_{i,t}$ are taxes. Taxes are administered in order to compensate quarantined households for their loss of labor income:

$$\tau_{i,t} = b_{i,t} + \mathbb{1}_C v_{i,t} + \mathbb{1}_U \tilde{v}_{i,t}$$

where $b_{i,t}$ is a standard unemployment benefit, while $v_{i,t}$ and $\tilde{v}_{i,t}$ represent the additional unemployment benefits and unconditional transfers implemented as a fiscal response to the pandemic, which will be more thoroughly discussed in the next section. It is important to note that in this baseline exercise households don't pay any taxes, and that $b_{i,t}$ and $v_{i,t}$ are only disbursed to quarantined agents.

Finally, in order to make sure that not all agents will be on the Euler equation, household debt is assumed to be bounded below:

$$a_{i,t+1} \geq \underline{a}$$

where $\underline{a}$ is the exogenous borrowing limit.

Recursively, the problem can be written as:

$$\begin{aligned} V_t(e,q,a) &= \max_{c,a'} u(c) + \beta \mathbb{E}_t[V_{t+1}(e',q',a')] \\ &\text{subject to} \\ &\quad c + a' = (1+r)a + y + \tau \\ &\qquad a' \geq \underline{a} \end{aligned}$$

3.2 Calibration

Table 1 summarizes the parameter values used in the calibration of the steady state in this partial equilibrium exercise. The time period is a month. The real interest rate is chosen to obtain a yearly interest of 1%, while the discount factor β is chosen to target a monthly average MPC of 25%. The coefficient of relative risk aversion is chosen so that the EIS is 0.5, a common value used in the literature (e.g. Auclert et al., 2021b). Flodén and Lindé (2001)'s estimates of US wages guide the choice of ρ and σ_ϵ. The parameters governing the probability of entering and exiting the quarantine state follow the work by Bayer et al. (2020), who have a similar set-up to the one presented in the

Table 1: Parameter Values

Parameter	Explanation	Value	Target/Source
r	Real interest rate	0.083%	Yearly rate of 1%
β	Discount factor	0.994	MPC of 25%
σ	Relative risk aversion	2	
ρ	Persistence of productivity shock	0.9915	Flodén and Lindé (2001)
σ_ϵ	Variance of productivity shock	0.40	Flodén and Lindé (2001)
$\pi_{q_{out},q_{in}}$	Transition prob. to quarantine	0.02%	Bayer et al. (2020)
$\pi_{q_{in},q_{out}}$	Transition to out of quarantine	50%	Bayer et al. (2020)
b	Unemployment benefit	40% income	
$\underline{a}$	Borrowing limit	0	

previous section. Unemployment benefits are chosen to replace 40% of labor income, in line with work by Guerrieri and Lorenzoni (2017) and Nekarda and Ramey (2020). The borrowing limit is set to zero, as in McKay et al. (2016).

3.3 Shocks

There are four shocks that will be sent to the calibrated steady state of the model. The first two aim to capture the main features of Covid: an increase in unemployment and a corresponding reduction in the households' desire to save due to fear of contagion. The second two relate to the fiscal response implemented by the US federal government to counteract the pandemic. Each of these shocks, and its calibration, are discussed in turn.

The first shock that is sent to the model endogenously increases the probability that households enter the quarantine state, following Bayer et al. (2020). To calibrate this shock, I proceed in the following manner. First, the unemployment shock is quantified as the excess unemployment in the months after the beginning of the pandemic relative to average unemployment rate in 2019. This is shown in panel A of figure 3, displaying both the actual excess unemployment and the one resulting from the model.

Second, since the incidence of the unemployment shock has differed widely across the income distribution (see, e.g., Chetty et al., 2020; Mongey et al., 2021), the intensity of the quarantine shock is income specific. Using the ASEC version of the CPS survey, it is possible to reconstruct the transition probabilities in and out of employment for different income levels in 2019 and in the first months of 2020 (until June). This way, I can compute the increase in the probability of entering

Figure 3

A. Excess Unemployment B. Excess Unemployment by Income

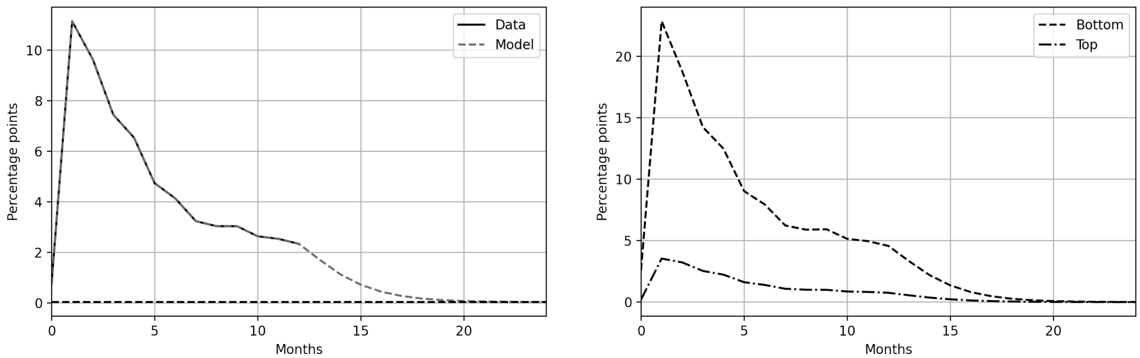

Notes: The solid black line in panel A displays the monthly increase in unemployment starting from March 2020 relative the the monthly average in 2019. The dashed grey line displays the monthly increase in unemployment relative to steady state in the model. Month zero corresponds to March 2020. The black line stops in March 2021, thereafter unemployment is assumed to follow an AR(1) process with persistence 0.5. Panel B shows the different incidence of the quarantine shock, focusing on the bottom and top income levels. According to the CPS-induced incidence ratios, top-income agents experienced a much smaller shock relative to bottom-income agents.

the unemployment state at different income levels after the beginning of the pandemic, expressing the increases relative to the probability increase of the highest income level. I then use these ratios to guide the incidence of the quarantine shock across households. Panel B of figure 3 shows the resulting unemployment level stemming out of the quarantine shock for the bottom and top income levels. As the chart shows, the resulting unemployment level is much higher for bottom-income households, since the CPS-induced incidence implies they were much more affected from this shock than top-income agents. It is important to stress that the calibration of this shock differs from the one in Bayer et al. (2020) along two dimensions. First, the shock is calibrated to match the aggregate increase in unemployment, rather than claims to unemployment. Second, the income-specific incidence of the shock follows evidence in the CPS survey, rather than the fit of a logistic function.

The second shock sent to the model aims at capturing the reduced desire to consume brought about by the pandemic. Following Carroll et al. (2020), this is achieved by reducing the utility that agents get out of consumption. Differently from the first shock, there is large uncertainty about how to calibrate this shock, both in terms of its magnitude and its length. To overcome these issues, I apply the following rationale. First, to choose the magnitude of the shock, I calibrate it so that the initial drop in consumption matches the one observed in the data (approximately 10% in the first quarter of 2020). Second, the magnitude of the shock is assumed to remain constant for 9 months, and to then decrease following an AR(1) process. The reason why the magnitude of the shock is constant for 9 months is to match the fact that, after approximately 9 months after the beginning of the pandemic, vaccines started to be rolled out in the US. The persistence parameter of the AR(1) process is subject to uncertainty: in the baseline calibration, its value is such that the shock is reabsorbed after 2 and a half years. However, given the uncertainty about the length of the reduced desire to consume, I will experiment with different values of this parameter in section 4.

The third and fourth shocks relate to the fiscal stimulus implemented by the US government to offset the negative effects of the pandemic. In particular, the stimulus provided to households consisted of two main components. On the one hand, starting with the CARES Act, the federal government stepped in to top-up the unemployment insurance provided by individual states. Until December 2020, the top-up consisted of $600 per week, disbursed for up to 13 weeks. It then dropped to $300 per week until September 2021, and was dismissed thereafter. Accordingly, my first fiscal shock is calibrated to reproduce the increase in unemployment insurance until September 2021. On the other hand, the federal government has provided US households with unconditional cash transfers. These have been implemented at three different times, each of which has differed with respect to the amount provided and the eligibility criteria[4]. In calibrating these transfers, I disregard the possibility of households having children (a condition that would entail the provision of additional resources) and, following Bayer et al. (2020) and Carroll et al. (2020), I assume each households is provided the amount promised to individuals. With respect to the eligibility criteria, the shock is calibrated to follow the ones that were set out for singles. Finally, in line with Bayer et al. (2020), the amount from these transfers is gradually provided to households in diminishing amounts, following an AR(1) process with persistence parameter 0.7.

3.4 Partial Equilibrium IRFs

Figure 4 shows the IRFs of consumption and savings to the four shocks discussed above, introduced one at the time. The solid black line represents the responses to a quarantine shock only. Panel A shows that, in line with standard expectations, consumption drops when more households are sent to quarantine and experience a wage loss. With respect to the response of assets, two effects are in place. On the one hand, households want to save more for precautionary motives: unemployment risk has increased, and all agents try to insure themselves against the possibility of being put in quarantine by saving more. On the other hand, more households are effectively put in quarantine,

[4]To be able to solve for the non-linear GE solution to my model in section 4, I disregard the second round of unconditional cash transfers provided by the Trump administration in December 2020. This is the smallest of the rounds in terms of amount provided to households, and was quickly followed by the much larger transfers by the Biden administration in March 2021.

Figure 4

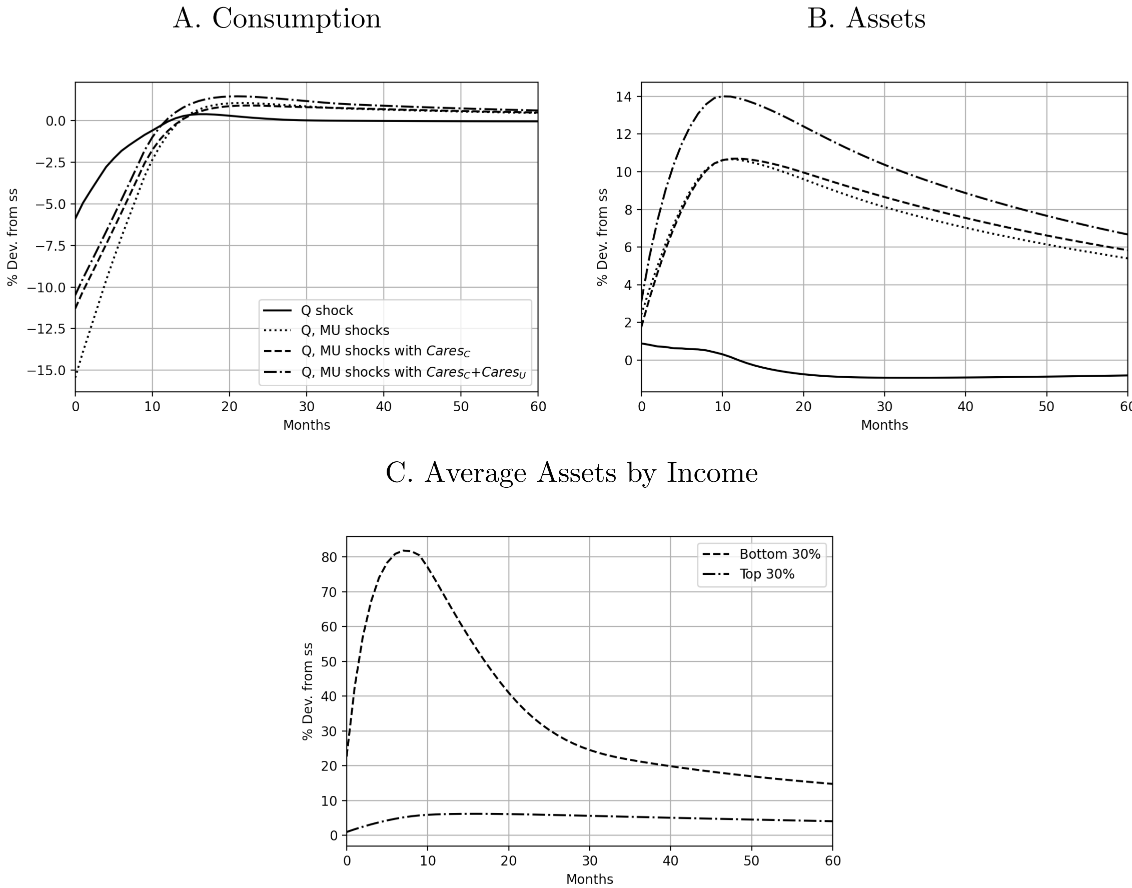

Notes: Panel A and B show the response of consumption and assets, respectively, to combinations of shocks. The solid black line represents the response to the quarantine shock. The dotted black line represents the response to the quarantine and marginal utility shocks. The dashed black line represents the response to the two previous shocks plus the conditional transfer shock, and the dashed-dotted line represents the response to all four shocks. Panel C shows the response of average assets held by the bottom (dashed line) and top (dashed-dotted line) 30% of the income distribution to the four shocks. Y-axis: percentage deviations from steady state values. X-axis: months after the shocks hit.

suffering a reduction in their income that reduces their possibility to save. Panel B shows that the first effect dominates initially, so that assets slightly increase. As time passes by, however, assets decrease as more agents become unemployment and as the risk of being quarantine diminishes.

The dotted black line shows the impulse responses once the second Covid shock is added to the analysis. In this case the drop in consumption is much larger than in the previous case, since households not only consume less because of precautionary motives and unemployment constraints, but also because their desire to consume decreases. Interestingly, and in line with Carroll et al. (2020), most of the drop in consumption is due to the marginal utility shock: the reason lies in the fact that this shock affects all households, while the former disproportionately affected lower income agents. The counterpart of the large reduction in consumption is the spike in personal savings, which increase widely relative to the previous shock. Disregarding additional general equilibrium feedbacks, the dotted black line represents the outcome of having Covid without fiscal intervention.

The dashed black line introduces the first component of the fiscal response to the pandemic. As shown in panel A, and in line with Bayer et al. (2020), conditional transfers greatly reduced the decline in consumption, decreasing the drop by approximately five percentage points on impact. Moreover, similarly to findings in Carroll et al. (2020), the fiscal response leads to an re-bound of consumption large enough for it to increase relative to the pre-pandemic steady state for an extended period of time. Panel B shows some interesting dynamics: relative to the case with no fiscal intervention, the path for savings is mostly unaffected. The reason underlying this movements is twofold. First, only conditional transfers have been introduced so far. Since most households becoming unemployed are low-income, and since low-income households have larger MPCs, the fiscal transfers are largely spent: consumption drops by less, while savings remain unchanged. Second, the conditional transfers implemented by the US government were so large that, for lower-income households, replacement ratios were above one. Consequently, conditional transfers were large enough to fully insure the poorer segments of the population against the unemployment risk. Accordingly, these households lose their precautionary savings motive, having less of an incentive to save.

Finally, the dashed-dotted line displays the response of consumption and assets when all four shocks are sent to the households. In panel A it can be seen that the additional resources provided to households lead to an even smaller drop in consumption. Moreover, consumption rebounds quickly and the increase relative to the pre-pandemic steady state is longer than 5 years. However, most of the action takes place on the savings side. Relative to the previous scenarios, assets increase greatly. The reason for this outcome is that, despite absolute top incomes were not eligible to receive these transfers, this stimulus was disbursed widely to middle and high income households. Since these households did not suffer much from the unemployment shock, they saved most of these additional resources to further insure themselves against the risk of being put in quarantine.

Panel C of figure 4 provides evidence about the increase of average assets for the top and bottom 30% income levels, relative to the pre-pandemic steady state[5]. The chart shows that lower-income households experienced a relative increase in private savings which was much larger than the one experienced by richer households. Importantly, this result is in line with the empirical evidence in

[5]These figures and the ones following refer to top and bottom 31.25%. For simplicity, I refer to them as top and bottom 30%.

Cox et al. (2020) and additional evidence from the JPMorgan Chase Institute[6]. The explanation for the dynamics shown in chart is that lower-income households, despite being disproportionately negatively affected from the unemployment shock, were compensated for their unemployment loss with a large top-up in their unemployment insurance, and additionally benefited from unconditional cash transfers. The reason why this chart is interesting is that it provides an explanation for the paths shown in panel A. We can see that, despite the initial increase in private savings, these are drawn down very quickly: low-income households have high MPCs, so that as soon as their desire to consume recovers (i.e., as the economy re-opens out of the pandemic), they rapidly make use of their accumulate assets to increase consumption. This fact explains the rebound of consumption above its steady state level in panel A: had savings been accumulated less in the hands of low-income households, the rebound would have been less sharp. Appendix A discusses the intuition behind this result further.

This exercise shows that, within a partial equilibrium framework, the combination of Covid and fiscal shocks produces results that are in line with the empirical evidence. In particular, the response of savings is mostly driven by the reduced desire to consume and expansionary fiscal policy, while the movements in consumption are sensitive to all of the four shocks. In particular, after dropping in the beginning of the pandemic, consumption is expected to recover and stay above its pre-pandemic steady state for an extended period of time. In order to make progress and understand whether this demand increase can lead to inflationary pressures, the next section builds the more comprehensive general equilibrium model.

4 Excess Savings and Inflation

This section shows the combined response of consumption, assets and inflation to the combination of Covid and fiscal shocks introduced in the previous section. The goal is to disentangle whether excess savings will be a source of inflationary pressures and, if so, to what extent. To do so, I will firstly discuss the different components of the model. Then, I will show the IRFs in the baseline scenario: this exercise shows that excess savings create inflationary pressures, but these are small in magnitude. Finally, given the great uncertainty surrounding the length of the pandemic and future fiscal policies, I will implement additional exercises to understand under what conditions excess savings are a more important source of inflation. This will turn out to be the case in which the US government decides to keep its debt level unchanged for long into the future — a plausible scenario.

[6]Detailed movements in liquid balances across different income levels of JPMorgan costumers can be found here.

4.1 Model

I build a Heterogeneous Agent New Keynesian (HANK) model with the minimal necessary features to capture the relation between excess savings and inflation.

Households. The household block slightly extends the one introduced in section 3. The preferences of household i are represented by the utility function

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \frac{c_{it}^{1-\sigma}}{1-\sigma} - \nu \frac{l_{it}^{1-\gamma}}{1-\gamma}\right]$$

where, relative to the utility in equation 1, households additionally receive disutility out of labor. The new budget constraint is

$$c_{i,t} + a_{i,t+1} = (1 + r_t)a_{i,t} + y_{i,t} + \tau_{i,t} + d_{i,t}$$

where there are three differences relative to the budget constraint in equation 3. First, $y_{i,t}$ is not anymore a fraction of constant output as in equation 2:

$$y_{i,t} = w_t \cdot q_{i,t} \cdot e_{i,t} \cdot n_{i,t}$$

with wage w_t being an equilibrium object. Second, households now have to pay taxes as well as possibly receive transfers from the government:

$$\tau_{i,t} = \begin{cases} -\tilde{\tau}_{i,t} + \mathbb{1}_U \tilde{v}_{i,t} & \text{if } q = 0 \\ b_{i,t} + \mathbb{1}_C v_{i,t} + \mathbb{1}_U \tilde{v}_{i,t} & \text{if } q = 1 \end{cases}$$

where $\tilde{\tau}_{i,t}$ represents taxes, which will be levied proportionally to the household's skill level $e_{i,t}$. $\tau_{i,t}$ dictates that only out-of-quarantine agents are subject to tax payments. Third, the new element in the budget constraint, $d_{i,t}$, captures dividends that are provided to households out of firms' profits. These dividends are assumed to be divided across all households, regardless of their quarantine status, proportionally to their skill level $e_{i,t}$.

Recursively, the problem can be written as:

$$V_t(e, q, a) = \max_{c,n,a'} \left\{ \frac{c^{1-\sigma}}{1-\sigma} - \nu \frac{l^{1-\gamma}}{1-\gamma} \right] + \beta \mathbb{E}_t V_{t+1}(e', q', a') \right\}$$

subject to

$$\begin{aligned} c + a' &= (1 + r)a + y + \tau + d \\ a' &\geq \underline{a} \end{aligned}$$

Firms. The firm sector follows the one in Auclert et al. (2021a), with sticky-prices à la Rotemberg and quadratic adjustment costs. In particular, final goods producers aggregated intermediate goods through a standard CES aggregator with elasticity of substitution $\mu/(\mu-1)$, while intermediate goods producers are monopolistically competitive with linear production function. Output produced by firm j is $y_{j,t} = n_{j,t}$, where $n_{j,t}$ is labor employed by intermediate producer j. Each firm j sets its price $p_{j,t}$ subject to quadratic adjustment costs

$$\xi_t = \frac{\mu}{2\kappa(\mu-1)}\left[\log\left(\frac{p_{j,t}}{p_{j,t-1}}\right)\right]^2 Y_t$$

where κ captures the slope of the Phillips curve and Y_t is aggregate output. With inflation being π_t, this set-up implies the following Phillips curve[7]:

$$\log(1+\pi_t) = \kappa(w_t - \mu^{-1}) + \frac{1}{1+r_{t+1}}\frac{Y_t}{Y_{t+1}}\log(1+\pi_{t+1}). \tag{4}$$

Dividends rebated to households are $d_t = Y_t - w_t N_t - \xi_t$: output net of costs.

Government. The government is composed of a fiscal and a monetary authority. The fiscal authority needs to satisfy its inter-temporal budget constraint:

$$B_t = (1+r_{t-1})B_{t-1} + \int (b_{it} + \mathbb{1}_C v_{it})q_{it}di + \int \mathbb{1}_U \tilde{v}_{it}di - T_t$$

where B_t is government debt and T_t are total taxes levied in period t. $\mathbb{1}_C$ and $\mathbb{1}_U$ are equal to one when the government implements conditional and unconditional transfers, respectively. Importantly, it is assumed that the government follows a rule for tax revenues:

$$T_t = \bar{T} + \phi_T(B_{t-1} - \bar{B})$$

where $\bar{T}$ represents the steady state value of tax revenues. This rule implies that the government does allow the debt level to move over the short run, but it uses taxes to make sure that the debt level returns to its steady state over time. In particular, the parameter ϕ_T governs the speed of adjustment: the higher is ϕ_T, the quicker the government will raise taxes after an increase in the debt level. Crucially, since there is wide uncertainty about the speed with which the US government will try to reduce its debt level, I will experiment with different values of this parameter.

The monetary authority is assumed to follow a simple Taylor (1993) rule, where the goal is to stabilize inflation: $i_t = r^* + \phi\pi_t$, where r^* is the steady state real interest rate. Following Bayer et al. (2020), the central bank does not react to deviations of output from its target value because of the complication of identifying output's target value when quarantine measures are in place.

Market clearing. The goods market clearing condition is:

$$Y_t = C_t + \xi_t$$

[7] See Appendix B for a detailed derivation.

Table 2: Parameter Values

Parameter	Explanation	Value	Target/Source
Households			
β	Discount factor	0.9948	$r = 1\%$
σ	Relative risk aversion	2	
γ	Inverse Frisch elasticity	2	
ν	Labor disutility coefficient	0.789	$Y = 1$
ρ	Persistence productivity shock	0.9915	Flodén and Lindé (2001)
σ_ϵ	Variance productivity shock	0.40	Flodén and Lindé (2001)
$\pi_{q_{out},q_{in}}$	Transition prob. to quarantine	0.02%	Bayer et al. (2020)
$\pi_{q_{in},q_{out}}$	Transition to out of quarantine	50%	Bayer et al. (2020)
b	Unemployment benefit	40% income	
$\underline{a}$	Borrowing limit	0	
Firms			
μ	Steady-state markup	1.2	
κ	Slope of Pillips curve	0.1	
Government			
ϕ	Taylor rule coefficient	1	
B	Bond supply	0.515	Liquid balances (FoF)

where $C_t = \int c_{i,t} di$. Simply, output needs is used to cover private consumption and adjustment costs. The bond market clearing condition is:

$$B_t = \int a_{i,t} di$$

which states that bonds held by households need to equal total government bonds. Finally, the labor market clears when labor demanded equals labor supply:

$$N_t = \int e_{i,t} n_{i,t} (1 - q_{i,t}) di.$$

Three prices are available to assure a solution: the price of the consumption good p_t, the price of labor w_t and the real rate r_t.

Calibration

Table 2 summarizes the parameter values used for the monthly steady state calibration. Most of the parameters referring to households have been discussed already in section 3. Exceptions are the new parameters related to the disutility of labor: ν is set with the goal of normalizing monthly output to 1; γ, the inverse Frisch elasticity, is set to 2, a common parameter in the literature (e.g. McKay et al., 2016; Auclert et al., 2021b). Similarly, the remaining coefficients referring to the other parameters in the firm and government blocks are commonly used in the literature. An important parameter to discuss, however, is B: total bond supply. The goal is to capture the distribution of savings across income levels as specified by the JPMorgan Chase Institute[8], since these represent the most detailed data available on the distribution of liquid balances since the beginning of the pandemic. Since the figures provided by the Institute refer to cash balances, B is calibrated accordingly, measured as the sum of checkable and savings deposits from the Flow of Funds[9].

4.2 Results

Figure 5 provides the IRFs to the combination of Covid and fiscal shocks. Consistently with the partial equilibrium analysis in section 3 and the empirical evidence, consumption drops and savings increase as a response to the shocks. In particular, consumption drops on impact by slightly less than 10%, and recovers to its previous steady state value in approximately one and a half years. This is a recovery which is slightly slower than the one observed in the data, where consumption recovered its pre-pandemic values after less than 18 months (the actual consumption path is shown in Appendix 9).

Panel B displays movements in aggregate assets. The chart shows two separate phases of the pandemic. The first one, lasting one year, is characterized by both a rapid increase and a consequent rapid decrease in assets. To understand this phase, it is important to consider the interest rate movements in panel D. In particular, as expected from the partial equilibrium exercise, the combination of Covid and fiscal shocks is such that households' desire to save increases massively (panel B in figure 4). Differently from that exercise, however, the increase in assets needs to be matched by a corresponding increase in bond supply by the government for the real rate to remain constant. Bond supply does increase, since the government needs to finance higher unemployment benefits to quarantined agents and its fiscal stimulus programs, but not as much to compensate the higher desire to save. Accordingly, the real rate drops in the months after the pandemic hits, reflecting the excess demand for bonds by private agents. The initial large drop and subsequent recovery are explained by the dynamics of the quarantine shock shown in panel A of figure 3. Excess unemploy-

[8] Data avaialble here.

[9] I use the figures in Table B.101 of the Flow of Funds, lines 11 (checkable deposits and currency) and 12 (time and savings deposits). This is a more restrictive sample than the one often used sample by Guerrieri and Lorenzoni (2017) and McKay et al. (2016), which also includes treasury securities, agency and GSE securities, municipal securities, corporate and foreign bonds, corporate equities and mutual fund shares. The resulting average MPC I obtain is 0.278.

Figure 5

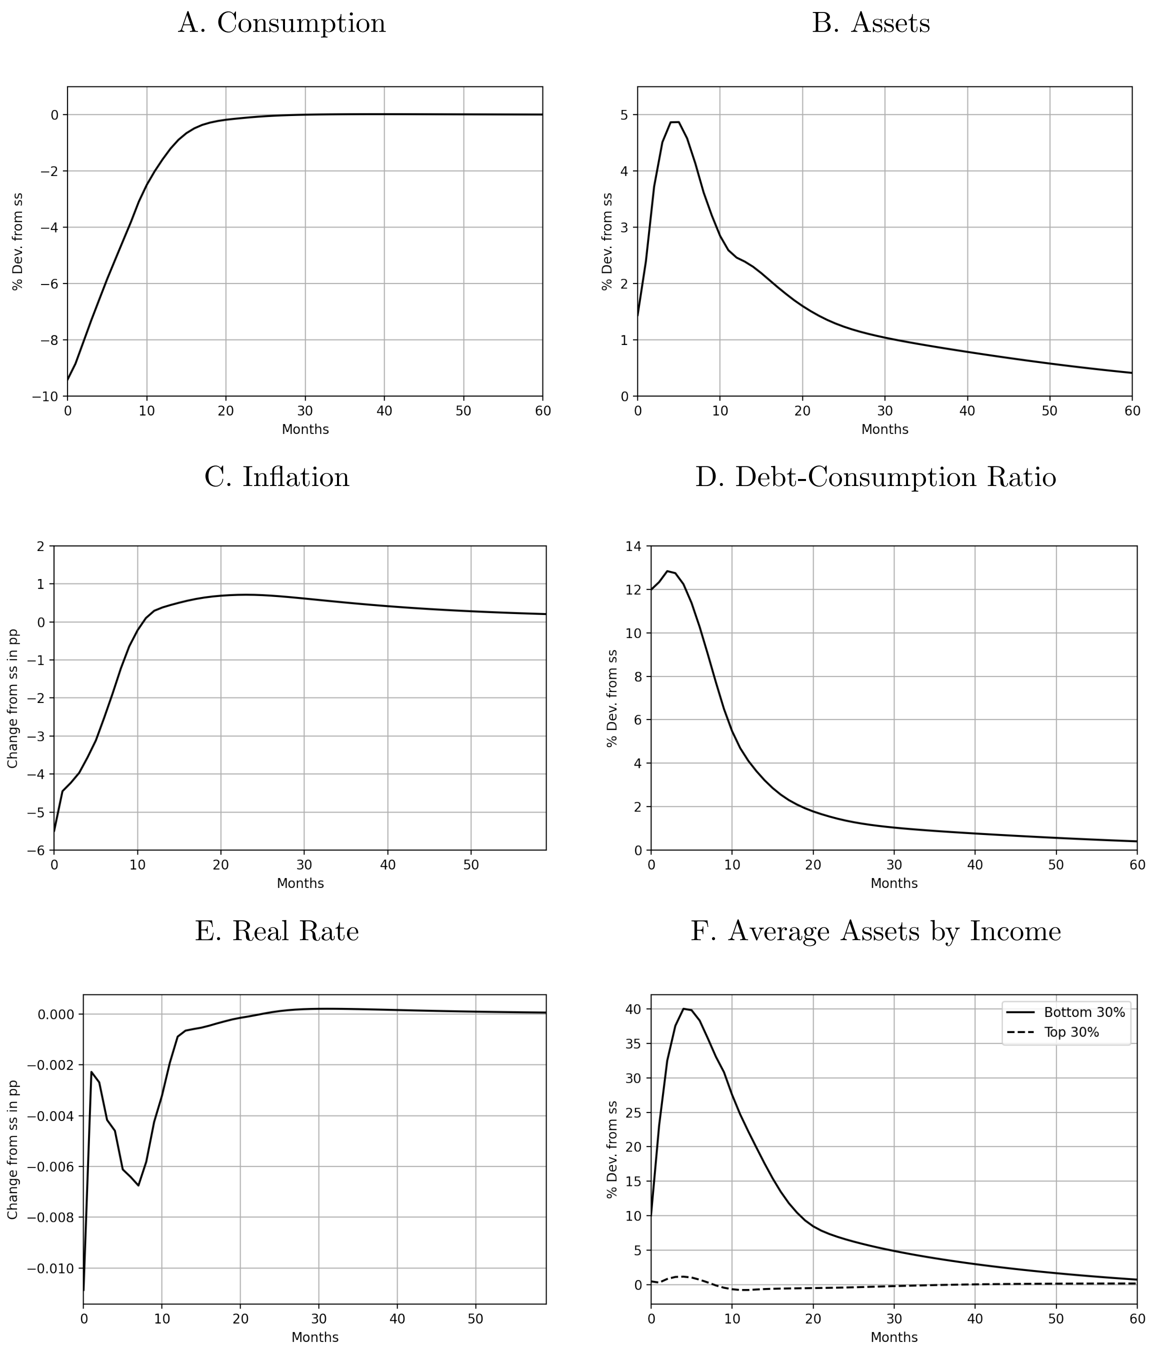

Notes: Impulse responses to the combination of Covid and fiscal shocks. Details on shocks and their calibration can be found in section 3. Panel D shows debt-to-consumption ratio, instead of debt-to-output ratio since, given the economy considered (without government expenditure, investment and exports), consumption is the most important output component. Y-axis: either percentage deviations from steady state values (panel A, B, D and F) or change from steady state expressed in percentage points (panel C and E). X-axis: months after shocks hit.

ment is initially low and, as a consequence, so are unemployment benefits, so that government debt increases slightly on impact. Debt then increases as more households are affected by the quarantine shock, driving the interest rate pick-up.

The reduction in the real rate especially reduces the increase in assets by richer segments of the population, for whom precautionary motives are less of a concern given the distribution of the quarantine shock (panel B in figure 3). Hence, as shown also in panel F, the first phase of the pandemic is characterized by a relative increase in assets by poorer households. Given their large MPCs, these households drawdown their assets as soon as their desire to consume recovers (i.e., as soon as the marginal utility shock decreases its magnitude) and this, together with a lower desire to accumulate assets by higher-income households due to lower rate, is what explains the rapid decrease in aggregate assets in this phase of the pandemic.

The second phase of the pandemic, starting after approximately one year, is characterized by a gradual decrease in aggregate savings. The reason behind this slow decrease is that savings are mostly left in the hands of higher income households since, on the one hand, poorer households have already drawn down much of their savings due to their high MPCs and, on the other hand, the recovery in the real rate incentivizes the increase in assets by richer agents. Higher income households only slowly use their assets over time because of their low MPCs, explaining the gradual decrease in private savings.

Panel C shows the inflation counterpart of these movements in aggregate savings. The first phase, lasting approximately one year, is characterized by a drop in inflation on impact, which recovers to pre-pandemic values in less than a year (this is in line with the empirical evidence, as shown in figure 2). Initial deflation is due to the reduced desire to consume by households, due both to the exogenous marginal utility shock and to the increase in precautionary motives. The magnitude of the decrease in inflation is very large, and more than the one observed in the data, suggesting that forces outside of the model assured lower deflationary pressures[10].

The second phase is characterized by a pick-up in inflation. During this period, inflation does increase relative to its pre-pandemic steady state, but the magnitude of this increase, less than one percentage point, is contained relative to the one observed in the data (shown in figure 2). The explanation for these movements lies in the dynamics of private savings. In the first phase, the rapid decrease in assets contributes to the fast recovery of inflation. However, the drawn down is not large nor fast enough to lead to an increase in consumption that leads to high inflationary pressures. In

[10] Bayer et al. (2020) predict lower deflation in a richer model with investment. Alternative forces outside of my model and the one in Bayer et al. (2020) that could explain the lower than expected reduction in actual inflation are international influences as argued by Forbes (2019), and the presence of nonlinear Phillip curves as in Gagnon and Collins (2019).

the second phase inflation is above its pre-pandemic steady state for long, approximately four years, but the magnitude of the increase is mild, since these pressures are mostly the consequence of the gradual reduction in assets by higher-income households.

Overall, this analysis shows that excess savings will be a source of persistent, albeit mild, inflationary pressures in the medium run.

4.3 Alternative Lengths of Marginal Utility Shock

This section and the next evaluate the extent to which previous results depend on two important assumptions, both of which are difficult to test: the length of the marginal utility shock and the reaction of the fiscal authority to the increase in debt.

With respect to the former assumption, figure 6 shows the alternative movements of inflation (panel A), aggregate assets (panel B), and the real rate (panel C) under two different scenarios. In the first one (dashed-black line), the marginal utility shock is assumed to decay with a higher persistence parameter (0.83) relative to the baseline scenario (0.80). This represents a scenario in which the reduction in the desire to consume, due to the fear of contagion, is reduced for longer. When this is the case, households want to increase their savings more. This has an important consequence on the real rate: relative to the baseline scenario, due to a more persistent excess demand for bonds, the real rate is depressed for longer and recovers its pre-pandemic steady state value in approximately two and a half years (panel C). Accordingly, while savings increase by more in the first phase (lower-income households accumulate larger amounts expecting a prolonged drop in utility from consumption), during the second phase higher-income households have less of an incentive to save relative to the baseline scenario. This leads to the faster drop in assets during the second phase, shown in panel B and taking place after 15 months from the beginning of the shocks.

The dashed-black line in panel A shows that inflation picks up more slowly than in the benchmark case. The reason for this is that the first phase, during which assets are high, is longer under this scenario. This reflects the more persistent reduction in the desire to consume. As a consequence to this, the pick-up of consumption is slower, and only happens once unemployment is closer to its pre-pandemic steady state value (i.e., when the quarantine shock is on its decaying phase). Overall, this implies lower inflationary pressures.

The third scenario considered is displayed by the dotted-black lines, where the persistence of the marginal utility shock is assumed to be lower than in the benchmark case, and equal to 0.77. In this scenario households are expected to recover their utility out of consumption faster than in

Figure 6

A. Inflation

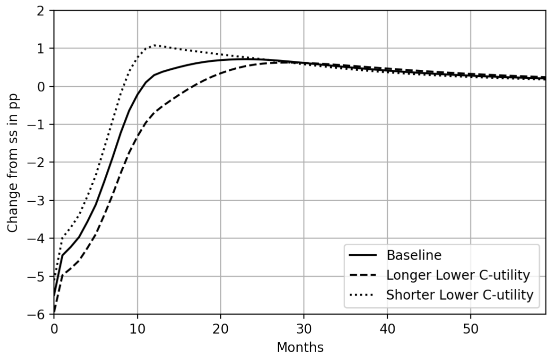

B. Assets

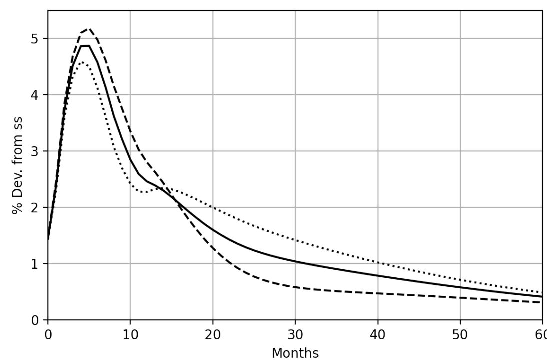

C. Real Rate

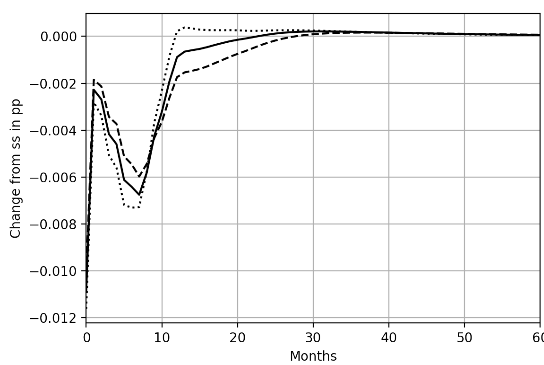

Notes: Responses under different assumptions about the length of the marginal utility shock. The solid line represents the baseline scenario, where the marginal utility shock is assumed to decay following an AR(1) process with persistence parameter equal to 0.8. The dashed line represents the responses when assuming a persistence parameter of 0.83. The dotted line represents the responses when assuming a persistence parameter of 0.77. Y-axis: either percentage deviations from steady state values (panel B) or change from steady state expressed in percentage points (panel A and C). X-axis: months after shocks hit.

the baseline scenario. Panels B and C show that this case is the mirror image of the previously considered one, represented by the dashed line. In particular, households want to hold assets for a shorter amount of time, so that the real rate reacts by increasing more than in the benchmark case: given the increase in debt due to the fiscal package, the government needs to incentivize households to hold bonds despite their lower desire to do so, which leads to a higher rate. As a consequence to this, while assets increase less in the first phase since the desire to save is lower, higher-income households quickly increase their assets during the second phase. Accordingly, households end up having more resources during the second phase of the pandemic, as shown in panel B.

Panel A shows that inflation recovers faster in this scenario, and it reaches higher levels than in the benchmark case (up to one percentage points). These movements are the consequence of the smaller accumulation of assets during the first phase of the pandemic, so that consumption picks up faster and while employment is still very much depressed due to the quarantine shock.

Overall, these simulations show that excess savings would lead to higher inflationary pressures when assuming shorter periods in which the desire to consume is shrunk. Yet, relative to the benchmark case, in which the marginal utility shock completely disappears after approximately two and a half years, the more empirically reasonable deviation seems to be one in which the persistence of this shock is longer, rather than shorter. This implies that the main conclusion from the previous section, where excess savings are expected to be a source of mild and prolonged inflation, is not much altered by the investigations of this section.

4.4 Alternative Fiscal Consolidations

As a consequence of the fiscal stimulus implemented to counteract the contractionary effects of the pandemic, the government faces an increase of its debt level after the Covid-19 breakthrough (panel D in figure 5). Assuming the government plans to have its debt eventually return to pre-pandemic values[11], it can choose how long it will take to achieve this goal. In the benchmark case, debt returns to its pre-pandemic level in approximately 12 years. By modifying the parameter ϕ_T, which controls the speed with which the government moves tax revenues after debt deviates from its steady state, it is possible to implement alternative scenarios with different lengths of the fiscal consolidation.

Figure 7 shows the IRFs under alternative assumptions about the paramter ϕ_T. The dashed-black line represents the scenario in which ϕ_T is smaller than in the baseline case. Since this corresponds to the case in which the government lets the debt level be higher than its steady state value for longer, panel B shows aggregate assets that are higher than in the benchmark case for an extended period of time. Importantly, this scenario implies a quicker, and much larger increase in inflation, as shown in panel A. The reason for this result lies in the fact that, when the government leaves debt high for long, taxes are raised very slowly over time. Accordingly, households have more resources available in the short run, which has the important consequence of allowing them to keep their consumption level constant at a lower labor supply level. This leads to the rapid and prolonged increase in inflation relative to the benchmark case, which is necessary to reduce households' consumption to levels that are compatible with their lower labor supply.

[11] This assumption is made throughout the analysis, where I am considering scenarios in which the economy returns to its pre-pandemic steady state. Alternatively, it is possible to analyze transition dynamics towards a new steady state but, in this case, assumptions about the new steady state would need to be made.

Figure 7

A. Inflation

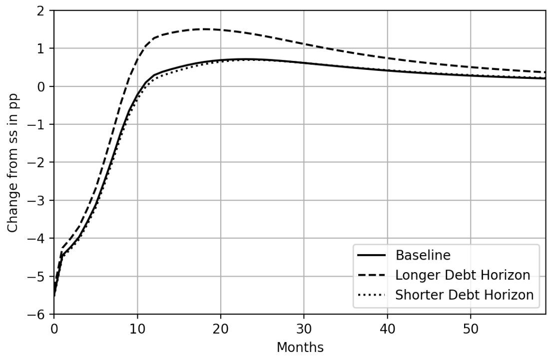

B. Assets

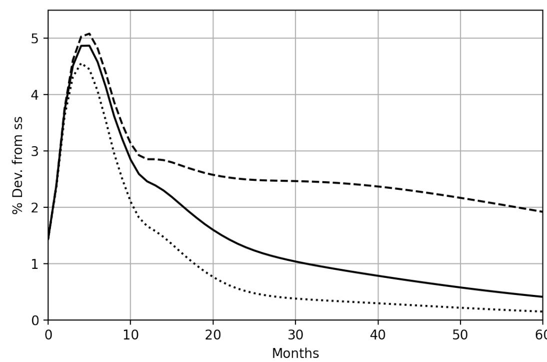

C. Real Rate

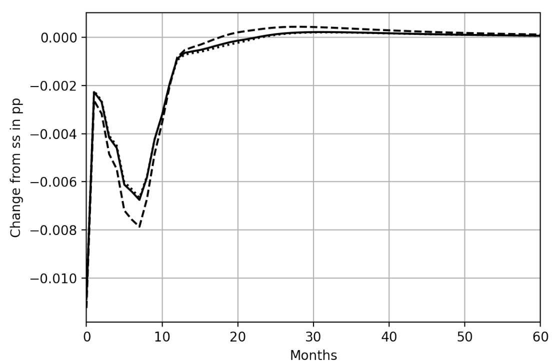

Notes: Responses under different assumptions about the fiscal response to the increase in debt. The solid line represents the baseline scenario, where the ϕ_T parameter is assumed to be equal to 0.05. The dashed line represents the responses when the ϕ_T parameter is assumed to be equal to 0.02. The dotted line represents the responses when the ϕ_T parameter is assumed to be equal to 0.08. Y-axis: either percentage deviations from steady state values (panel B) or change from steady state expressed in percentage points (panel A and C). X-axis: months after shocks hit.

The dotted-black line represents the scenario in which ϕ_T is larger than in the benchmark case, representing the case in which the government lets its debt level deviate from its steady state value for a shorter period of time (panel B). In this case, panel A shows that inflation does not deviate much from the baseline scenario: more taxes make households accumulate less savings, without altering their consumption and labor supply schedule. As a consequence, this scenario does not entail particularly interesting consequences on the relation between savings and inflation relative to the benchmark scenario.

5 Conclusion

The outbreak of the Covid pandemic has massively affected the world economy, stimulating a literature that tries to understand the consequences that it will have once the health crisis winds up. This paper adds to this literature by analyzing the consequence that one of the most important legacies of the pandemic, the increase in private savings, will have on future inflation dynamics. The characteristics of the Covid shock, together with the fiscal response implemented by the US government, have led to an increase in assets which is especially relevant for lower-income households. Through the lenses of a HANK model, I showed that the distribution of savings across households with different income levels is of crucial importance for future dynamics. Low-income households, characterized by high MPCs, quickly use their assets as soon as the economy re-opens, leading to a fast recovery of consumption and inflation after their initial drop. Afterwards, assets are left in the hands of richer households, who draw them down slowly. As a consequence, inflation recovers fast and then displays a mild and persistent increase relative to its pre-pandemic level. This is the most important result of this paper: in the baseline scenario, excess savings are expected to be a source of moderate inflationary pressures for approximately four years. Nonetheless, I showed that there exists an empirically plausible scenario in which excess savings will be a more important source of inflationary pressures. This corresponds to the case in which the government lets the debt level be higher than its pre-pandemic steady state for a prolonged period (more than 12 years). Since the US government has not yet discussed the extent to which it intends to bring its debt level back to its previous values, this scenario may be a realistic one, leading to extended and large inflationary pressures from excess savings.

References

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021a). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Rognlie, M., and Straub, L. (2021b). The Intertemporal Keynesian Cross. *Working paper.*

Bayer, C., Born, B., Luetticke, R., and Müller, G. J. (2020). The Coronavirus Stimulus Package: How large is the transfer multiplier? *Working paper.*

Bewley, T. (1977). The Permanent Income Hypothesis: A Theoretical Formulation. *Journal of Economic Theory*, 16(2):252–292.

Bilbiie, F., Eggertsson, G., Primiceri, G., and Tambalotti, A. (2021). "Excess Savings" Are Not Excessive. *Liberty Street Economics - Blog Piece - April 5, 2021.*

Blanchard, O. (2021). In defense of concerns over the $1.9 trillion relief plan. *Peterson Institute for International Economics - Blog Piece - February 18, 2021.*

Carroll, C. D., Crawley, E., Slacalek, J., and White, M. N. (2020). Modeling the Consumption Response to the CARES Act. *NBER Working paper.*

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and Team, O. I. (2020). The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data. *Working paper.*

Cox, N., Ganong, P., Noel, P., Vavra, J., Wong, A., Farrell, D., and Greig, F. (2020). Initial Impacts of the Pandemic on Consumer Behavior: Evidence from a Linked Income, Spending, and Savings Data. *Working paper*, (82).

Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2021). Inequality in Life and Death. *Working paper.*

Faria-e-Castro, M. (2021). Fiscal Policy during a Pandemic. *Journal of Economic Dynamics & Control*, 125.

Flodén, M. and Lindé, J. (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? *Review of Economic Dynamics*, 4(2):406–437.

Forbes, K. (2019). Inflation Dynamics: Dead, Dormant, Or Determined Abroad? *Brookings Papers on Economic Activity*, pages 257–319.

Fornaro, L. and Wolf, M. (2021). The Scars of Supply Shocks. *Working paper.*

Gagnon, J. and Collins, C. (2019). Low Inflation Bends the Phillips Curve. *PIIE Working Paper No. 19-6.*

Ganong, P., Greig, F., Liebeskind, M., Noel, P., Sullivan, D. M., and Vavra, J. (2021). Spending and Job Search Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data. *Working paper.*

Guerrieri, V. and Lorenzoni, G. (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap. *Quarterly Journal of Economics*, 132(3):1427–1467.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2020). Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *NBER Working paper.*

McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–58.

Mongey, S., Pilossoph, L., and Weinberg, A. (2021). Which Workers Bear the Burden of Social Distancing? *NBER Working paper*, (27085).

Nekarda, C. J. and Ramey, V. A. (2020). The Cyclical Behavior of the Price-Cost Markup. *Journal of Money, Credit and Banking*, 52(S2):319–353.

Summers, L. H. (2021). Opinion: The Biden stimulus is admirably ambitious. But it brings some big risks, too. *The Washington Post, February 4th.*

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnagie-Rochester Conference Series on Public Policy*, 39:195–214.

Appendices

A Implications of Different Asset Distributions

This section provides a simple example to elucidate the importance of the distribution of assets across households with different income levels. To achieve this goal, I implement a simple exercise. Starting from the partial equilibrium model of section 3, I shock the household block twice with a marginal utility shock. The first time (solid-black lines and black bar in figure 8), the shock is sent so that it only affects bottom-income households. Accordingly, as it can seen in panel C, which shows the percentage of the increase in assets that is accounted by households at the bottom and top 30% of the income distribution, the increase in assets displayed in panel A is mostly accounted by households at the bottom of the income distribution. The second time (dashed-grey line and grey bars), the shock is sent so that it only affects top-income households. Accordingly, panel C shows that the increase in assets is mostly due to larger assets held at the top of the income distribution.

Panel A shows that the two shocks are calibrated so that the peak increase in assets is the same in the two scenarios, despite the different cohorts directly affected by the shock. Nonetheless, they have very different consequences for future outcomes. When assets accumulate at the bottom of the income distribution, they are quickly drawn down as the shock is reabsorbed. This is because low-income households have high MPCs. As a consequence to this, panel B shows that consumption would be expected to recover quickly and increase relative to its steady state value significantly. Differently, when assets accumulate at the top of the income distribution, they are expected to be drawn down slowly (panel A), with a corresponding mild increase in consumption relative to steady state.

B Phillips Curve Derivation

The final goods producer produces output Y_t combining intermediate goods according to a CES aggregator:

$$Y_t = \left[\int_0^1 y_{i,t}^{\frac{1}{\mu}} di \right]^{\mu}$$

Figure 8

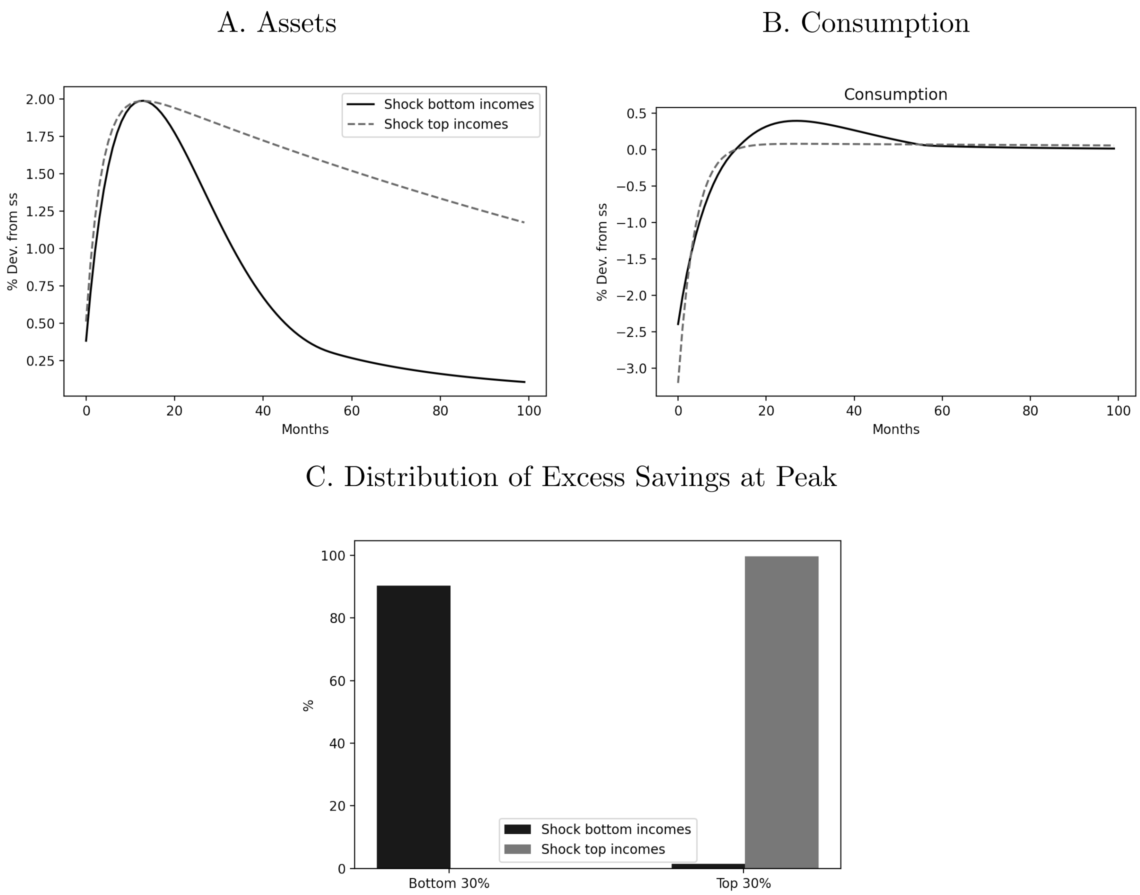

Notes: Responses to a marginal utility shock calibrated to affect bottom incomes (solid-black line and black bar) or top incomes (dashed-grey lines and grey bar). Panel C shows the percentage of the increase in assets from steady state at peak (i.e, after 18 months as displayed in panel A) that is accounted by households at the bottom (left-side bars) and top (right-side bars) 30% of the income distribution when the shock is calibrated to affect bottom-income households (black bars) and top-income households (grey bars).

where $Y_{i,t}$ are intermediate goods from each intermediate producer i. Accordingly, the maximization problem of the final goods producer is:

$$\max_{y_{i,t}} \left\{ P_t Y_t - \int_0^1 P_{i,t} y_{i,t} di \right\}$$
$$\text{s.t. } Y_t = \left[\int_0^1 y_{i,t}^{\frac{1}{\mu}} di \right]^{\mu}$$

where P_t is the price of the final good, and $P_{i,t}$ is the price of intermediate good i. The solution is given by:

$$yi,t = \left(\frac{P_{i,t}}{P_t} \right)^{\frac{\mu}{1-\mu}} Y_t$$

Intermediate goods producers take the demand from final producers as given and maximize their profits ($\Pi_{i,t}$), subject to quadratic adjustment costs:

$$\Pi_{i,t} = P_{i,t}y_{i,t} - P_t mc_{i,t}y_{i,t} - \frac{\mu}{2\kappa(\mu-1)}\left[\log(\frac{p_{i,t}}{p_{i,t-1}})\right]^2 P_t Y_t$$

Where, given the linearity of the intermediate goods producers, $mc_{i,t}$ will be equal to w_t. The associated maximization problem is:

$$\max_{P_{i,t}} P_{i,t}y_{i,t} - P_t mc_{i,t}y_{i,t} - \frac{\mu}{2\kappa(\mu-1)}\left[\log(\frac{p_{i,t}}{p_{i,t-1}})\right]^2 P_t Y_t + \frac{1}{(1+\pi_{t+1})(1+r_{t+1})}\left[P_{i,t+1}y_{i,t+1}P_{t+1}mc_{i,t+1}y_{i,t+1}\right]$$

Imposing symmetry, and plugging in the demand from final goods producers, the solution of this maximization problem boils down to the NKPC of equation 4.

C Consumption Path

Figure 9: PCE Consumption

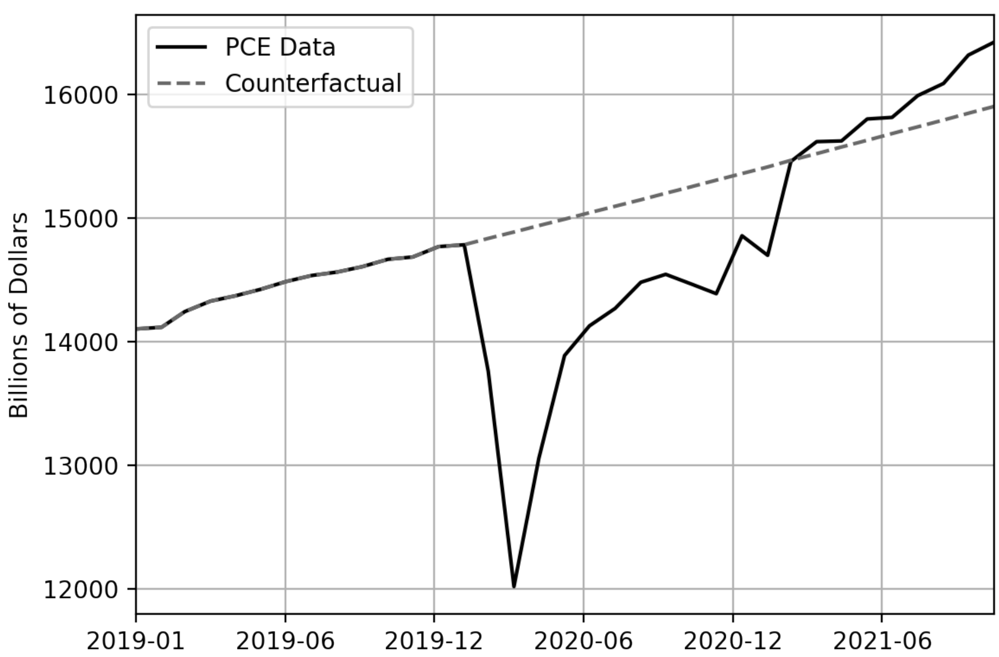

Notes: Data from FRED. The counterfactual path assumes that consumption growth had been constant constant to its 3-year average before 2020. X-axis: monthly values from January 2019 to November 2021.