

Precautionary Savings and Stabilization Policy in a Present-Biased Economy

Adriano Fernandes*

Rodolfo Rigato †

Job Market Paper

December 1, 2022

[Click here for the latest version](#)

Abstract

The business cycles literature has recently embraced heterogeneous-agent (HA) models, which generate large and dispersed marginal propensities to consume (MPCs), in line with the data. In this paper, we focus on the precautionary saving implications of these models, a less-studied feature. We first show that two calibrated versions of the model, one with standard and another with present-biased preferences, can both match MPC profiles as well as a host of other moments, but differ in their predictions for precautionary saving. We then measure the precautionary saving channel in the data by studying the response of asset accumulation to variation in unemployment insurance (UI) schedules across U.S. states as well as over time. We find small, statistically non-significant effects. Reproducing our empirical design using model-simulated data, the empirical estimates reject the standard model but are in line with the present-biased model. To illustrate the implications of this difference, we study the stabilization properties of UI in an estimated HA New Keynesian model. In standard HA models, UI affects aggregate consumption largely by reducing precautionary saving. By weakening this effect, a model with present bias predicts a fiscal multiplier of temporary UI extensions 40% smaller than a standard model. Moreover, it predicts UI to have a smaller effect in reducing aggregate consumption volatility, being therefore a less powerful automatic stabilizer as well.

*Harvard University. Email: adrihfernandes@g.harvard.edu.

†Harvard University. Email: dinisrigato@g.harvard.edu.

We are very grateful to Gabriel Chodorow-Reich, David Laibson, and Ludwig Straub for their invaluable guidance during this project. We thank Samuel Hanson and Adi Sunderam for their detailed comments, as well as Adrien Bilal, Michael Blank, Veronica De Falco, Xavier Gabaix, Benny Goldman, Omeed Maghzi, Peter Maxted, Guilherme Neves, and numerous seminar participants for their helpful insights. Any errors are our own.

1 Introduction

There is a growing consensus that heterogeneity plays an important role in macroeconomics. The business cycles literature has recently embraced heterogeneous-agent (HA) models, which generate large and dispersed marginal propensities to consume (MPCs), changing the transmission channels of fiscal and monetary policies (McKay et al. 2016, Kaplan et al. 2018, Auclert et al. 2018). In these models, however, agents engage in precautionary saving, a less-studied feature. How strong is this precautionary behavior in the data? Is it in line with the predictions of standard HA models? What are its implications for business cycles and stabilization policy? These are the questions we address in this paper.

Our first contribution is to document that in standard HA models, the precautionary savings response to changes in income risk is large and inconsistent with the data, but that the introduction of present bias bridges this gap. We do so in the following steps. First, we set up two different HA models: one with standard preferences and another with present bias. Next, we show both numerically and analytically that the present-biased model predicts smaller precautionary savings responses to changes in UI generosity, a policy that affects the income risk households face. Then, we use variation in UI schedules across U.S. states as well as over time to measure this response empirically, finding small, statistically non-significant effects. Finally, we show that this evidence can be used to discriminate across the two models: when running the same regression in the standard model, we obtain a large coefficient that falls outside the estimated confidence interval. The data, therefore, soundly reject the standard model. The present-biased model, on the other hand, delivers a coefficient well within confidence intervals, and so is not rejected by the data.

Our second contribution is to explore the implications of reduced precautionary saving for business cycles and stabilization policy. In particular, we focus on the stabilization properties of UI in an estimated Heterogeneous-Agent New Keynesian (HANK) model. We study the role of UI both as a discretionary policy tool and as an automatic stabilizer.¹ We find that in standard HANK models, UI affects consumption largely through precautionary saving. Facing a more generous UI policy, agents in these models understand they do not need to save as much to self-insure against future unemployment risk. When precautionary behavior is attenuated by present bias, UI is therefore less effective. As a discretionary policy tool, UI extensions have fiscal multipliers approximately 40% smaller when agents are present-biased. As an automatic stabilizer, the effect of UI in reducing aggregate consumption volatility is also around 40% smaller in this case if compared to a standard HANK model.

We start by introducing a canonical model of precautionary saving which features heterogeneous workers subject to uninsurable income and unemployment risk. Unemployed workers have access to UI benefits that expire after a fixed number of months and take-up is assumed to be partial: a fraction of the unemployed do not collect their benefits, in accordance with empirical evi-

¹According to McKay and Reis (2021), “fiscal stabilizers are rules in law that make fiscal revenues and outlays relative to total income change over the business cycle.”

dence.² These features ensure that we are correctly capturing the effects of UI on income risk and precautionary saving.

We calibrate two versions of our model: one with standard preferences and one with present bias. While most macroeconomic models are populated by rational, forward-looking agents who discount the future exponentially, an ever-growing body of evidence challenges this model of intertemporal choice.³ Present-biased agents strongly discount future utility relative to the present, more strongly than the additional discount between two consecutive future dates, deviating, therefore, from the exponential discounting model. In particular, we focus on present bias of the naïve type: agents are not aware of their present bias, leading them to make dynamically inconsistent choices. Naïve present-biased agents, therefore, have the desire to save to self-insure against future income risk, but they procrastinate: they do not start saving in the present because they wrongly believe their future selves will do so. Both versions of the model – exponential and present-biased – are calibrated to match the same moments.

Our first exercise is to compute the increase in precautionary savings as UI generosity is reduced. We entertain two cases: one in partial equilibrium, in which interest rates and taxes are kept constant, and another in general equilibrium, in which these objects are allowed to vary. In partial equilibrium, eliminating UI generates an increase in total savings of more than 30% with standard preferences, but of only approximately 15% with present bias. In general equilibrium, annualized interest rates drop by almost 2 p.p. in the former case, but by less than 0.5 p.p. in the latter. Therefore, present bias substantially attenuates the precautionary saving response to changes in income risk. The usefulness of the partial equilibrium exercise is that later we compare the predictions of both models to estimates of precautionary savings responses to state-level UI policies. When simulating UI reforms in the model, we assume that state-level policies do not affect national interest rates, similar to the partial equilibrium case.

To shed some light on these results, we analyze what happens in a tractable parameterization of our model. We consider a special case in which, among other simplifying assumptions, agents cannot borrow and the supply of assets converges to zero – the so-called *zero liquidity equilibrium*.⁴ Under this assumption, we can analytically characterize the sensitivity of the equilibrium real interest rate to changes in social insurance. Our formula allows us to decompose the effects of present bias into (i) the extra discounting that applies to future utility and (ii) the time inconsistency that arises from naïveté. When we evaluate our expression using standard parameter values, we find that the reduced sensitivity of interest rates – and therefore of precautionary saving – comes almost entirely from the latter.

²See Blank and Card (1991) and Auray et al. (2019).

³Ericson and Laibson (2019) provide an up-to-date survey on intertemporal preferences. Beshears et al. (2018) survey the evidence with implications for household finance. DellaVigna (2018) provides a survey of structural estimates of the parameter governing the degree of present bias in quasi-hyperbolic models (typically denoted by β). Most recently, Blow et al. (2021) pursue a nonparametric approach: using data from consumer expenditure surveys, they show that pass rates for choice rationalizability are *twenty* times higher under naïve quasi-hyperbolic discounting relative to exponential discounting. Our calibrated time preference parameters are quite likely under the densities of $\beta - \delta$ parameters they recover.

⁴See Werning (2015), Challe (2020), and McKay and Reis (2021) for other examples of zero liquidity models.

The reduced sensitivity of precautionary saving under present bias raises the question of whether this is consistent with the data. To address this empirical issue, we exploit variation in UI generosity across US states and over time to estimate how strongly precautionary savings respond to social insurance reforms. Using Survey of Income and Program Participation (SIPP) data, we expand on Engen and Gruber (2001) by leveraging a simulated instrument strategy in a canonical two-way fixed effect design to estimate the response of household liquid asset holdings to changes in state-level UI replacement rates. We find a small, statistically non-significant effect. This finding is robust across a number of wealth measures and different empirical strategies, such as two-stage least squares and stacked regression designs. In our baseline two-way fixed effect specification, with 95% confidence, a 10 p.p. reduction in UI generosity is associated with an average increase in liquid asset holdings no greater than 0.48% of annual income.

We then reproduce the exact same regression using model-simulated data. In the standard model (i.e., with exponential discounting), the same 10 p.p. reduction in replacement rates is associated with an increase in asset holdings-to-income ratio of 0.64% – a value well outside the estimated confidence interval. The data, therefore, reject this prediction of the standard HA model. The present-biased model, on the other hand, predicts an average increase of 0.21%, well within the empirical confidence intervals.

After showing that the standard model generates precautionary savings responses too large in comparison to the data, we study the implications of this finding for stabilization policy. To do so, we augment our model with wage rigidity, monetary and fiscal rules, and aggregate shocks. We estimate the additional parameters using Bayesian methods and time series data on aggregate consumption, inflation, and interest rates. We then use the estimated model to study two UI-related policies.

First, we analyze the effects of discretionary UI extensions on aggregate consumption, as in Kekre (2021). UI directly stimulates the consumption of unemployed workers, but it also affects the consumption of *employed* workers – typically, the vast majority of the economy – via precautionary saving. In the standard HA model, a 12-month increase in UI duration has a fiscal multiplier of 1.05. When precautionary behavior is attenuated by present bias, the multiplier is 0.63 – a substantially smaller value.

We then turn to the role of UI as an automatic stabilizer. Unemployment rates rise during recessions and UI helps displaced workers sustain their consumption. Therefore, UI has the effect of reducing the volatility of aggregate consumption over time. In the spirit of McKay and Reis (2016), we reduce the generosity of UI and compute the associated increase in this volatility, which we take as a measure of the efficacy of UI as an automatic stabilizer. To understand the role of precautionary saving in this exercise, we decompose the consumption response to labor demand fluctuations into two terms. First, an increase in unemployment has a direct negative effect on households' disposable income, which we call the *displacement effect*. Second, if the increase in unemployment is persistent, it leads households to save in order to self-insure against increased income risk. We call this the *anticipation effect*. In a standard HANK model, cutting UI generosity

by half increases the standard deviation of aggregate consumption by 17.2%, mostly for increasing anticipation effects. In a present-biased economy, this effect is dampened, and the same policy leads to a 10.1% increase in the volatility of consumption. UI is therefore a less effective automatic stabilizer in this case.

Related literature. Our work relates to several strands of recent research in macroeconomics. First, a growing literature that studies the effects of household heterogeneity for the transmission of macroeconomic shocks and stabilization policy, recent examples including Kaplan and Violante (2014), Werning (2015), McKay et al. (2016), Auclert et al. (2018), and Auclert et al. (2020). Closer to our setting are papers that focus explicitly on precautionary saving and countercyclical income risk. McKay (2017) and Schaab (2020) study the role of increased precautionary saving in exacerbating the aggregate consumption decline during recessions. In a related paper, Ravn and Sterk (2021) study how this feedback effect arises endogenously in a model with search and matching. Cho (2020) uses an estimated HANK model to assess the relative importance of MPC heterogeneity and precautionary saving in explaining the dynamics of aggregate consumption. These papers, however, employ models with standard preferences. We extend their analyses by first showing that standard models generate predictions about precautionary savings that contradict the data, looking at aggregate predictions of these models only after bridging this gap using present-biased preferences. Challe (2020), Acharya and Dogra (2020), and Bilbiie and Ragot (2021) also feature countercyclical income or unemployment risk in standard models, but focus on monetary policy, while we look at UI extensions.

Two papers in the aforementioned literature deserve special attention for focusing on UI or other automatic stabilizers. The first is McKay and Reis (2016), who study the role of automatic stabilizers in reducing the volatility of aggregate output and consumption. They focus, however, on a broader set of stabilizers, including, for example, consumption and progressive income taxes. By focusing on UI exclusively, we are able to incorporate realistic features that are absent in their analysis, such as UI expiration and partial take-up. Moreover, our contribution is not only computing the efficacy of UI as a stabilization tool, but also understanding how standard macroeconomic models overstate efficacy by featuring too strong precautionary behavior. In related work, McKay and Reis (2021) characterize the optimal design of automatic stabilizers in a similar setting. We abstain from normative analyses, focusing only on the positive effects of UI. The second paper is Kekre (2021), who studies UI extensions in a similar model with exponential preferences, finding fiscal multipliers close to 1. We find a similar value in our specification with standard preferences. However, the present-biased specification that matches the sensitivity of precautionary savings in the data delivers a fiscal multiplier around 40% lower, suggesting that the strength of precautionary saving is an important force behind his result.

Another related strand of literature studies how behavioral frictions interact with stabilization policy, with recent examples including Gabaix (2020), Erceg et al. (2021), and Vimercati et al. (2021). Closer to our setting is Laibson et al. (2021), who explore the implications of present bias for the

transmission of monetary shocks. In particular, they focus on how procrastination generated by present bias can slow the mortgage refinancing channel of monetary policy. We differ from them by focusing on a different implication of present bias, namely reduced precautionary behavior in liquid asset accumulation. We also focus on fiscal policy. Another related work is by Maxted (2020), who studies the implications of present bias for several outcomes of heterogeneous-agent models, such as wealth distributions. He does not, however, focus on business cycles.

Last, a body of empirical work studies the effects of social insurance – including UI – and income risk on consumption-savings choices. Closest to us is Engen and Gruber (2001), who also estimate the response in asset holdings associated with variation in UI schedules across U.S. states and over time. We extend their analysis both by incorporating more data and by employing more modern econometric methodologies, such as stacked regressions. We also build upon Ganong and Noel (2019), who estimate the spending of workers along unemployment spells. They show that the consumption drop upon UI expiration – a perfectly predictable event – is inconsistent with the consumption smoothing behavior implied by standard preferences, but that present bias can match this empirical moment. We use their estimates of consumption during unemployment as a calibration target for our model.

Outline. The paper proceeds as follows. In section 2, we introduce a standard model of precautionary savings, both with standard and present-biased preferences. In section 3, we study the predictions of both models for precautionary savings and contrast it with empirical evidence. In section 4, we introduce the model ingredients necessary for business cycle analysis and estimate the additional parameters. Section 5 shows the implications of present-biased preferences for stabilization policy. Section 6 then presents our concluding remarks.

2 A Standard Model of Precautionary Savings

In this section, we introduce a standard model of consumption and savings decisions in the presence of uninsurable income risk, which motivates agents to accumulate precautionary savings. This risk is, however, partly mitigated by government transfers. Our goal is to understand how precautionary savings respond to changes in income risk and we do so by varying the generosity of such transfers. We later show that present bias attenuates this response, so we allow for present-biased preferences from the start. For now, we focus only on what is strictly necessary for understanding precautionary saving, leaving features such as monetary policy and nominal rigidity for section 4. Until then, all variables are expressed in real terms.

2.1 Setup

Time is discrete, indexed by t , and each period corresponds to a month. There is a unit mass of households indexed by i .

Preferences. Households have preferences over streams of consumption and labor supply given by

$$\mathcal{V}_{it} = u(c_{it}) - v(n_{it}) + \beta_i \mathbb{E}_t \sum_{k=1}^{\infty} \delta^k (u(c_{i,t+k}) - v(n_{i,t+k})). \quad (1)$$

Above, c_{it} and n_{it} denote, respectively, consumption and hours worked in period t . The functions $u(\cdot)$ and $v(\cdot)$ take the following forms:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \quad v(n) = \psi \frac{n^{1+\varphi}}{1+\varphi}.$$

Agents discount future utility exponentially according to the discount factor $\delta \in (0, 1)$. There is, however, a *present bias* parameter $\beta_i \in (0, 1]$ that further discounts all future utils relative to present ones. This functional form is usually called *quasi-hyperbolic* discounting (Laibson, 1997). In the particular case in which $\beta_i = 1$ only the exponential discounting remains, as in most macroeconomic models. We refer to this as the *exponential*, or *standard*, preferences. We follow Ganong and Noel (2019) in allowing heterogeneous degrees of present bias: $\beta_i \in \{\beta^L, \beta^H\}$, for $\beta^L < \beta^H$. We denote by ω_L and $\omega_H = 1 - \omega_L$ the shares of households with low and high betas, respectively.

To fully specify time preferences, one additional assumption is necessary. Preferences of the form (1) are *time-inconsistent* whenever $\beta_i < 1$. This means that absent commitment devices, the agent's future selves will deviate from the consumption plans laid out in the present. We assume that agents are *naïve* (O'Donoghue and Rabin, 1999), falsely believing their future selves will not deviate from currently established plans.⁵ Intuitively, naïve present-biased agents have the desire to save to self-insure against future income risk, but they procrastinate in doing so.

Earnings and employment. Besides present bias, agents are heterogeneous with respect to employment status, labor productivity z_{it} , and asset holdings a_{it} . Each period, households face the following budget constraint:

$$a_{it} + c_{it} = (1 + r_t + \chi \cdot 1(a_{it} < 0)) a_{i,t-1} + y_{it} \quad (2)$$

$$a_{it} \geq \underline{a},$$

where y_{it} is the agent's disposable income, \underline{a} is an exogenous borrowing limit, and r_t is the real interest rate. It incorporates a borrowing spread χ whenever $a_{i,t-1} < 0$, as in Kaplan et al. (2018) and Laibson et al. (2021).

Agents in the model can be either employed or unemployed. At the beginning of period t , a fraction s_t of employed workers is separated and becomes unemployed, and a fraction f_t of unem-

⁵Alternatively, one could assume that agents are *sophisticates*: they understand their future selves will deviate from their current plans, and therefore strategically respond to their own time-inconsistency. Harris and Laibson (2001) derive a generalized Euler equation allowing for sophistication. Their key insight is that the consumption behavior of sophisticates is an intermediate case between the naïve and exponential cases: the closer (further) agents are from the borrowing constraint, the more they behave as in the naïve (exponential) case.

employed agents finds employment. We denote by U_t and $E_t = 1 - U_t$ the shares of unemployed and employed agents, respectively. Our timing convention is that a newly unemployed worker cannot immediately find a new job and stays unemployed for at least one period. Employed agents earn labor income

$$y_{it}^{labor} = (1 - \tau_t)w_t n_{it} e^{z_{it}},$$

where w_t is the real wage, n_t is the number of hours worked, and τ_t is a linear tax rate. For now, we assume that all employed workers inelastically supply one unit of labor, i.e., $n_{it} = 1$. In section 4, we relax this assumption and introduce nominal wage rigidity.

Unemployed workers receive UI benefits. In practice, however, not all the unemployed are eligible for it, and many who are do not apply for it (Auray et al. 2019, Chodorow-Reich and Coglianesi 2019). To account for this fact, we follow Kekre (2021) in assuming that only a randomly selected fraction ζ of newly unemployed agents receive UI. Benefits are given by

$$UI_{it} = (1 - \tau_t) \min\{bwne^{z_{it}}, \bar{b}\}.$$

They consist of a constant replacement rate b up to a maximum level \bar{b} and expire after N^{ui} periods. UI benefits are taxed according to the same tax rate τ_t that applies to labor income, approximating U.S. law. We omit time subscripts when referring to steady-state values of variables. Benefits, therefore, depend on steady-state real wages and labor supply, rather than current values. Moreover, we assume that the idiosyncratic labor productivity evolves according to

$$z_{it} = \begin{cases} \mu^z + \rho^z z_{it} + \sigma^z \varepsilon_{it} & \text{if employed} \\ z_{i,t-1} & \text{if unemployed} \end{cases}$$

where $\varepsilon_{it} \sim NID(0,1)$. These assumptions ensure that UI benefits are constant during an unemployment spell, while still allowing us to summarize labor productivity with a single state variable z_{it} . The parameter μ^z is set to normalize $\int \exp(z_{it}) di = 1$.

Unemployed workers also have other sources of income. First, they have access to private insurance of the form

$$y_{it}^{priv} = y^{priv} wne^{z_{it}},$$

which can be interpreted as home production or other sources of income, such as spousal insurance (Bardóczy 2022). This additional income is necessary for matching the decrease in consumption when an agent becomes unemployed, as statutory UI replacement rates alone would generate a drop that is too large compared to the data. For the same reason, we follow McKay and Reis (2016) in assuming that households that do not receive UI, either because it has expired or because they are not eligible, receive safety-net government transfers T^{sn} . These transfers can be interpreted as means-tested programs such as SNAP and are calibrated to match the same consumption drop upon UI expiration as observed in the data.

Disposable income can be summarized as

$$y_{it} = \begin{cases} y_{it}^{labor} & \text{if employed} \\ y_{it}^{priv} + UI_{it} & \text{if unemployed receiving UI} \\ y_{it}^{priv} + T^{sn} & \text{if unemployed not receiving UI} \end{cases}$$

Labor income is taxed linearly at a rate τ_t , as are UI benefits, approximating U.S. law.

Government. The government collects taxes that are used to finance interest payments on public debt B_t , a constant amount of purchases G , as well as social security spending S_t , given by

$$S_t = \int UI_{it} 1_{it}^{ui} di + T^{sn} \int 1_{it}^{sn} di,$$

where 1^{ui} and 1^{sn} are indicator functions for the receipt of UI and safety net transfers, respectively. Total government revenue is given by

$$R_t = \tau_t \left(\int y_{it}^{labor} 1_{it}^e di + \int UI_{it} 1_{it}^{ui} di \right), \quad (3)$$

The government budget constraint is, therefore,

$$B_t = (1 + r_t)B_{t-1} + G + S_t - R_t. \quad (4)$$

Firms. There is a perfectly competitive representative firm that produces final output Y_t^{firms} according to the linear production function

$$Y_t^{firms} = Z \cdot N_t, \quad (5)$$

where Z is a fixed total factor productivity (TFP) term and N_t is total labor input. Given that average labor productivity is normalized to equal 1, it is possible to express N_t as

$$N_t = E_t \cdot n_t. \quad (6)$$

Market clearing The model has two market clearing conditions. The goods market clearing condition is

$$C_t + G - \chi \int a_{i,t-1} 1_{it}^{debt} di = Y_t, \quad (7)$$

where 1_{it}^{debt} is an indicator function for agents with $a_{i,t-1} < 0$. Total output Y_t is given by

$$Y_t = Y_t^{firms} + Y_t^{priv},$$

where Y_t^{priv} is the total private insurance income:

$$Y_t^{priv} = \int y_{it}^{priv} 1_{it}^u di.$$

Last, we assume that the government is the sole issuer of bonds in the economy, which gives us the following asset market clearing condition

$$\int a_{it} di = B_t.$$

2.2 Calibration

Now we turn to the calibration of our model. We set up two versions of it, one with exponential and another with present-biased preferences. We calibrate several parameters externally using standard values, and then calibrate internally the remaining parameters to match two sets of empirical moments. The first is *spending during unemployment* for UI recipients. The left panel of figure 1 shows this set of empirical moments, reproduced from Ganong and Noel (2019) and normalizing pre-unemployment consumption to 1. The figure also shows the values generated by each of our models.

Ganong and Noel (2019) compute these moments using the universe of Chase consumer checking and credit card accounts. The authors show that these moments are inconsistent with the exponential model, as the strong consumption smoothing behavior generated by standard preferences cannot account for the sizable drop in spending upon the perfectly predictable expiration of UI benefits, which happens between months 7 and 8. They also show that present-biased preferences are consistent with this feature of the data. The usefulness of this set of moments in our context is therefore threefold. First, it captures the extent to which unemployment represents a risk for consumers. Second, it gives us a clear way of calibrating our preference parameters. Third, the consumption drop upon unemployment is also related to the response of aggregate consumption to employment fluctuations.

The second set of moments we use is the *intertemporal marginal propensities to consume* (iMPCs), defined as the dynamic responses of consumption to unanticipated shocks to disposable income. As shown by Auclert et al. (2018), the iMPCs constitute a set of sufficient statistics for the consumption responses to such shocks, and will therefore play an important role in understanding stabilization policy later on. The right panel of figure 1 shows the same iMPC estimates used by Auclert et al. (2018), obtained from the consumption increase of lottery winners in Norway, as reported in Fagereng et al. (2021). The figure shows that both exponential and present-biased models fare well in matching these moments.

The top part of table 1 shows the parameters that are calibrated externally, and which are common to both specifications. We fix the steady-state real interest rate to $r = 4\%$ annually and let the discount factors adjust to match this value. The coefficient of risk aversion is set to $\sigma = 2$, as in Carroll (1997). The AR(1) parameters of the idiosyncratic labor productivity process are taken

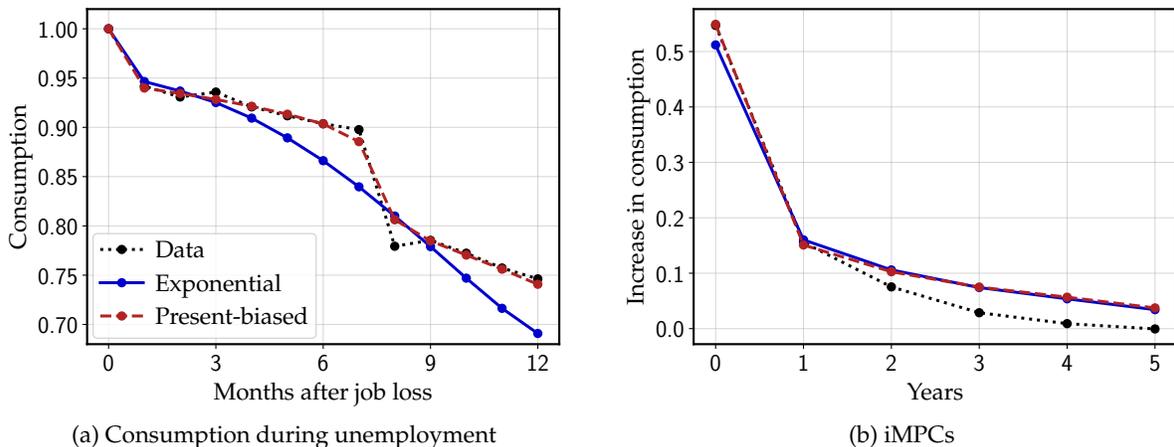


Figure 1: Calibrated consumption during unemployment and iMPCs.

from Floden and Lindé (2001), adjusted for our monthly setting.⁶ The public debt B is set at 26% of aggregate income to match the aggregate liquid asset holdings of households, as estimated by Kaplan et al. (2018).

Regarding the UI system, we calibrate $b = 0.5$ and $\bar{b} = 0.48$. As we normalize $n = w = 1$ in steady state, the latter value corresponds to 48% of the average pre-tax labor income. The UI parameter values are set to match the distribution of replacement rates in the SIPP data, explained in more detail in section 3.2, as shown in figure 2. The black dotted line shows replacement rates by income decile in the data, while the red dashed line shows it for the income distribution in our model, corresponding to the seven points of the discretized version of the AR(1) process for idiosyncratic skill. The UI duration is $N^{ui} = 7$, as in Ganong and Noel (2019).

We choose the take-up rate $\zeta = 0.77$ to match the estimated UI take-up among eligible workers from Auray et al. (2019). In the data, however, not all workers are eligible to receive UI. The reason why we choose this value is that, in section 3.2, in order to better capture the effects of UI on asset accumulation, some sample restriction assumptions are required, leaving us almost exclusively with a sample of eligible households. This value, therefore, allows us to better compare our model with the empirical evidence. The separation and job finding probabilities are set to $s = 1.2\%$ and $f = 19\%$ to jointly match a steady-state unemployment rate of 5.97% and average duration of unemployment spells of 17 weeks, the average values for the U.S. economy in the 1960-2019 period. The wage markup is set at $\mu^w = 1.1$. The TFP Z is used to normalize $w = 1$. The bottom part of table 1 shows internally calibrated parameters, which differ in the exponential and present-biased specifications. In the latter, the degree of present bias β^L and β^H , the mass of each type ω_L and ω_H , the private insurance parameter y^{priv} and transfers T^{sn} , as well as the borrowing constraint \underline{a} and spread χ are all jointly determined to match iMPCs and consumption paths. The value $\chi = 0.0098$ in our calibration delivers an annualized credit spread of around 12%, close to the 10.3% value

⁶Floden and Lindé (2001) use annual data. We adjust the persistence parameter by setting $\rho = \rho_{annual}^{\frac{1}{12}}$ and choose the variance of the shock to generate the same stationary variance.

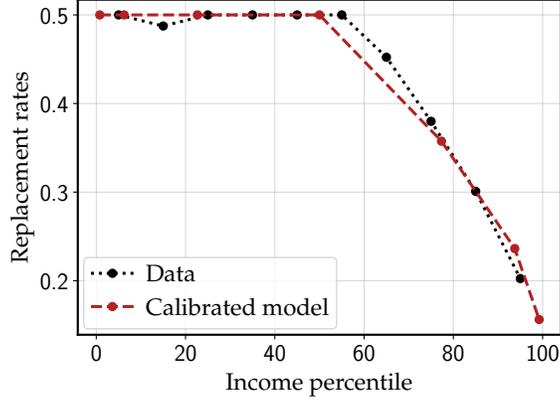


Figure 2: Replacement rates along the income distribution.

computed by Laibson et al. (2021).

The present-biased model generates an average consumption of the unemployed equal to 80.4% of the average consumption of the employed, in line with empirical estimates (McKay and Reis, 2021). As the drop in consumption upon unemployment is an important determinant of precautionary saving, as emphasized by Schaab (2020), we calibrate the exponential model to generate the same 80.4% ratio. It is important to notice that the value $\underline{a} = 0$ for the exponential model is not a restriction – it is rather a value necessary for matching iMPCs and the relative consumption of employed and unemployed workers. Even with $\underline{a} = 0$, the exponential model is not able to generate such a large consumption difference between unemployed and employed workers, delivering a ratio of 85.7%. In section 3, we relax the borrowing constraint of the exponential model as a robustness check. The discount factor δ is chosen to match the calibrated steady-state interest rate r . Note that the present-biased model requires weaker discounting to match the same interest rate. Last, the labor tax τ adjusts to keep the government budget balanced in the steady state.

The present-biased model generates a distribution of liquid asset holdings similar to others found in the literature. For example, Laibson et al. (2021) use data from the Survey of Consumer Finances to compute a share of 60% of households with credit card debt and use this value as a target for the mass of households with negative liquid wealth in their model, which generates a value of 47%. They also report that average credit card debt corresponds to 9% of annual income, which their model matches exactly. Our model generates 62% of agents with debt and an average debt of 9.8% of annual income, even though we do not target these moments in our calibration.

3 Precautionary Savings and Income Risk

We now turn to the following question: if there is an increase in income risk, by how much do precautionary savings increase? We start by characterizing this sensitivity of savings to changes in social insurance policies in our model, both numerically and analytically for a tractable special case. Then, we provide estimates of similar objects using SIPP data and show that the exponential

Parameter	Meaning	Value		Source/Target
Externally calibrated parameters				
σ	CRRRA	2		Carroll (1997)
ρ^z	AR(1) persistence	0.993		Floden and Lindé (2001)
$\sigma^z / \sqrt{1 - (\rho^z)^2}$	AR(1) std.	0.51		
s	Separation rate	0.012		$U = 5.97\%$
f	Job finding rate	0.19		Avg. weeks unemp. 17
b	Replacement rate	0.50		US UI system
\bar{b}	UI cap	0.48		
N^{ui}	Duration	7		
ξ	Take-up	0.77		Auray et al. (2019)
r	Interest rate	4% p.a.		Standard
B/Y	Liq. assets	0.26		Kaplan, Moll and Violante (2018)
Internally calibrated parameters				
		Exponential	Present-biased	
δ	Discounting	0.988	0.991	Target interest rate
$\{\beta^L, \beta^H\}$	Present bias	1	{0.5, 0.9}	
ω_H	Share of high β	–	0.74	
y^{priv}	Non-UI income	0.320	0.320	iMPC and consumption
T^{sn}	Safety net transfers	0.087	0.087	during unemployment
\underline{a}	Borrowing constraint	0	–1.48	
χ	Borrowing cost	–	0.0098	
τ^ℓ	Labor income tax	0.217	0.217	Government budget balance

Table 1: Calibrated parameters.

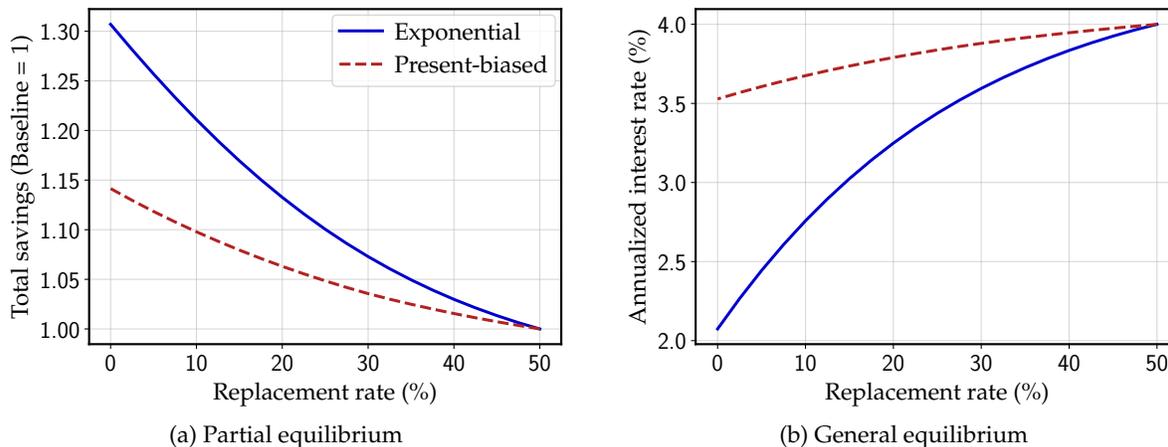


Figure 3: The response of precautionary savings to changes in UI.

model generates too strong responses, while the present-biased version is much closer to the data. Importantly, the model is not calibrated using this information, so this is a key untargeted moment that our model matches.

3.1 Theory

We start by computing the response of precautionary savings to changes in UI in our model. Two parameters govern transfers to the unemployed: the replacement rate b and the maximum benefit \bar{b} . The experiment we perform is to scale down both by the same factor until reaching zero.

Figure 3 shows the results. The left panel shows what happens to the total amount of savings in the economy in a partial equilibrium setting in which interest rates and taxes are kept constant at their baseline levels. The right panel shows what happens if these variables adjust to restore equilibrium. The usefulness of the partial equilibrium exercise is that later we compare the predictions of both models to estimates of precautionary savings responses to changes in UI policies in the data. When simulating UI reforms in the model, we assume that state-level policies do not affect national interest rates.

The most salient feature of this exercise is that the exponential calibration features much more sensitive precautionary saving than the present-biased one. In partial equilibrium, savings increase by more than 30% in the standard model, while only by less than 15% when agents are present-biased. In general equilibrium, annualized interest rates decrease from 4% to approximately 2% in the exponential economy, while they fall by less than 0.5 p.p. in the model with present bias. To understand the mechanism behind this effect, we now turn to a tractable calibration that allows us to derive closed-form expressions for this result in the latter case.

A tractable parameterization. Heterogeneous-agent models are often not analytically tractable because it is difficult to keep track of the distribution of agents over the associated state space. To overcome this issue, we consider here a *zero liquidity* parameterization (Werning 2015): we assume

that agents cannot borrow, i.e., $\underline{a} = 0$, and consider the case in which the supply of liquid assets shrinks to zero: $B \rightarrow 0$. For further simplicity, we assume that there is no heterogeneity in labor productivity ($e^{z_{it}} = 1$ for all agents), no UI expiration, full take-up, and no private insurance ($y^{priv} = 0$), which makes UI is the sole source of income for the unemployed. Moreover, we assume no heterogeneity in present bias $\beta^H = \beta^L \equiv \beta$. The tractability here comes from the fact that the Euler equation of employed agents still holds with equality as $B \rightarrow 0$, which allows us to characterize analytically the interest rate in this economy.

Given the assumptions above, all employed agents consume the same amount, given by their income $y^e = 1 - \tau_t$, while unemployed agents consume $y^u = (1 - \tau_t)b$. Due to naïveté, present-biased agents believe their future selves will consume an amount $c_e^{exp} < y_e$ in the future if they remain employed – they erroneously believe their future selves will start accumulating precautionary savings. More details on this model, as well as the proof of the proposition below, can be found in Appendix A. We have the following result:

Proposition 1. *In the zero liquidity parameterization, the response of the interest rate to changes in social insurance b can be decomposed as*

$$\frac{dr}{db} = \underbrace{\beta}_{\text{discounting}} \delta(1+r)^2 \left[s\sigma b^{-\sigma-1} + \underbrace{(1-s)\sigma \left(\frac{c_e^{exp}}{y_e} \right)^{-\sigma-1} \frac{d}{db} \left(\frac{c_e^{exp}}{y_e} \right)}_{\text{time inconsistency}} \right]. \quad (8)$$

Above, we highlight the terms that are exclusive to the present-biased calibration. First, in this case, there is an extra discount factor β that reduces the sensitivity of r . Second, changes in UI policies may lead to changes in expected consumption in the future, what we dub the *time inconsistency* effect. Both these effects are absent with exponential preferences.

We can plug in numbers into the above expression to understand what causes the reduced sensitivity of precautionary saving under present bias. We set $\beta = 0.8$, an intermediate case between the high- and low-beta agents in the previous calibration, and set $b = 0.804$ so as to generate the average consumption gap between employed and unemployed agents as before. We recalibrate δ in each case to generate the same 4% annualized interest rate, but this is not important as δ is always very close to 1, and so it plays a negligible role in equation (8).

For the exponential calibration, we have

$$\frac{dr}{db} = \underbrace{\delta(1+r)^2}_{\approx 1} s\sigma b^{-\sigma-1} = 0.046.$$

This is a massive effect, given that r is a *monthly* interest rate. A decrease in the replacement rate of 10 p.p. would generate a decrease in monthly interest rates of 0.46 p.p., approximately 5.5 p.p. in annualized terms. Importantly, the effect would be large even with log utility, corresponding to

$\sigma = 1$. On the other hand, with present-biased preferences, we have

$$\frac{dr}{db} = \underbrace{\beta}_{0.8} \delta(1+r)^2 \left[\underbrace{s\sigma b^{-\sigma-1}}_{0.046} + \underbrace{(1-s)\sigma \left(\frac{c^{exp}}{y^e}\right)^{-\sigma-1} \frac{d}{db} \left(\frac{c^{exp}}{y^e}\right)}_{-0.042} \right] = 0.0037.$$

Interest rates are now much less sensitive, with the bulk of the difference attributable to time inconsistency.⁷

Our theoretical analyses highlighted two points. First, exponential agents in standard models exhibit sizable precautionary behavior. Second, introducing present bias attenuates precautionary behavior: households naively believe that their future selves will start self-insuring for unemployment, but, in equilibrium, never do so. In what follows, we measure the sensitivity of precautionary savings to UI generosity and show that the implications of the present-biased model are more aligned with the data.

3.2 Empirical evidence

We measure the effect of varying the generosity of social insurance on the stock of liquid assets of households. In particular, we leverage an institutional feature of the US UI system: each US state can set its own UI schedule, both for eligibility and for benefits paid conditional on being eligible. We explore cross-sectional, time series variation in UI schedules to estimate the effect of UI generosity on precautionary savings. Contrary to what theory predicts, we obtain statistical zeros of modest economic magnitudes in the data. Then, we run the same regressions in the model as in the data, both for our exponential and present-biased calibrations. Based on the bounds of our confidence bands in the data, we rule out the exponential model as a data generating-process, but fail to reject the present-biased model.

Data. We use the 1984-2008 panels from the Survey of Income and Program Participation. The SIPP is a rotating panel that allows us to observe household balance sheets, demographics, earnings, and state of residence. Wealth information is collected only in particular survey waves. For the most part, we do not observe within-household time series, which limits longitudinal analyses. Thus, we build repeated cross-sections.

We build our sample analogously to Engen and Gruber (2001). In short, we focus on households with stable (labor) earnings histories – as our theoretical mechanism focuses on the behavior of *employed households* who might be at risk of job loss – for which we observe UI generosity and assets. Our final sample contains households observed in all 50 states (plus DC) for 22 years over the 1984-2011 period. See the data appendix for a more detailed description of the SIPP and sample construction.

⁷It is important to notice how magnitudes are larger in the zero liquidity calibration and how interest rates are more sensitive in this case. This is mostly because agents are not allowed to self-insure, which fixes the marginal utility gap between employed and unemployed agents.

We leverage earnings histories to construct measures of UI generosity. We use the UI calculator from Kuka (2020), which builds on previous calculators used in the literature. The calculator contains semesterly information on the particular formulae each state uses to compute benefits or eligibility. These formulae were hand-collected by previous researchers from statutory files. Using information on earnings history and number of children (the two key inputs for determining eligibility and benefits), the calculator outputs how much (and whether) each household head would make in benefits if she were to lose her job in her current state of residence. We convert weekly benefits into a replacement rate by dividing by a (weekly) measure of previous earnings used in many states' UI formulae. We refer to the data appendix for details on the calculator, variable construction, and other institutional aspects of UI.

Although individual-level replacement rates are very granular, identification concerns hamper us from taking them as exogenous measures of policy generosity. Individuals may select into states based on policy generosity. If there are unobserved factors that influence liquid asset accumulation and location decisions, a regression on individual-level replacement rates will provide confounded estimates. More broadly, if there are unobservables that influence an individual-level input of the UI schedule and our outcomes, our estimates will be confounded. At the core of the issue is that individual-level replacement rates leverage information that contaminates our variation of interest, namely *statutory* policy variation.

We address this issue by constructing *simulated policy instruments* (Currie and Gruber, 1996). We start with our sample of household heads and strip these observations from state and time information. This is a national, cross-time sample of individuals. We assume that all such heads lose their jobs in a particular state-time cell (e.g. Georgia in 2002), and compute their eligibility and their benefits using the UI schedule from that particular state-time cell. We then calculate median eligibility and replacement rates for each state-time-number of kids cell. These moments serve as our simulated policy instrument for that particular cell. We repeat this for all states and time periods from 1984:H1 to 2011:H2. The simulated policy instrument solves each concern raised by using individual-level replacement rates. First, we use a national counterfactual sample, thus circumventing the issue of spatial selection. Second, by pooling multiple individuals, idiosyncratic unobservables that might bias the individual measure of policy generosity wash out.

An additional benefit of proxying generosity by simulated policy instruments is to reduce the likelihood that our empirical exercises explore undesirable variation. Since at least Chetty (2008), the literature on UI has noted that individual-level simulated replacement rates can be noisy, i.e. due to measurement error in the inputs used in the calculator. Hence, an additional benefit of our approach is that, by taking medians of the distribution of simulated replacement rates, we considerably reduce noise and avoid undue influence of outlier observations when computing our measures of generosity. In a sense, our simulated policy instrument aids with measurement error concerns in the spirit of Griliches (1986).

Specification and Identifying Variation. Following Engen and Gruber (2001), we estimate the effect of UI generosity on liquid asset holdings by OLS using a canonical two-way fixed effects (TWFE) design:

$$\left(\frac{\text{Liquid Assets}}{\text{Earnings}}\right)_{ist} = \alpha_s + \tau_t + \gamma \times RR_{st} + X'_{ist}\beta + \varepsilon_{ist} \quad (9)$$

In the above, the dependent variable is a measure of the stock of net liquid assets held by household i living in state s at period t , normalized by household earnings. Our coefficient of interest is γ , which measures the change in the levels of liquid assets to income ratios following a 100 p.p. change in replacement rates RR_{st} . Going forward, unless there is potential for confusion, we refer to our dependent variable simply as liquid assets for ease of exposition. We use simulated instruments as our measure of policy generosity. X_{ist} collects any time-varying controls at the state or household level. Last, α_s and τ_t are state and time effects, respectively, and ε_{ist} is a residual. Standard errors are clustered at the state level.

To understand our identifying variation, consider two scenarios. First, suppose there is only cross-sectional variation: states differ dramatically in their UI schedules, but there are no changes in generosity over time. If there are omitted factors correlated both with the decision of a state to have more generous UI and for households in that state to accumulate liquid assets (e.g., richer states lean Democratic, and Democrats favor generous UI), estimates of simple cross-sectional regressions of liquid assets on UI generosity will be confounded. Our inclusion of state effects controls for these time-invariant state-level confounders.

Second, suppose all states have the same UI schedule. If there are omitted trends or cyclical factors affecting the time series of UI and liquid assets simultaneously, estimates of time series regression of liquid assets on UI generosity will be confounded. That would be the case, for instance, if households' asset values fall in recessions, while generosity rises due to federal-level UI extensions. Our inclusion of time effects controls for these state-invariant time-level confounders.⁸ Controlling for time and state effects is essential to isolate unobservable confounders that are present in the data, but that we do not include in the model.

As standard in TWFE designs, our inclusion of state and time effects implies that our source of identifying variation is within-state time series variation. In plain English, our estimates rely neither on comparing California and Missouri in 2001 nor on comparing trends or fluctuations in aggregate liquid asset buffers and UI generosity. Rather, the TWFE estimator considers settings where California changed UI generosity in 2001 (relative to its time-series mean), but Missouri did not; and compares how liquid asset buffers changed in California versus Missouri (relative to each state's time-series means). The estimator averages many such comparisons across states and over time to arrive at a single point estimate.

Besides state and time effects, we include household head-level covariates to control for life-cycle dynamics, earnings potential, and household composition. We include a quadratic polyno-

⁸Time effects are defined at the calendar year-survey wave level. The latter also accounts for survey wave-specific unobservables. The results are essentially unchanged from using solely calendar year effects.

mial on age and a quartic one on log weekly wages, the interaction of age with log weekly wage, dummies for female, caucasian, and married, number of kids, and years of education for the head.⁹ All specifications control for simulated eligibility.

Another concern one may have is that state-time variation in UI generosity is not random. Rather, there could be unobserved factors at the state-time level that simultaneously influence the generosity of UI and households' precautionary savings. For instance, UI changes could be driven by local business cycles or could be correlated with changes in other social insurance programs that also affect precautionary saving. In some specifications, therefore, we also include time-varying state-level controls: the unemployment rate to control for time-varying local economic slack that may be correlated with UI generosity and asset accumulation; measures of welfare spending to control for the generosity of other policies households may have access to;¹⁰ measures of UI spending (regular state payments, regular federal payments, and emergency payments) and net UI reserves to control for budgetary and political considerations; and dummies for whether the state extended UI discretionarily, and for whether Federal emergency UI payments were in place to further control for discretionary UI generosity not captured in our model.

Summary Statistics. Appendix E collects summary statistics for our sample and provides an extensive discussion. We summarize some key points below.

Our sample's balance sheets are qualitatively consistent with the evidence presented in Campbell (2006) based on broader data from the SCF. Our liquid assets measure is *Net Financial Assets* (NFA), which sums assets in banking institutions (interest-bearing or not), assets in other financial institutions (interest-bearing or not), equity in stocks or mutual funds, and other financial assets, and then subtracts unsecured debts. This is the clearest mapping between liquid assets in the model and in the data. As we seek to capture liquid assets, NFA excludes retirement accounts and illiquid assets (e.g., real estate or vehicles).¹¹

The average household holds about \$33,000 in NFA, the bulk of which comes from interest-bearing deposits. Other components of NFA only add mass to the tails of the distribution. Up to the 25th percentile, households in our sample are net borrowers in liquid assets. The median household only has \$0 in NFA. Only at the 75th percentile, we see NFA reaching about \$4,000,

⁹We also add total household social insurance payments, as well as dummies for having ever received UI or non-UI social insurance, primarily to control for other sources of social insurance. However, given our sample characteristics, these controls have limited to no impact on the point estimates.

¹⁰For this purpose, we include state-level expenditures in EITC, Medicare and Medicaid, retirement and disability insurance programs, SNAP, AFDC/TANF, and supplemental security income, normalized by the size of the resident population in a state-time cell. These controls proxy well for state-specific linear trends, but retain a clear economic motivation while not aggressively eliminating state-time variation that can help us identify our effect of interest. We also control for statutory EITC rates, as well as maximum values for EITC and AFDC that households would be eligible for based on their state, time period, and number of children. Last, we include a staggered adoption dummy for the welfare reforms adopted in the 1990s. Bitler and Hoynes (2010) show how there were substantial changes in social insurance generosity following these reforms.

¹¹Although there could be a response of illiquid assets due to precautionary reasons, we abstain from this channel. First, our model only speaks to precautionary savings in liquid assets. Second, we focus on assets that the household can most easily disinvest from, since adjustment costs may be sizable (Kaplan and Violante, 2014). Third, illiquid assets usually yield consumption flows, a motive for accumulation orthogonal to our channel of interest.

about 6% of median yearly earnings. Once we normalize by household earnings,¹² the average household has a ratio of liquid assets to income of 0.06. There is considerable dispersion, with ratios ranging from -0.19 in the 10th percentile to 0.35 in the 90th percentile. But even these contain outliers, which arise due to measurement error or due to the extreme tail behavior of the wealth distribution. To safeguard against the undue influence of these outliers, we symmetrically trim NFA over earnings at the 2.5 and 97.5 percentiles, obtaining liquid asset-to-income ratios ranging roughly from -0.8 to 2.

The median head is eligible and faces a replacement rate of 42%, ranging from 32% at the 10th percentile to 52% at the 90th percentile.¹³ Appendix E shows the policy variation we leverage by state. In particular, for each state, we plot median replacement rates (by number of children) for the 1984:H1-2011:H2 period. First, there is substantial cross-sectional variation in average replacement rates. Some states have generous UI schedules (e.g., Rhode Island), whereas others are less so (e.g., Alabama). Most of the variation in UI generosity is indeed cross-sectional: state effects absorb 62% of the variation in median replacement rates. Second, replacement rates have been either stable or slightly increasing over time. Time effects absorb 15% of the variation in median replacement rates. As a reminder, our TWFE specification leverages neither pure cross-sectional nor pure time-series variation to estimate the coefficient of interest. Third, most changes in UI generosity are incremental, with states often tweaking schedules frequently, mostly by changing minimum or maximum weekly benefits. Nonetheless, there have been episodes with substantial changes in replacement rates within-state, a source of variation we later explore.

TWFE Results. Table 2 reports OLS estimates of specification (9) across different control sets. In all regressions, the dependent variable is the NFA divided by annual household earnings, and UI generosity is measured with the median replacement rate. Throughout, we emphasize lower bounds of 95% confidence bands, which illustrate the strongest effects one could expect under conventional significance levels, and therefore are relevant figures for ruling out competing models.

In column 1, we show the results of a regression featuring time and state fixed effects, as well as head characteristics. We fail to reject a null of zero effects. The two-standard error lower band rules out decreases in asset-income ratios smaller than -0.006 per 10 p.p. increase in median replacement rates. To contextualize the magnitude of these effects, we compute the distribution of first differences in yearly median replacement rates, after residualizing by state and time effects and winsorizing at the 1st and 99th percentile. On average, median replacement rates change by 0.26 p.p. year-over-year, with a standard deviation of 2.5 p.p. For this average change in replace-

¹²More precisely, we average monthly total household earnings over all waves prior and including the wave of the wealth observation, and then annualize the measure. In this way, we hope to smooth any short-lived fluctuations in earnings driven by e.g. short-term unemployment.

¹³Consistent with our sample selection efforts, few household heads ever received UI. Amounts received in social or private insurance transfers are negligible. The replacement rate and the social insurance reciprocity in our sample might seem low compared to our calibration. They are not. Social or private insurance reciprocity in the model is zero among the employed, consistent with our sample. And the replacement rates we find are driven by higher earnings among the employed, and consistent with take-up adjusted replacement rates in the model. When calibrating replacement rates in the model, we match replacement rates across the income distribution, a more demanding task.

	Net Financial Assets		
	(1)	(2)	(3)
Median RR	-0.0051 (0.0299)	-0.0391 (0.0261)	0.0109 (0.0298)
Outcome mean	0.06	0.06	0.06
N	178119	178119	178119
Individual controls	✓	✓	✓
State controls		✓	✓
Time FE	✓		✓
State FE	✓		✓

Table 2: Two-way fixed effects regressions

Notes: this table reports estimates of regressions of net financial assets (normalized by annual household earnings) on simulated unemployment insurance median replacement rates. Net financial assets is measured as the sum of interest-bearing or non-interest bearing assets on banking institutions, interest-bearing or non-interest bearing assets on other financial institutions, and stocks and mutual fund holdings, net of unsecured debts. Simulated replacement rates are defined at the state-year-semester-number of kids level. Column (1) reports estimates of regressions that include state and time fixed effects, as well as individual-level controls for the household head (see the main text for a list). Column (2) removes the fixed effects, but includes household head-level controls and time-varying state controls (see main text). Column (3) adds back state and time FE. Standard errors are clustered at the state level.

ment rates, using the lower bound of our confidence bands, the corresponding change in liquid asset-income ratios would be -0.0001, or 0.28% of the mean.¹⁴

We remove fixed effects in column 2, adding instead state-level time-varying controls to our individual controls. Most importantly, we control for state-level unemployment rates and several measures of welfare spending. The former controls for potential cyclical confounders related to local economic slack, whereas the latter controls primarily for trends in the generosity of other welfare programs. For example, it could be the case that states differentially alter the generosity of UI in recessions, so that our measure of generosity would be confounded by cyclical variation that is not quasi-exogenous. Alternatively, it could be that other welfare programs are expanded in these scenarios, leading us to understate the effect of UI generosity in mitigating consumption risk. Our state-time controls mitigate the risk that these alternative stories are driving our estimates.

In column 2, we obtain economically modest, but still statistically non-significant effects. However, state-level controls alone do not absorb unobserved aggregate shocks or time-invariant state-specific characteristics. Hence, we take column 3 as our preferred specification, in which we add back state and time fixed effects. We find that the two-standard error lower band rules out decreases in asset-income ratios smaller than -0.005 per 10 p.p. increase in median replacement rates. In absolute terms, assuming no effects on earnings (ie. a constant denominator), a 10 p.p. increase in replacement rates would lower liquid assets by at most 0.5% of income under conventional significance levels, or about \$350 at the sample mean for household earnings.

¹⁴Without residualizing by the fixed effects, the standard deviation rises to 3.5p.p. However, residualized changes better reflect the within-state time variation we leverage for estimation.

We assess the robustness of our main results in table 3. Column 1 repeats our baseline from Table 2. In column 2, we consider gross financial assets as our measure of liquid assets (i.e., not subtracting unsecured debts), confirming our results are not driven by focusing on net measures. Column 3 excludes stocks and mutual funds, which are risky assets that are not included in our model. Restricting to safe liquid instruments for self-insurance does not alter our conclusions. Column 4 adjusts our replacement rate measure for eligibility by assigning a zero replacement rate if ineligible. The result is essentially unchanged, reflecting the fact that almost all households in our sample are eligible for UI benefits. Column 5 considers a less granular measure of statutory UI generosity (see Appendix D.2 for details). While using a less granular instrument increases the standard errors of our point estimates, the results remain statistical zeros.

Column 6 reproduces column 1, but without trimming the asset-income ratio distribution at all. As can be seen, the influence of extreme outliers leads to meaningless point estimates and standard errors. If anything, column 6 reinforces the need to trim the outcome distribution in the presence of such extreme outliers. Column 7, on the other hand, reports results if we trim more aggressively, discarding both the bottom and top 5% of the asset-income ratio distribution. The results are consistent with those in column 1. Last, column 8 double-clusters standard errors at the state and year levels. This alternative inferential assumption does not alter our conclusions.

	Baseline	GFA	Excl. risky	Elig. adj.	< Granular	No trim	Trim 10%	Time cluster
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Median RR	0.0109 (0.0298)	0.0229 (0.0238)	-0.0066 (0.0233)	0.0100 (0.0293)	0.0010 (0.0364)	-0.2432 (0.2031)	0.0132 (0.0241)	0.0109 (0.0294)
Outc. mean	0.06	0.10	0.01	0.06	0.06	0.17	0.04	0.06
N	178119	177812	176024	178119	178119	187504	168731	178119

Table 3: TWFE Robustness

Notes: this table reports estimates of regressions of net financial assets (normalized by annual household earnings) on simulated unemployment insurance median replacement rates. Each column on the table corresponds to a different robustness check. Column (1) corresponds to the results in Column (3) of Table 2, our baseline specification. Column (2) replaces NFA with Gross Financial Assets as the outcome. Column (3) excludes stocks and mutual funds from NFA. Column (4) adjusts replacement rates for eligibility. Column (5) instead uses a less granular, state-time level measure of UI generosity rather than a state-time-number of kids measure. Column (6) reports results without any trimming of the dependent variable. Column (7) repeats Column (6), but trimming the top and bottom 5% of the outcome distribution. Column (8) reports results with double-clustered standard errors at the state and year level. Except for column (8), standard errors are clustered at the state level.

In sum, we find no statistically significant relationship between UI generosity and liquid asset holdings. Our estimated effects are modest in magnitudes: as we will see, the lower bound of our confidence bands rule out the effects predicted by exponential models but are consistent with the effects predicted by the behavioral model.

There may be individual-level variation in UI generosity that is baked into the policy schedule (e.g., income), but that is not captured by our simulated instrument. In the appendix, we explore this variation, and estimate by two-stage least squares the exactly-identified IV model:

$$\left(\frac{LiquidAssets}{Earnings}\right)_{ist} = \rho_s + \psi_t + \varphi \times RR_{ist} + X'_{ist}\xi + \epsilon_{ist}$$

$$RR_{ist} = \phi_s + \omega_t + \nu \times RR_{st} + X'_{ist}\lambda + \varrho_{ist}$$

The only novelty relative to specification (9) – the reduced form of the system above – is that we now use individual-level simulated replacement rate RR_{ist} , but instrumented with our simulated policy RR_{st} .

In interpreting the OLS versus IV coefficients, it is paramount to conceptualize the policy exercise we are interested in. When we observe variation in replacement rates, there are two possibilities. First, a reform has occurred, and the parameters of the UI schedule have changed. This is a *statutory policy* variation. At the same time, parameters may have remained fixed, but inputs into benefits calculations may have changed. This is *not* a policy variation. Our model makes statements about *policy* variation. In plain English, when confronting our model with the data, the relevant identified moment is γ . Therefore, we focus on our OLS estimates, presenting the IV results briefly for completeness.

Panel B of Appendix table 12 reports first-stage results. In our preferred specification (column 3), a 1 p.p. increase in the median replacement rate leads to a 0.68 p.p. increase in the individual-level replacement rate (Panel B). The effects are all significant at the 1% level. Both the magnitudes and the statistical significance of the estimates are consistent across control sets. There is little concern about weak instruments: Kleibergen and Paap (2006) robust F-statistics are all large, much above conservative thresholds.¹⁵ Panel A reports two-stage least squares estimates. First, relative to table 2, magnitudes are larger: in the exactly-identified case, 2SLS estimates are numerically equivalent to the ratio of reduced form to first-stage estimates. Expectedly, our lower than 1-1 passthrough estimates inflate our IV estimates. Second, the sensitivity to controls parallels exactly that observed in Table 2. To show the robustness of our results, we now turn to a different empirical strategy.

Stacked Regressions A recent econometrics literature has raised concerns about interpreting the canonical TWFE specification in a range of settings (see de De Chaisemartin and D’Haultfoeuille, 2022 for a review). First, we discuss how these concerns apply to our case. Second, we assess the sensitivity of our TWFE estimates to these concerns by leveraging an alternative specification: stacked regressions (Gormley and Matsa, 2011, Cengiz et al., 2019). Overall, we find that our estimates remain similar once we move to this alternative specification, giving credence to the results presented above.

Under the assumption that changes in liquid asset ratios for households in states with smaller changes in generosity provide reasonable counterfactuals for changes in liquid asset ratios that

¹⁵Since we are in an exactly-identified setting, the Kleibergen and Paap (2006) F-statistic is identical to the Olea and Pflueger (2013) F-statistics robust to heteroscedasticity, clustering, and autocorrelation. This implies that maximal bias sizes in the IV estimates are also small.

would have been observed for states with larger changes in generosity (had their generosity changed similarly), our TWFE specification can be interpreted as estimating the average treatment effect on the treated through a generalized difference-in-difference design (Angrist and Pischke, 2009). Previous papers in applied microeconomics have emphasized this design-based interpretation, especially in settings with clean policy reforms as sources of quasi-random variation, for which the “strong” parallel trends (Callaway and Sant’Anna, 2021) assumption above can be defended soundly. We have avoided labeling our exercises as such. Instead, we have sought to explore most state-time variation in UI generosity in our data, finding a robust moment that could refute (or not) a theoretical property of a class of models.

Regardless of interpretation, the canonical TWFE specification estimates our coefficient of interest by performing many comparisons of the following form: suppose a state changes replacement rates, whereas some other states do not. What is the relative change in liquid asset ratios across households in these states? TWFE averages these effects across comparisons to arrive at a single estimate. The interpretability issues raised by the literature arise precisely due to *how* TWFE *mechanically* performs such comparisons and averages across them.

Given our setting, we see two potential concerns. First, our treatment is continuous. As a result, some of the comparisons are between states whose replacement rates have not changed much or between states whose replacement rates have changed sizably. In turn, informative comparisons are those in which “treated” states saw a sizable change in replacement rates, whereas “control” states did not. Observing modest changes in liquid assets would then be more compelling evidence of a small effect. The TWFE averages all comparisons, which may affect estimates. Second, there is variation in dosage over time. Ideal experiments compare states with a one-off change in replacement rates to states with no changes at all. However, Appendix D shows that states often change UI generosity, with a much more complex timing than simple one-off or staggered changes. This can give rise to “forbidden comparisons”, in which the TWFE estimator uses early-treated states as controls for later-treated states (Callaway and Sant’Anna, 2021).

To address these concerns, we leverage stacked regressions (Cengiz et al., 2019). The central idea is to explore clean comparisons to estimate γ . These are “events” in which (i) there is a simple, one-off change in replacement rates in a treated state, and (ii) there is little to no change in replacement rates in control states. The first requirement alleviates concerns about treatment timing variation, and the second is about the cleanliness of the comparisons made by the estimator. The econometrician then builds datasets of matched treatment-control states for each event, stacks them together, and re-estimates the model with event-specific state and time effects. Relative to the canonical TWFE, stacked regressions only leverage within-event variation present in these cleaner comparisons, taking a variance-weighted average of the effects estimated from each individual event. If reforms are indeed quasi-exogenous conditional on controls, one can interpret the exercise in a design-based sense as estimating a generalized difference-in-difference design.

First, we identify states that experience substantial increases in replacement rates in a given year. Starting from a state-year panel of simulated replacement rates, we mark all state-year pairs

that experienced a change in weekly benefits and replacement rates of over 5 p.p. year-over-year. This isolates 125 state-year pairs – the unconditional probability of an event is 8.75%. Hence, the events we isolate are sizable and infrequent reforms. We verify that, in all cases, the jumps we identify were preceded by some change in the formulae for the UI schedule. Henceforth, we refer to these events simply as reforms.

We fix a window of 6 years prior to 10 years after the reform – the event window. For each reform, we check whether the state had another reform within the event window. If so, we discard the event.¹⁶ Otherwise, we include the event in the dataset and mark the state as treated (e.g. California in 2002). We then construct a clean control set for the event. For a state to be a clean control, it must not have been treated within the event window.¹⁷ With treated and control units selected, we bring in wealth and demographic variables at the household level from our SIPP sample.

We repeat this procedure for all substantial reforms we have identified. Out of the initial 125 state-year pairs, 39 are clean events that can be used for estimation. We stack all datasets, assigning to each an identifier. We then estimate (9) by OLS, but now replacing state and time effects with state-by-dataset and time-by-dataset effects. This ensures we use variation within each event h , averaging across events to estimate γ . Last, in line with previous literature, we cluster standard errors conservatively at the state level. Formally, we estimate

$$\left(\frac{LiquidAssets}{Earnings} \right)_{isth} = \alpha_{sh} + \tau_{th} + \gamma \times RR_{sth} + X'_{isth} \beta + \varepsilon_{isth} \quad (10)$$

Table 4 presents the results. Column 1 adds individual controls, as well as state and time fixed effects. Contrary to theory, we pick up a significant positive correlation between replacement rates and liquid asset-income ratios. A natural concern is that this result is driven by omitted variables at the state-time level. In column 2, we replace the fixed effects with state controls and find negative effects, but statistically non-significant. Column 3 is our preferred specification. It includes all controls present in Column 3 of Table 2. The two-standard error lower bound rules out decreases in asset-income ratios smaller than -0.0048 per 10 p.p. increase in median replacement rates.

In Appendix Table 13, we show our results are robust to a battery of robustness tests. First, dropping all state-year observations with less than 25 observations leaves our baseline result unchanged. Second, discarding all events that happen during NBER recessions (and may therefore violate the quasi-exogeneity of reforms underlying a difference-in-differences interpretation of our design) barely moves our point estimates and standard errors. Our results survive if we exclude stocks or mutual funds from NFA, or trim more of the distribution of NFA. Consistent with the results of table 3, results are not sensible in the specification where we don't trim the outcome whatsoever, highlighting how extreme outliers may drive results and reinforcing the need for some

¹⁶A exception to this rule is when we pool reforms. For example, California experienced large increases in generosity both in 2000 and 2001. These events are treated as a single reform.

¹⁷That is, it must not have experienced a reform in UI where both weekly benefits and replacement rates jumped over 5p.p. year-over-year.

	Net Financial Assets		
	(1)	(2)	(3)
Median RR	0.1063** (0.0398)	-0.0227 (0.0429)	0.0577 (0.0530)
Outcome mean	0.07	0.07	0.07
N	1129022	1129022	1129022
Individual controls	✓	✓	✓
State controls		✓	✓
Time × event FE	✓		✓
State × event FE	✓		✓

Table 4: Stacked regressions

Notes: this table reports estimates of regressions of net financial assets (normalized by annual household earnings) on simulated unemployment insurance median replacement rates in our stacked differences design. There are 39 reforms in our sample. Column (1) reports estimates of regressions that include time-by-event and state-by-event fixed effects, as well as household head controls. Column (2) replaces fixed effects by state-level controls. Column (3) adds fixed effects back. Standard errors are clustered at the state level. *, **, *** denote significance at the 10, 5 and 1% levels.

trimming. Our results are sensitive to two robustness tests: when we use a less granular measure of replacement rates or when we consider gross financial assets. In the former, state-level controls soak up too much variation in UI for us to be able to detect any reasonable effects, as evidenced by the inflated standard errors. On the latter, the distribution of GFA to earnings trimmed at the 95th percentile is too narrow: the max is only 1.17, which is small compared to top percentiles in the model. Trimming the GFA-earnings distribution at 98th percentile instead would increase the max to 2.5, more in line with model-based wealth distributions, and also imply estimates quantitatively consistent with our other results.

In sum, the lower bounds of the 95% confidence bands implied by the point estimates in Column 3 of Table 4 are consistent with those implied by our baseline TWFE specification. The similarity of the results produced by the canonical TWFE and the stacked TWFE reinforces our belief that our estimated effects are not artificially lower due to some of the mechanics of the canonical TWFE design.

3.3 Bringing together model and data

We now show that we can use the evidence above to discriminate across models. To do so, we run the same reduced-form regressions in model-generated data using both the exponential and present-biased calibrations. As our identification relies on variation in UI generosity over time, it is first necessary to allow for time-varying replacement rates in the model and postulate a law of motion for them, which we do for this experiment only. For all other exercises in the paper, replacement rates are assumed to be constant. To guide our modeling, figure 4 shows median replacement rates in California and Louisiana from 1984 to 2011, each line corresponding to different

numbers of children in the household. There is a clear pattern: they often decrease over time, as nominal payments are eroded by inflation, but are subject to infrequent reforms that substantially change the generosity of benefits.

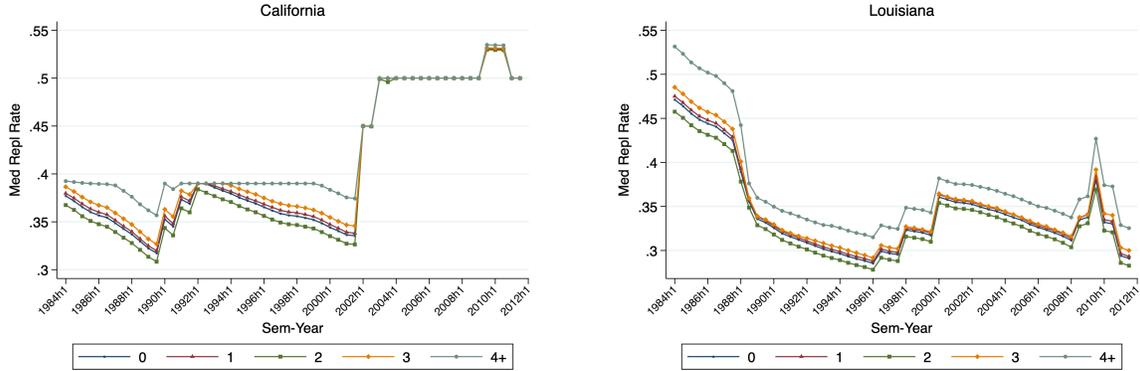


Figure 4: Median replacement rates in California and Louisiana by number of children.

Motivated by these patterns, and constrained by the fact that replacement rates in the model must follow a finite-state Markov chain for tractability, we assume the following stochastic process. The replacement rate is assumed to take the values

$$b_t \in \{0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55\}.$$

We keep b_t bounded between 0.25 and 0.55 as our time series for median state-level replacement rates rarely fall outside this interval (see Appendix E). Each period, with probability $p^{ref} > 0$ there is a UI reform and the replacement rate jumps to a neighboring point on its grid of possible values, increasing or decreasing with equal probabilities (except at the boundaries of the grid, in which there is a single direction to move). Moreover, we assume that the UI cap \bar{b}_t , now also time-varying, is scaled accordingly, i.e., if $b_t = 0.55$ (10% above its baseline value), then \bar{b}_t will also be 10% above its baseline value. We set the reform probability $p^{ref} = 1/120$, which corresponds to a mean interval between reforms of 10 years, or a median interval of approximately 7 years. This delivers an annual autocorrelation of replacement rates of 0.987, close to what is observed in the data.¹⁸ All simulations here are done in partial equilibrium, emulating small states of negligible size compared to the national economy.

As in the SIPP, we draw repeated cross-sections of households every year – the frequency according to which most wealth observations are made available. We simulate 51 states for 22 years and, for each state, sample 180 households. This delivers a number of observations similar to our empirical setup. We repeat this procedure 300 times and report the average estimated coefficient and standard errors across simulations.

Table 5 shows the results. The exponential model delivers coefficients that are substantially

¹⁸We estimate an AR(1) model for median replacement rates for each state in our sample. The average semestery coefficient is 0.984, which implies an annual autocorrelation of 0.968.

	Data			Exponential	Present-biased
	NFA	GFA	Excl. risky		
Median RR	0.0109 (0.0298)	0.0229 (0.0238)	-0.0066 (0.0233)	-0.0636 (0.0267)	-0.0207 (0.0525)
Lower 95% confidence band	-0.0475	-0.0237	-0.0523	-	-

Table 5: Regressions in model-simulated data.

larger than our empirical estimates, and always outside confidence intervals. This model is, therefore, rejected by the data. On the other hand, the model with present-biased households delivers much lower coefficients that lie within confidence intervals. Moreover, for our simulated sample size, the present-biased model generates coefficients that are not statistically different from zero, matching another feature of our empirical estimates. We now turn to a robustness exercise that shows that this is indeed a feature of many models with exponential discounting, not only our specific one.

Robustness. When comparing the coefficients generated by the exponential model with the confidence intervals obtained empirically, we are not testing the null hypothesis of standard preferences alone. Rather, it is a joint test of standard preferences and all other underlying assumptions in the model, such as functional forms and parameter values. As a robustness test of our results, table 6 shows the coefficients obtained from many variants of the exponential model. For all variants, we recalibrate the discount factor δ to keep the same real interest rate in equilibrium. This experiment is therefore conservative, as the resulting models are not guaranteed to match the same empirical moments we target in our calibration, i.e., they may generate unrealistic marginal propensities to consume or consumption drops upon unemployment. Importantly, in all variants results lie outside the 95% confidence bands obtained empirically.

Column 1 corresponds to a model in which agents have log utility, corresponding to a coefficient of relative risk aversion $\sigma = 1$. In column 2, we generalize our baseline model by allowing the job finding rates of unemployed agents that receive UI to be functions of the replacement rate. We assume that the job finding rate f_t of UI recipients is given by

$$f_t = f - 0.02 \cdot \log \left(\frac{b_t}{b} \right),$$

where b and f are the baseline calibration values. We set a slope of -0.02 , as estimated by Ganong et al. (2021). In column 3 we reduce the supply of liquid assets in the economy by half, while column 4 shows what happens as we increase it by 50%. In column 5, follow a different calibration strategy for the exponential model. We set $\underline{a} = -1.48$ and $\chi = 0.0098$, as in the present-biased calibration, but recalibrate y^{priv} the same ratio of consumption between employed and unemployed agents. It shows that our results do not rely on the specific assumption of $\underline{a} = 0$ in our baseline

	Exponential model							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\sigma = 1$	Varying f	Low liq.	High liq.	Borrowing	$y^{priv} = 0$	$N^{ui} = \infty$	$\xi = 1$
Median RR	-0.0551	-0.0589	-0.0611	-0.0607	-0.0802	-0.1885	-0.0842	-0.0903
	(0.0315)	(0.0267)	(0.0122)	(0.0423)	(0.0295)	(0.0294)	(0.0285)	(0.0285)

Table 6: Regressions for variants of the exponential model.

calibration.

So far, most variations of the model delivered slightly smaller coefficients, but still outside the confidence intervals estimated in the data. In columns 6 through 8, we show the response of precautionary savings predicted by less conservative models, but which are nevertheless commonly employed in the literature (McKay and Reis 2016, Schaab 2020). Column 6 shows a model without private insurance, in which the only source of insurance is the statutory UI schedule. Columns 7 and 8 show the cases without UI expiration and with full take-up. In these cases, precautionary savings are more sensitive to changes in UI generosity, as UI is a more important source of insurance for the unemployed, either because more unemployed agents receive benefits or because they have no other sources of income.

The experiments above show that the standard heterogeneous-agent model delivers strong precautionary savings responses to changes in income risk. These responses are larger than the ones observed in the data and can be rejected with at least 95% confidence. The present-biased model, on the other hand, delivers predictions that are in line with and not rejected by the data. The natural next step is to understand the consequences of this for the dynamics of aggregate consumption and stabilization policy. Before doing that, however, it is necessary to increment our model with the necessary features for business cycle analysis, such as monetary policy and nominal rigidity. This is the goal of the next section.

4 Introducing Business Cycles

So far, our analysis has been based either on comparing steady states with different UI policies, or by simulating a panel of independent states that behave like small open economies – without general equilibrium feedback effects of aggregate shocks. In this section, we start by introducing business cycle features in the model. Then, we estimate the model parameters using Bayesian methods.

4.1 Closing the model

Wage stickiness. Stabilization policy operates by stimulating aggregate demand during recessions. For aggregate demand to have a meaningful role in our model, however, it is necessary to

introduce some type of nominal rigidity. We do so by introducing nominal wage rigidity as in Auclert et al. (2018). There is a continuum of unions indexed by j that set nominal wages. Each household in the model supplies labor to all unions. Unions then maximize a utilitarian aggregator of the welfare of employed workers, given by

$$\mathbb{E}_t \sum_{k=0}^{\infty} \delta^k \frac{1}{E_{t+k}} \int [u(c_{i,t+k}) - v(n_{t+k})] 1_{i,t+k}^e di, \quad (11)$$

where 1_{it}^e is an indicator function for employed households. For simplicity, we make the assumption that unions are not present-biased. Conceptually, one can think of unions as acting to partly mitigate the dynamic inconsistency issues that arise from present bias, acting as a commitment device on wage-setting decisions. Moreover, we follow Auclert et al. (2018) in assuming that unions demand the same number of hours n_t from each household.

As in Calvo (1983), in each period a fraction $1 - \theta$ of unions is able to optimally reset nominal wages. The remaining fraction θ can partially index their wages to past inflation, adjusting according to the rule

$$W_{jt} = W_{j,t-1} \left(\frac{P_{t-1}}{P_{t-2}} \right)^\iota, \quad (12)$$

where W_{jt} is the nominal wage charged by union j , P_t is the price level, and ι is the indexation parameter. Indexation generates inflation inertia, which is necessary to match the time series behavior of inflation in the data, as we do later on.

Firms are still perfectly competitive and operate the same linear technology (5). Total labor input N_t is now a composite of the labor supplied by each different union N_{jt} , as in Dixit and Stiglitz (1977):

$$N_t = \left[\int N_{jt}^{1-1/\nu_t} \right]^{\frac{\nu_t}{\nu_t-1}}. \quad (13)$$

Above ν_t is the elasticity of substitution across the types of labor supplied by different unions.

As firms are perfectly competitive, the price level is given by

$$P_t = \frac{W_t}{Z}, \quad (14)$$

where the aggregate wage index is given by

$$W_t = \left[\int W_{jt}^{1-\nu_t} \right]^{\frac{1}{1-\nu_t}}.$$

As firms are perfectly competitive, the real wage is constant at $\frac{W_t}{P_t} = Z$.

As shown in Appendix B, as a result of this setup, inflation evolves as follows. First, define the aggregate labor wedge

$$\mu_t = \frac{v'(n_t)}{\frac{1}{E_t} \int u'(c_{it}) (1 - \tau_l) w_t e^{z_{it}} 1_{it}^e di'}$$

which is a measure of labor market tightness (when μ_t is high, the disutility of labor is large compared to the average marginal utility of consuming the proceeds of this labor). Inflation π_t then follows a standard New Keynesian Phillips Curve with indexation:

$$\pi_t = \frac{\delta}{1 + \delta\iota} \mathbb{E}_t \pi_{t+1} + \frac{\iota}{1 + \delta\iota} \pi_{t-1} + \frac{(1 - \theta)(1 - \delta\theta)}{\theta(1 + \delta\iota)} d \log \mu_t + \varepsilon_t^w, \quad (15)$$

where the d operator denotes deviations from steady-state levels and ε_t^w is a wage mark-up shock that comes from shocks to the elasticity ν_t .

Labor demand. As shown in equation (6), total labor demand can be decomposed as a product of employment E_t and hours per worker n_t . To correctly capture the role of unemployment risk in driving aggregate consumption, we must specify to what extent E_t and n_t vary over the cycle. We postulate the following relationship between deviations from steady-state values:

$$d \log E_t = \lambda \cdot d \log N_t$$

$$d \log n_t = (1 - \lambda) \cdot d \log N_t,$$

where λ is a parameter that captures the cyclical nature of employment. If total labor demand N_t is 1% above its steady-state level, employment E_t and hours per worker n_t will be $\lambda\%$ and $(1 - \lambda)\%$ above their steady-state values, respectively. This is possible due to wage rigidity since, as usual in these models, unions commit to supplying whatever quantity n_t firms demand at the prevailing nominal wage.

To pin down employment E_t in equilibrium, we have one degree of freedom in excess, as both the separation rate s_t and the job finding rate f_t are allowed to vary. We assume that

$$ds_t = -\eta \cdot df_t, \quad (16)$$

where the parameter η captures the relative importance of separation versus job finding rates in driving employment fluctuations.

Monetary and fiscal policies. There is a monetary authority that follows the Taylor (1993) rule

$$i_t = \phi^i i_{t-1} + (1 - \phi^i)(i + \phi^\pi \pi_t) + \varepsilon_t^i,$$

where ϕ^π determines how aggressive the central bank is when reacting to inflation, ϕ^i is a parameter that introduces inertia in nominal interest rates, and ε_t^i is a monetary policy shock. As with inflation, introducing inertia is necessary for matching the time series behavior of interest rates. The real interest rate r_t that appears in the consumer and government budget constraints (2) and (4) is given by

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t}.$$

We postulate that the fiscal authority sets labor taxes τ_t each period to enforce the following fiscal rule:

$$dR_t = \phi^B \cdot dB_{t-1}.$$

This rule means that government revenues R_t , defined in (3) responds positively to the past debt level: the government increases revenues by ϕ^B dollars for every dollar past debt is above its steady-state value.

4.2 Bayesian estimation of model parameters

In section 2.2, we determined values for the parameters that affect the steady state of the model. Before proceeding, however, it is necessary to determine values for the remaining parameters – the ones that govern the dynamic response of the economy to aggregate shocks. We do so using a mix of calibration and Bayesian estimation.

First, we calibrate λ – the degree to which employment drives labor demand variation as opposed to hours per worker – using aggregate data on total hours, average weekly hours, and employment in the U.S. nonfarm business sector for the 1960-2019 sample. We take logarithms of these quarterly time series and extract their cyclical components using the Hodrick and Prescott (1997) filter with the standard smoothing parameter value of 1600. The left panel of figure 5 shows results, suggesting a value $\lambda = 0.817$.

Regarding η , the degree to which job finding rates drive employment fluctuations as opposed to separation rates, we use data from Shimer (2012). The right panel of figure 5 shows the job finding probability on the horizontal axis and the employment exit probability on the vertical axis. Both series correspond to the cyclical components obtained from the Hodrick-Prescott filter as described above. The sample period is 1960-2007. The best-fitting line is obtained from the first principal component, which avoids the problem that ordinary least squares coefficients depend on which variable is placed on the right-hand side. This analysis suggests a value $\eta = 0.0142$. The last two parameters we calibrate are the Taylor rule coefficient $\phi^\pi = 1.5$, a standard value in the literature, and the fiscal rule parameter $\phi^B = 0.1$, a value close to the one used in Auclert et al. (2020).¹⁹

We estimate the remaining parameters using aggregate time series data and Bayesian inference, computing the log-likelihood for the linearized model using the methods developed in Auclert et al. (2021). Importantly, all impulse response functions considered in this section are with respect to linearized shocks around the deterministic steady state of the economy. The efficacy of UI as an automatic stabilizer depends, as will be shown below, on the nature of the shocks that affect the economy, hence the usefulness of this estimation exercise: it ensures an empirically realistic mix of shocks. We use three time series: (i) aggregate consumption growth, measured by the log change of the real Personal Consumption Expenditures (PCE) index, (ii) inflation, measured by the PCE

¹⁹Auclert et al. (2020) have a slightly different fiscal rule in which the government adjusts taxes directly, not revenues, in response to the level of debt. However, the average labor income (in the steady state) in the model is normalized to 1, so equation (3) implies that $R_t \approx \tau_t$.

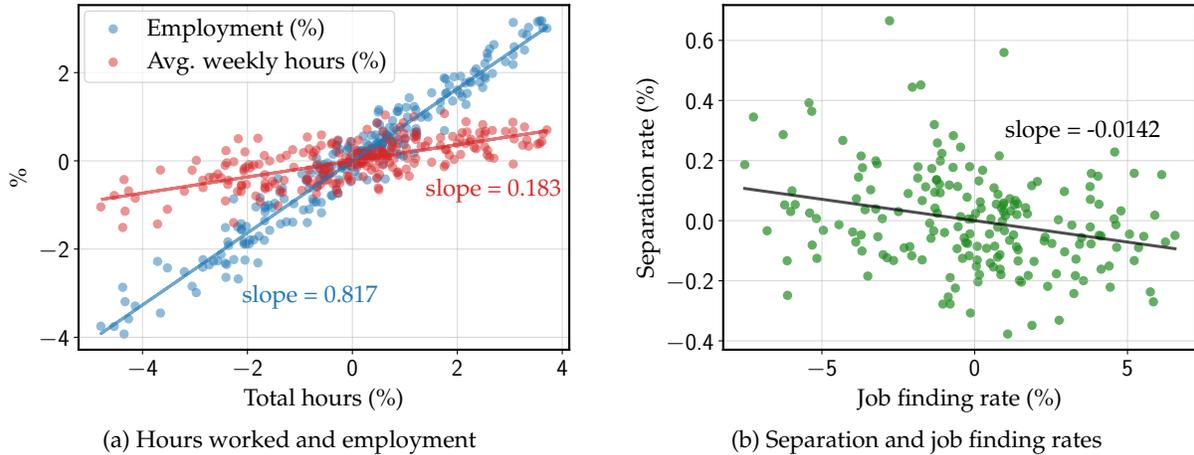


Figure 5: Parameters governing labor demand variation.

deflator, and (iii) the Fed Funds rate. We take our sample to be 1984-2007, starting after inflation was conquered in the 1980s and ending when monetary policy hit the zero lower bound.

To avoid stochastic singularity, we need therefore at least three shocks in our model. We so far have a wage markup shock ε_t^w , which we assume to follow an AR(1) process with persistence ρ^w and standard deviation of innovations σ^w , and the monetary shocks ε_t^i , also assumed to be given by an AR(1) process with parameters ρ^i and σ^i . We introduce here a discount factor shock ε_t^d , often taken to be a proxy for shocks to consumption demand (Schaab, 2020), in such a way that agents discount factors are given by $\delta \cdot \exp(\varepsilon_t^d)$. This last shock also follows an AR(1) process with parameters ρ^d and σ^d .

We have nine estimated parameters: six governing the stochastic process of shocks, the interest rate persistence ρ^i , the wage rigidity parameter θ , and wage indexation ι . Table 7 shows the results. The prior distributions we employ are standard in the literature (Smets and Wouters, 2007). We draw 200,000 parameter values from the posterior distribution using the Metropolis-Hastings algorithm. To keep the exponential and present-biased models comparable, we use the same parameter values for both. Estimation is conducted using the present-biased calibration only. Finally, Appendix C shows impulse response functions for the estimated model.

5 Implications for Stabilization Policy

We now turn to the implications of dampened precautionary saving for stabilization policy. As a stabilization tool, unemployment insurance acts not only by increasing the disposable income of the unemployed but also by shifting the balance of risk faced by employed households – an increase in UI affects the consumption of employed agents as well, as long as they engage in precautionary saving against unemployment risk. We explore two different roles of unemployment insurance. First, as a discretionary policy tool used to increase consumption. For this experiment, we compute the fiscal multipliers of temporary UI extensions, as in Kekre (2021). Then we turn to

		Prior Distribution			Posterior Distribution		
		Distr.	Mean	St. Dev.	Mode	Mean	St. Dev.
ϕ^i	Taylor rule inertia	Beta	0.85	0.1	0.930	0.927	0.010
θ	Wage stickiness	Beta	0.85	0.1	0.982	0.982	0.003
ι	Wage indexation	Beta	0.5	0.2	0.589	0.591	0.064
ρ^d		Beta	0.5	0.2	0.978	0.977	0.050
ρ^i	Shock persistence	Beta	0.5	0.2	0.545	0.477	0.082
ρ^w		Beta	0.5	0.2	0.343	0.336	0.087
σ^d		Inv. Gamma	0.05	1	0.025	0.027	0.005
σ^i	Shock s.d.	Inv. Gamma	0.05	1	0.020	0.022	0.003
σ^w		Inv. Gamma	0.05	1	0.032	0.030	0.007

Table 7: Priors and posteriors of estimated parameters.

the role of UI as an automatic stabilizer. Unemployment is countercyclical, so to the extent that UI subsidizes the consumption of unemployed workers, it should reduce the volatility of aggregate consumption over time. We measure the efficacy of UI as an automatic stabilizer by the increase in aggregate consumption volatility in a counterfactual economy with less generous UI, in the spirit of McKay and Reis (2016).

5.1 Temporary UI extensions

The policies we implement here are as follows. Starting from the steady state, at time $t = 0$ the government announces a transfer of amount T to unemployed agents that lasts for 12 months, i.e., for periods $t = 0$ through $t = 11$, unemployed agents may collect T dollars from the government. We consider two variations of UI extensions: one in which *all* unemployed agents receive these transfers and another in which only unemployed agents with *expired benefits* are eligible. The first is an overall increase in UI generosity that affects replacement rates, eligibility, and duration simultaneously. This policy is useful for illustrating the precautionary saving mechanism: as we show below, a large part of the increase in output, in this case, comes from the consumption of *employed* agents, who do not receive benefits. The second policy corresponds to a 12-month increase in UI duration,²⁰ similar to the ones studied in Kekre (2021). As we work with linearized responses, the value of T only scales the effects but does not change the fiscal multiplier of these policies. For simplicity, we set T such that the government would spend an additional \$1 dollar each period if the unemployment rate was kept constant at its steady-state level.²¹

Figure 6 shows how output responds. The effect is substantially larger in the exponential econ-

²⁰The way the policy is set up, i.e., transferring resources to *all* agents with expired UI, may cause some workers with UI expired for more than 12 months to also be eligible for these transfers. They correspond, however, to less than 0.1% of the total population in the model, and are therefore negligible from a quantitative perspective.

²¹This is without loss of generality, as any increase in spending due to changes in unemployment would be second-order, but we only solve for aggregate dynamics to a first-order approximation.

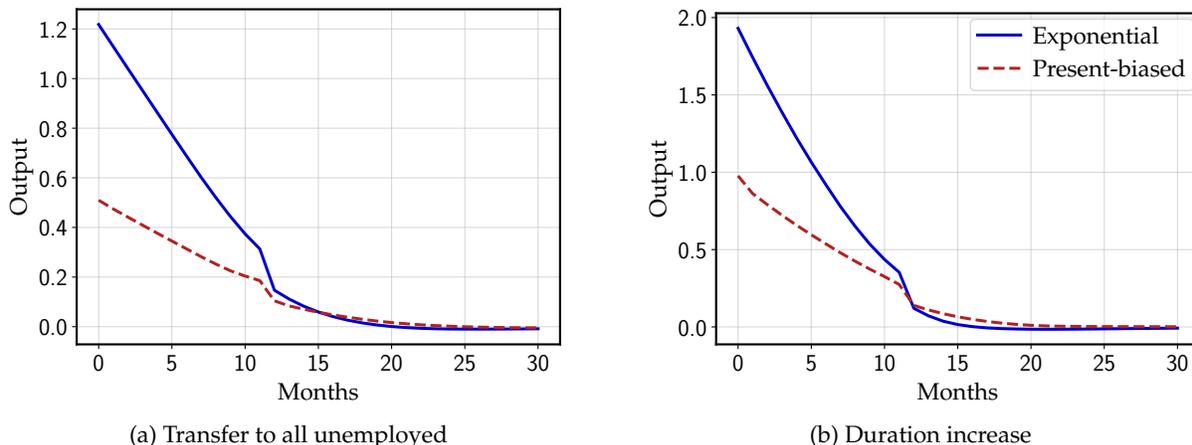


Figure 6: Output response to UI extension

omy. To measure this more precisely, we compute the fiscal multipliers of these transfers by dividing the net present value (NPV) of the output response by the NPV of the additional spending, according to the formula:

$$\text{multiplier} = \frac{\sum_{t=0}^{\infty} (1+r)^{-t} Y_t}{\sum_{t=0}^{11} (1+r)^{-t}},$$

where the denominator already takes into account that total extra spending in each period is normalized to 1. In the case in which all unemployed agents receive transfers, multipliers are 0.77 for the exponential economy and 0.35 with present bias. For the increase in UI duration, they are 1.05 and 0.63, respectively. This is in line with Kekre (2021), who finds multipliers close to 1 for similar increases in duration in a model with exponential households. Multipliers are larger for duration extensions, consistent with the fact that long-run unemployed agents have larger marginal propensities to consume.

The difference between the two cases is driven primarily by precautionary saving. To make this point clear, figure 7 shows the partial equilibrium consumption response split into consumption of employed (left panel) and unemployed (right panel) agents when the UI extension targets all unemployed workers. Surprisingly, the largest difference appears in the consumption of employed agents. In the exponential model, UI extensions are more effective in stimulating output because employed agents reduce precautionary saving in response to them, as it reduces the income risk that comes from unemployment. Note that this does not mean that the consumption of each individual employed agent responds strongly: employed workers are much more numerous, so small individual responses can be sizable at the aggregate level. In the present-biased calibration, on the other hand, precautionary behavior is attenuated, and employed agents respond much less strongly to UI extensions. The same qualitative pattern holds for the consumption of the unemployed, although the difference is much smaller in this case, as most of the effect comes directly from the increase in disposable income.

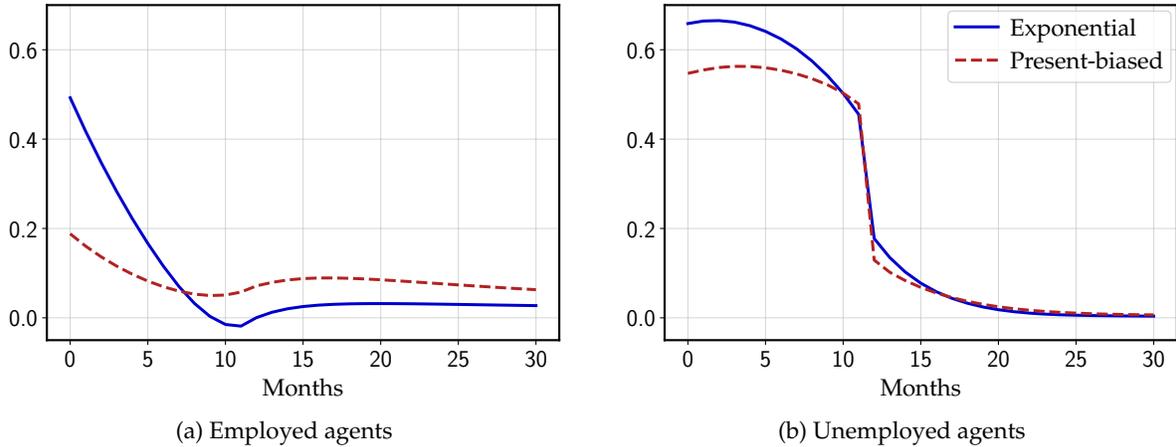


Figure 7: Consumption responses to transfers to all unemployed agents in partial equilibrium

	Increase in the standard deviation of $\log(C_t)$				Annualized r
	All shocks	Only monetary	Only markup	Only discount	
Exponential	17.2%	9.0%	14.1%	19.6%	3.10%
Present-biased	10.1%	6.8%	10.9%	11.2%	3.72%

Table 8: The effects of reduced automatic stabilizers on aggregate consumption volatility.

5.2 UI as an automatic stabilizer

Now, we reduce the generosity of UI and compute the increase in the volatility of aggregate consumption. More specifically, for each exponential and present-biased models, we compute a new steady state with parameters b and \bar{b} reduced by half, and then compute the increase in aggregate consumption volatility as the economy is hit by the estimated shocks above, but around this new steady state. We also reduce safety net transfers T^{sn} by the same factor for three reasons. First, it keeps results comparable to McKay and Reis (2016), who perform this same experiment. Second, these transfers also act as automatic stabilizers. Third, it prevents a situation in which the income of an unemployment agent increases upon UI expiration, which would occur for very low income levels. As these transfers in our model are automatic stabilizers, it is expected that aggregate consumption becomes more volatile when they are reduced.

Table 8 shows the results. The first column shows how the standard deviation of the time series of aggregate consumption changes when automatic stabilizers are reduced in the full model. To understand if and how automatic stabilizers interact differently with each of the three shocks in the model, the subsequent columns show the results of the same experiment, but when only one shock is present at a time. For example, the second column shows the increase in consumption volatility in an economy that is hit only by monetary shocks. The last column shows the equilibrium real interest rate in the reduced-transfers economy.

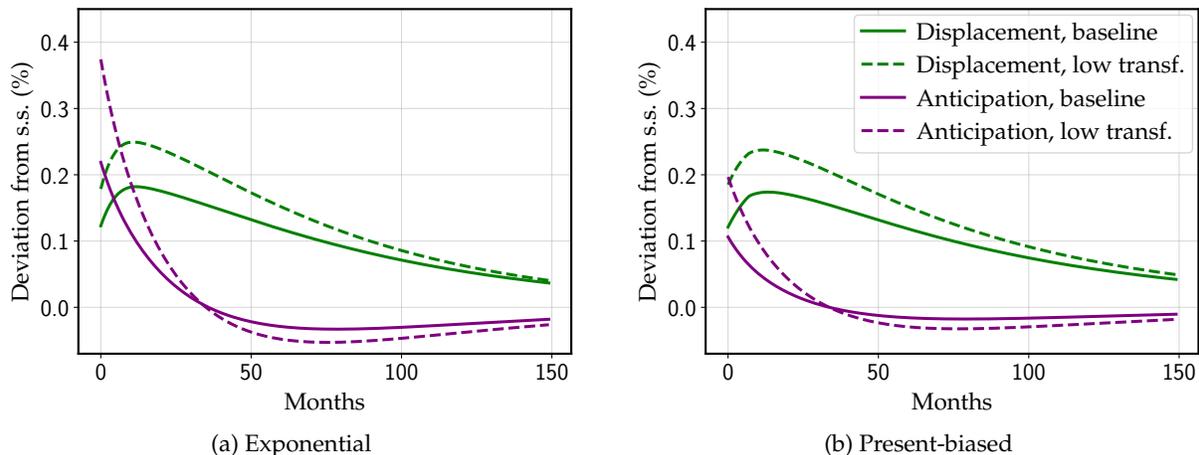


Figure 8: Consumption response to employment shocks in partial equilibrium.

In all cases, UI is less effective in reducing the volatility of aggregate consumption in the present-biased economy. To understand what drives this result, we study the partial equilibrium response of aggregate consumption to a change in employment E_t . Starting from the steady state, we feed into the model an increase in employment given by $d \log E_t = 0.978^t$, with persistence 0.978 corresponding to the autocorrelation of nonfarm employment in the data. This shock is implemented by choosing sequences for separation and job finding probabilities $\{ds_t, df_t\}_{t=0}^{\infty}$, also expressed in deviations from steady-state values, that generate the $d \log E_t$ path subject to the relationship (16).

Figure 8 shows the results. We decompose the consumption response to this increase in employment into two terms. In green, we show the consumption response to changes in transition probabilities s_t and f_t , but keeping fixed at their steady-state levels households' expectations of these probabilities. This term, therefore, captures the effect on consumption of households moving across employment statuses but filtering out any changes in consumption due to changes in precautionary saving in response to the shock. We call this term the *displacement effect*. In purple, we show the response of consumption to *expected* changes in f_t and s_t , but keeping their actual values fixed at their steady-state levels. This term, which we call the *anticipation effect*, captures the change in consumption that is exclusively attributable to the precautionary saving response to shifts in unemployment risk. To a first-order approximation, the total consumption response is the sum of displacement and anticipation effects, as following the one-time shock agents have correct perceptions about the future.²²

Two features emerge from figure 8. First, the anticipation effect is larger in the exponential economy. This is intuitive. The employment shock generates a persistent reduction in unemployment risk. As exponential agents have precautionary behavior more sensitive to such changes in income risk, they display stronger anticipation effects. Second and most importantly, the anticipa-

²²Present-biased agents have incorrect beliefs about their future consumption decisions only, but they correctly perceive movements in aggregate variables.

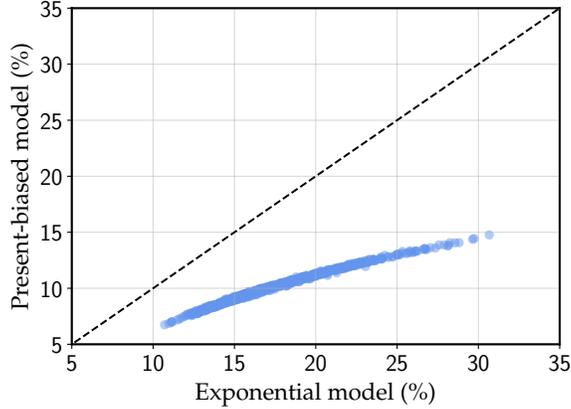


Figure 9: Increase in consumption volatility in the exponential and present-biased models for various parameter values

tion effects are more sensitive to a reduction in automatic stabilizers in the exponential economy: in a low-transfers setting, exponential agents engage in much stronger cyclical precautionary behavior. As a measure of this effect, the anticipation effect at $t = 0$ increases by 0.09 in the present-biased case and by 0.155 in the exponential case – a value more than 60% higher, which explains why automatic stabilizers are more effective in the latter case.

Finally, the scenarios above were constructed using parameter values that correspond to the posterior modes in table 7. As is sometimes the case in these models, there is substantial uncertainty regarding parameter values. To assess whether our result is robust to this type of parameter uncertainty, figure 9 shows the increase in consumption volatility from reducing automatic stabilizers in the exponential and present-biased economies for 1000 parameter values drawn from the joint posterior distribution, as well as a 45-degree line. All points lie below the line, meaning that for all sampled parameter values the exponential economy predicts stronger automatic stabilizers.

6 Conclusion

The business cycles literature has recently incorporated heterogeneous-agent models for generating large marginal propensities to consume, an empirically realistic prediction that is also a key determinant of the transmission of shocks to aggregate consumption. With heterogeneity, however, comes uninsurable risk, which engenders precautionary saving, a less well-studied feature of these models.

In this paper, we study the strength of precautionary saving in heterogeneous-agent models and their implications for stabilization policy. We measure the sensitivity of precautionary savings to changes in income risk by leveraging variation of unemployment insurance replacement rates across U.S. states and over time. Lower replacement rates are associated with more income risk, which the model predicts should lead to more asset accumulation for self-insurance purposes. Empirically, however, we find a small, statistically non-significant effect.

Next, we contrast our empirical findings to the predictions of two HA models: one with standard and another with present-biased preferences. By running the same two-way fixed effects design in model-simulated data, we find large precautionary savings responses in the exponential model that fall outside the estimated confidence intervals. This model is therefore rejected by the data. Moreover, this is robust to many commonly used variations of our baseline model. The present-biased calibration, on the other hand, delivers smaller effects that fall well within the estimated confidence interval.

Last, we look at the implications of weakened precautionary saving for the efficacy of UI as a stabilization tool. We find that UI operates largely through precautionary saving: facing increased UI generosity, households understand they do not need to save as much to self-insure against future unemployment risk. Present bias, by attenuating precautionary saving, reduces the fiscal multipliers of temporary UI extensions by almost 40%. For the same reason, this model predicts UI to be a less powerful automatic stabilizer as well, having a smaller effect in reducing the volatility of business cycles.

Going forward, our work also leaves open questions. First, does the fact that households display weak precautionary behavior change the optimal mix of stabilization tools that should be used during recessions? Should policymakers rely less on UI relative to other tools, or does this mean that UI extensions should be even more aggressive to reach the desired results? Second, we focus solely on the positive implications of UI, leaving normative considerations aside. In a naïve present-biased economy, should UI be more generous because agents do not self-insure as much as in an exponential model, even though it is a less effective automatic stabilizer? We believe these are important questions to be addressed by future research.

References

- Acharya, S. and Dogra, K. (2020). Understanding hank: Insights from a prank. *Econometrica*, 88(3):1113–1158.
- Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.
- Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.
- Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.
- Auclert, A., Rognlie, M., and Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research.
- Auray, S., Fuller, D. L., and Lkhagvasuren, D. (2019). Unemployment insurance take-up rates in an equilibrium search model. *European Economic Review*, 112:1–31.
- Bardóczy, B. (2022). Spousal insurance and the amplification of business cycles.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2018). *"Behavioral Household Finance"*, volume 1. Elsevier.
- Bilbiie, F. O. and Ragot, X. (2021). Optimal monetary policy and liquidity with heterogeneous households. *Review of Economic Dynamics*, 41:71–95.
- Bitler, M. and Hoynes, H. W. (2010). The state of the safety net in the post-welfare reform era. Technical report, National Bureau of Economic Research.
- Blank, R. M. and Card, D. E. (1991). Recent trends in insured and uninsured unemployment: Is there an explanation? *The Quarterly Journal of Economics*, 106(4):1157–1189.
- Blow, L., Browning, M., and Crawford, I. (2021). Non-parametric Analysis of Time-Inconsistent Preferences. *The Review of Economic Studies*, 88(6):2687–2734.
- Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398.
- Campbell, J. Y. (2006). Household finance. *The journal of finance*, 61(4):1553–1604.

- Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics*, 112(1):1–55.
- Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.
- Challe, E. (2020). Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy. *American Economic Journal: Macroeconomics*, 12(2):241–83.
- Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of political Economy*, 116(2):173–234.
- Cho, D. (2020). Unemployment risk, mpc heterogeneity, and business cycles. *Manuscript, University of Melbourne*.
- Chodorow-Reich, G. and Cogley, J. M. (2019). Unemployment insurance and macroeconomic stabilization. *Unemployment Insurance and Macroeconomic Stabilization.* In *Recession Ready*, ed. Heather Boushey, Ryan Nunn, and Jay Shambaugh.
- Currie, J. and Gruber, J. (1996). Saving babies: The efficacy and cost of recent changes in the medicaid eligibility of pregnant women. *Journal of political Economy*, 104(6):1263–1296.
- De Chaisemartin, C. and D’Haultfoeuille, X. (2022). Difference-in-differences estimators of intertemporal treatment effects. Technical report, National Bureau of Economic Research.
- DellaVigna, S. (2018). *“Structural Behavioral Economics”*, volume 1. Elsevier.
- Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American economic review*, 67(3):297–308.
- Engen, E. M. and Gruber, J. (2001). Unemployment insurance and precautionary saving. *Journal of monetary Economics*, 47(3):545–579.
- Erceg, C. J., Jakab, Z., and Lindé, J. (2021). Monetary policy strategies for the european central bank. *Journal of Economic Dynamics and Control*, 132:104211.
- Ericson, K. M. and Laibson, D. (2019). *“Intertemporal Choice”*, volume 2. Elsevier.
- Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54.
- Floden, 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.
- Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327.
- Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

- 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. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-19).
- Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American economic review*, 109(7):2383–2424.
- Gormley, T. A. and Matsa, D. A. (2011). Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies*, 24(8):2781–2821.
- Griliches, Z. (1986). Economic data issues. *Handbook of econometrics*, 3:1465–1514.
- Harris, C. and Laibson, D. (2001). Dynamic choices of hyperbolic consumers. *Econometrica*, 69(4):935–957.
- Hodrick, R. J. and Prescott, E. C. (1997). Postwar us business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, pages 1–16.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.
- Kekre, R. (2021). Unemployment insurance in macroeconomic stabilization. Technical report, National Bureau of Economic Research.
- Kleibergen, F. and Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of econometrics*, 133(1):97–126.
- Kuka, E. (2020). Quantifying the benefits of social insurance: unemployment insurance and health. *Review of Economics and Statistics*, 102(3):490–505.
- Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting*. *The Quarterly Journal of Economics*, 112(2):443–478.
- Laibson, D., Maxted, P., and Moll, B. (2021). Present bias amplifies the household balance-sheet channels of macroeconomic policy. Technical report, National Bureau of Economic Research.
- Maxted, P. (2020). Present bias in consumption-saving models: A tractable continuous-time approach. Technical report, Mimeo.
- McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics*, 88:1–14.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.

- McKay, A. and Reis, R. (2016). The role of automatic stabilizers in the us business cycle. *Econometrica*, 84(1):141–194.
- McKay, A. and Reis, R. (2021). Optimal automatic stabilizers. *The Review of Economic Studies*, 88(5):2375–2406.
- O’Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.
- Olea, J. L. M. and Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3):358–369.
- Ravn, M. O. and Sterk, V. (2021). Macroeconomic fluctuations with hank & sam: An analytical approach. *Journal of the European Economic Association*, 19(2):1162–1202.
- Schaab, A. (2020). Micro and macro uncertainty. *Available at SSRN 4099000*.
- Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148.
- Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier.
- Vimercati, R. B., Eichenbaum, M. S., and Guerreiro, J. (2021). Fiscal policy at the zero lower bound without rational expectations. Technical report, National Bureau of Economic Research.
- Werning, I. (2015). Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research.

A The Zero Liquidity Model

In this appendix, we present the zero liquidity model (Werning 2015, Challe 2020, McKay and Reis 2021) from section 3. The key assumption here is that agents cannot borrow ($\underline{a} = 0$) and the supply of liquid assets converges to zero ($B \rightarrow 0$). We also there is no UI expiration ($N^{ui} = 0$), no sources of income other than UI for the unemployed ($y^{priv} = 0$), no labor productivity heterogeneity ($z_{it} = 0$), and no heterogeneity in the degree of present bias ($\beta^L = \beta^H \equiv \beta$). The later assumptions only deliver tractability and do not change results qualitatively. We omit time subscripts for now, as here we only compare steady states.

As an immediate consequence of the assumptions above, all employed agents have income $y_e = 1 - \tau_t$, while unemployed agents have income $y_u = (1 - \tau_t)b$. For any given interest rate r , we define the value function on employed and unemployed exponential agents, denoted by V_e^{exp} and V_u^{exp} , respectively, as the solutions of the following Bellman equation:

$$V_e^{exp}(a) = \max u(c) + \delta [(1 - s)V_e^{exp}(a') + sV_u^{exp}(a')] \quad (17)$$

$$\text{s.t. } c + a' = (1 + r)a + y_e, a' \geq 0$$

$$V_u^{exp}(a) = \max u(c) + \delta [fV_e^{exp}(a') + (1 - f)V_u^{exp}(a')] \quad (18)$$

$$\text{s.t. } c + a' = (1 + r)a + y_e, a' \geq 0.$$

The naïve present-biased agent then solves

$$V_e(a) = \max u(c) + \beta\delta [(1 - s)V_e^{exp}(a') + sV_u^{exp}(a')] \quad (19)$$

$$V_u(a) = \max u(c) + \beta\delta [fV_e^{exp}(a') + (1 - f)V_u^{exp}(a')],$$

subject to the same constraints as above. Note the two differences between the optimization of present-biased and exponential agents: the extra discount factor β and the fact that naïve present-biased agents believe they will behave like exponential discounters in the future. Clearly, if $\beta = 1$, then $V_e(\cdot) = V_e^{exp}(\cdot)$ and $V_u(\cdot) = V_u^{exp}(\cdot)$. Moreover, let $c_k(a)$ and $c_k^{exp}(a)$ be the optimal consumption policy function for present-biased and exponential agents with employment status k , respectively. We have the following lemma:

Lemma 1. *The unconstrained present-biased agent consumes according to the Euler equation*

$$u'(c_e(a)) = (1 + r)\beta\delta [su'(c_e^{exp}(a')) + (1 - s)u'(c_u^{exp}(a'))] \quad (20)$$

Proof. Combine the first-order condition of the maximization problem in (19), together with the first order and envelope conditions of (17) and (18). \square

The equation above is similar to a standard Euler equation, except for the fact that the present-biased agents have incorrect expectations about her future consumption. The zero liquidity equi-

librium is then characterized by the interest rate r such that the employed agents with zero assets choose to consume all their income, i.e., $c_e(0) = y_e$, but the Euler equation still holds with equality. Now we are ready to prove proposition 1.

A.1 Proof of proposition 1

Assume that the unemployed exponential agent with no assets is also constrained, i.e., $c_u^{exp}(0) = y_u$. This assumption always holds when $\beta = 1$. Intuitively, the reason is that employed agents have incentives to save to self-insure against unemployment risk in this model, so unemployed agents can consume out of their savings. If the interest rate is low enough so that the employed agent optimally decides not to save, then it is also optimal for the unemployed agent not to save. To the extent that the equilibrium interest rate varies continuously with β , this also holds in the present-biased case for β close enough to 1. We verify numerically that this is indeed the case for our parameterization, as can be done easily by solving equation (18). Therefore, the Euler equation (20) can be written as

$$1 + r = \frac{1}{\beta\delta \left[s \left(\frac{c_e^{exp}(0)}{y_e} \right)^{-\sigma} + (1-s)b^{-\sigma} \right]}, \quad (21)$$

where σ is the coefficient of relative risk aversion.

Now we can differentiate (21) with respect to b to obtain

$$\frac{dr}{db} = \frac{1}{\beta\delta} \frac{1}{\left[s \left(\frac{c_e^{exp}(0)}{y_e} \right)^{-\sigma} + (1-s)b^{-\sigma} \right]^2} \left[\sigma s \left(\frac{c_e^{exp}(0)}{y_e} \right)^{-\sigma} \frac{d}{db} \left(\frac{c_e^{exp}(0)}{y_e} \right) + \sigma(1-s)b^{-\sigma-1} \right]$$

Using (21) again, we obtain

$$\frac{dr}{db} = \beta\delta(1+r)^2 \left[\sigma s \left(\frac{c_e^{exp}(0)}{y_e} \right)^{-\sigma} \frac{d}{db} \left(\frac{c_e^{exp}(0)}{y_e} \right) + \sigma(1-s)b^{-\sigma-1} \right],$$

which is exactly equation (8).

B Derivation of the Phillips Curve

In this appendix, we provide a derivation of the Phillips Curve (15). First, note that prices are flexible and firms' nominal marginal cost is proportional to the nominal wage from (14), so price and wage inflation are the same in our model and can be used interchangeably.

There is a continuum of unions indexed by j , each one operating in a monopolistically competitive setting. All households supply labor to each union. As is standard in models with nominal rigidity, each union j commits to supplying whatever amount of labor N_{jt} is demanded at its prevailing nominal wage W_{jt} . Firms' total labor input is given by (13), which gives rise to the standard

constant-elasticity demand function for the labor supplied by union j :

$$N_{jt} = \left(\frac{W_{jt}}{\bar{W}_t} \right)^{-\nu} N_t. \quad (22)$$

Moreover, with a slight abuse of notation, we can write

$$N_{jt} = E_t n_{jt},$$

where n_{jt} is the number of hours employed agents supply to union j . Now define

$$n_t = \frac{N_t}{E_t}. \quad (23)$$

By combining (22) and (23), we obtain

$$\int n_{jt} dj = n_t \int \left(\frac{W_{jt}}{\bar{W}_t} \right)^{-\nu} dj.$$

The wage dispersion term on the right-hand side of the above equation can be shown to be second order (Galí 2015). Therefore, we have

$$\int n_{jt} dj = n_t + \text{second order terms}. \quad (24)$$

Unions maximize the objective function (11). Thus, when able to reset wages, unions choose the wage W_t^* that solves

$$\max_{W_t^*} \mathbb{E}_t \sum_{k=0}^{\infty} (\delta\theta)^k \frac{1}{E_{t+k}} \int [u(c_{i,t+k}) - v(n_{t+k})] 1_{i,t+k}^e di.$$

Recall that θ is the probability that a union cannot reset its wage in a given period and that all households are assumed to work the same number of hours n_t . Moreover, note that the optimal reset wage W_t^* is the same for all adjusting unions. Unions that are not able to optimally reset their wages can partially index them to past inflation as in (12).

The first order condition is

$$\mathbb{E}_t \sum_{k=0}^{\infty} (\delta\theta)^k \frac{1}{E_{t+k}} \int \left[u'(c_{i,t+k}) \frac{\partial c_{i,t+k}}{\partial W_{jt}^*} - v'(n_{t+k}) \frac{\partial n_{t+k}}{\partial W_{jt}^*} \right] 1_{i,t+k}^e di = 0 \quad (25)$$

As in Auclert et al. (2018), the envelope theorem gives us

$$\frac{\partial c_{i,t+k}}{\partial W_{jt}^*} = \frac{\partial y_{i,t+k}}{\partial W_{jt}^*}.$$

Due to indexation wage indexation, the real income of employed workers in period $t+k$ is given

by

$$\begin{aligned} y_{i,t+k} &= (1 - \tau_{t+k}) \frac{\Pi_{t-1,t+k-1}^l W_{jt}^*}{P_{t+k}} n_{t+k} z_{i,t+k}, \\ &= (1 - \tau_{t+k}) z_{i,t+k} \frac{N_{t+k}}{E_{t+k}} \frac{\Pi_{t-1,t+k-1}^l W_{jt}^*}{P_{t+k}} \left(\frac{\Pi_{t-1,t+k-1}^l W_{jt}^*}{W_{t+k}} \right)^{-\nu} \end{aligned}$$

where $\Pi_{t-1,t+k-1} = P_{t+k-1}/P_{t-1}$. The second line uses equations (22) and (23). Therefore, for employed workers we have

$$\begin{aligned} \frac{\partial y_{i,t+k}}{\partial W_{jt}^*} &= (1 - \nu)(1 - \tau_{t+k}) z_{i,t+k} n_{t+k} \frac{\left(\frac{\Pi_{t-1,t+k-1}^l}{P_{t+k}} \right)^{1-\nu}}{\left(\frac{1}{W_{t+k}} \right)^{-\nu}} \left(W_{jt}^* \right)^{-\nu} \\ &= (1 - \nu)(1 - \tau_{t+k}) z_{i,t+k} n_{t+k} \frac{W_{t+k}}{P_{t+k}} \left(\frac{\Pi_{t-1,t+k-1}^l}{W_{t+k}} \right)^{1-\nu} \left(W_{jt}^* \right)^{-\nu} \\ &= (1 - \nu)(1 - \tau_{t+k}) z_{i,t+k} n_{t+k} w_{t+k} \left(\frac{\Pi_{t-1,t+k-1}^l}{W_{t+k}} \right)^{1-\nu} \left(W_{jt}^* \right)^{-\nu}. \end{aligned}$$

Regarding hours worked, we can use (24) to obtain

$$\frac{\partial n_{t+k}}{\partial W_{jt}^*} = -\nu n_{t+k} \left(\frac{\Pi_{t-1,t+k-1}^l}{W_{t+k}} \right)^{-\nu} \left(W_{jt}^* \right)^{-\nu-1}.$$

After some simplification, the first order condition (25) then becomes

$$\mathbb{E}_t \sum_{k=0}^{\infty} (\delta\theta)^k \left(\frac{\Pi_{t-1,t+k-1}^l}{W_{t+k}} \right)^{-\nu} \left\{ \frac{\Pi_{t-1,t+k-1}^l}{W_{t+k}} W_{jt}^* \mathcal{U}_{t+k} - \frac{\nu}{\nu-1} \mathcal{V}_{t+k} \right\} = 0, \quad (26)$$

where

$$\begin{aligned} \mathcal{U}_t &= \frac{1}{E_t} \int [u'(c_{it}) (1 - \tau_t) z_{it} w_t] 1_{it}^e di \\ \mathcal{V}_t &= v'(n_t). \end{aligned}$$

From (26), we see that in a zero-inflation steady state

$$\mathcal{U} = \frac{\nu}{\nu-1} \mathcal{V}.$$

Now log linearizing (26) around the zero-inflation steady state, we get

$$\mathbb{E}_t \sum_{k=0}^{\infty} (\delta\theta)^k \left(1 + (1 - \nu) \iota \pi_{t-1,t+k-1} - (1 - \nu) \widehat{W}_{t+k} \right) \left\{ \widehat{W}_t^* + \iota \pi_{t-1,t+k-1} - \widehat{W}_{t+k} + \widehat{\mathcal{U}}_{t+k} - \widehat{\mathcal{V}}_{t+k} \right\} = 0,$$

where

$$\pi_{t-1,t+k-1} = \log P_{t+k-1} - \log P_{t-1}$$

and hats denote log deviations from the steady state. Discarding second-order terms, we obtain

$$\widehat{W}_t^* = (1 - \delta\theta)\mathbb{E}_t \sum_{k=0}^{\infty} (\delta\theta)^k \left(\widehat{W}_{t+k} + \widehat{V}_{t+k} - \widehat{U}_{t+k} - \iota\pi_{t-1,t+k-1} \right) = 0,$$

which can be written recursively as

$$\widehat{W}_t^* = (1 - \delta\theta) \left(\widehat{W}_t + \widehat{V}_t - \widehat{U}_t \right) + \delta\theta\mathbb{E}_t \left(\widehat{W}_{t+1}^* - \iota\pi_t \right). \quad (27)$$

The last equation we need to obtain the Phillips curve is the law of motion for the aggregate nominal wage, given by

$$W_t^{1-\nu} = (1 - \theta) (W_t^*)^{1-\nu} + \theta (W_{t-1}\Pi_{t-1}^t)^{1-\nu},$$

which, linearized around a zero-inflation steady state, becomes

$$\pi_t = (1 - \theta)(\widehat{W}_t^* - \widehat{W}_{t-1}) + \theta\iota\pi_{t-1} \quad (28)$$

Combining (27) and (28), we obtain the Phillips curve (15).

C Impulse Response Functions

This appendix contains impulse response functions of the estimated dynamic model from section 4. The parameter values used here correspond to the posterior modes from table 7. For each shock in the model, we compute the impulse response functions of aggregate consumption, unemployment, nominal interest rates, and inflation. Aggregate consumption is expressed in percent deviations from the steady state, while the other variables are expressed in percentage point deviations. The nominal interest rate and the inflation rate are annualized, i.e., multiplied by 12. Results are shown in figures 10, 11, and 12.

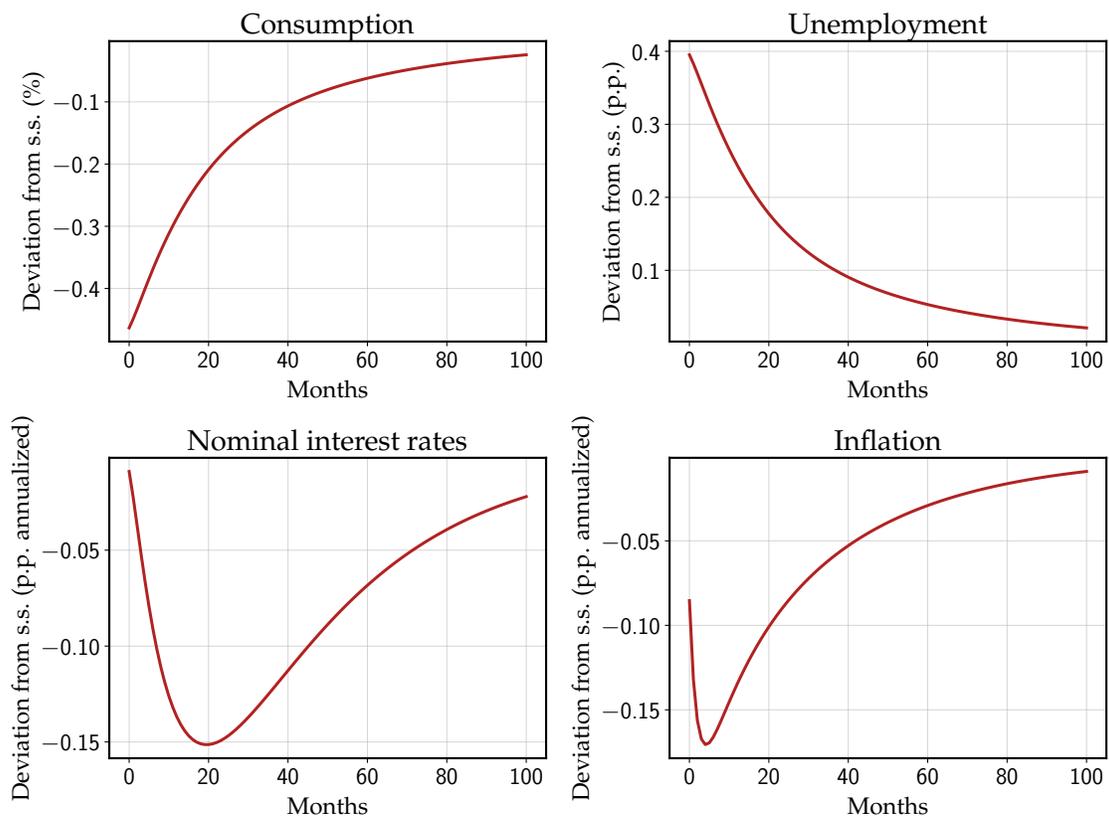


Figure 10: IRF to discount rate shock ϵ_t^d

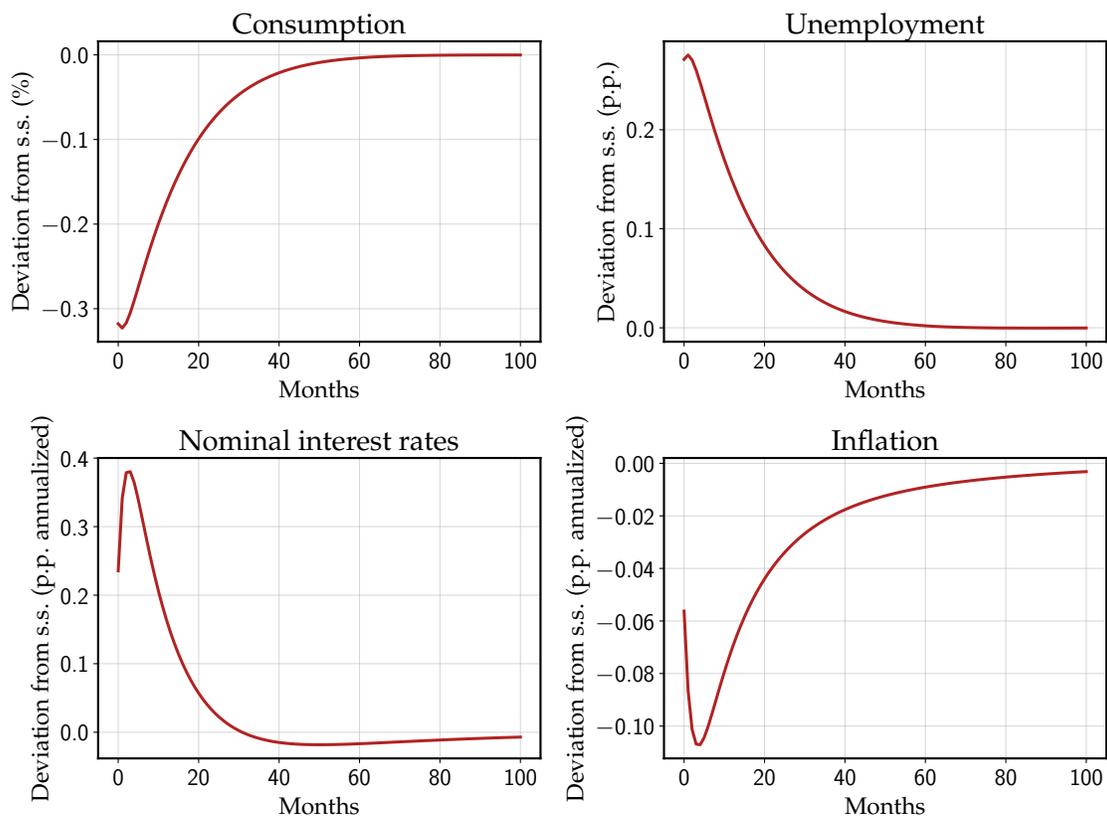


Figure 11: IRF to monetary shock ϵ_t^i

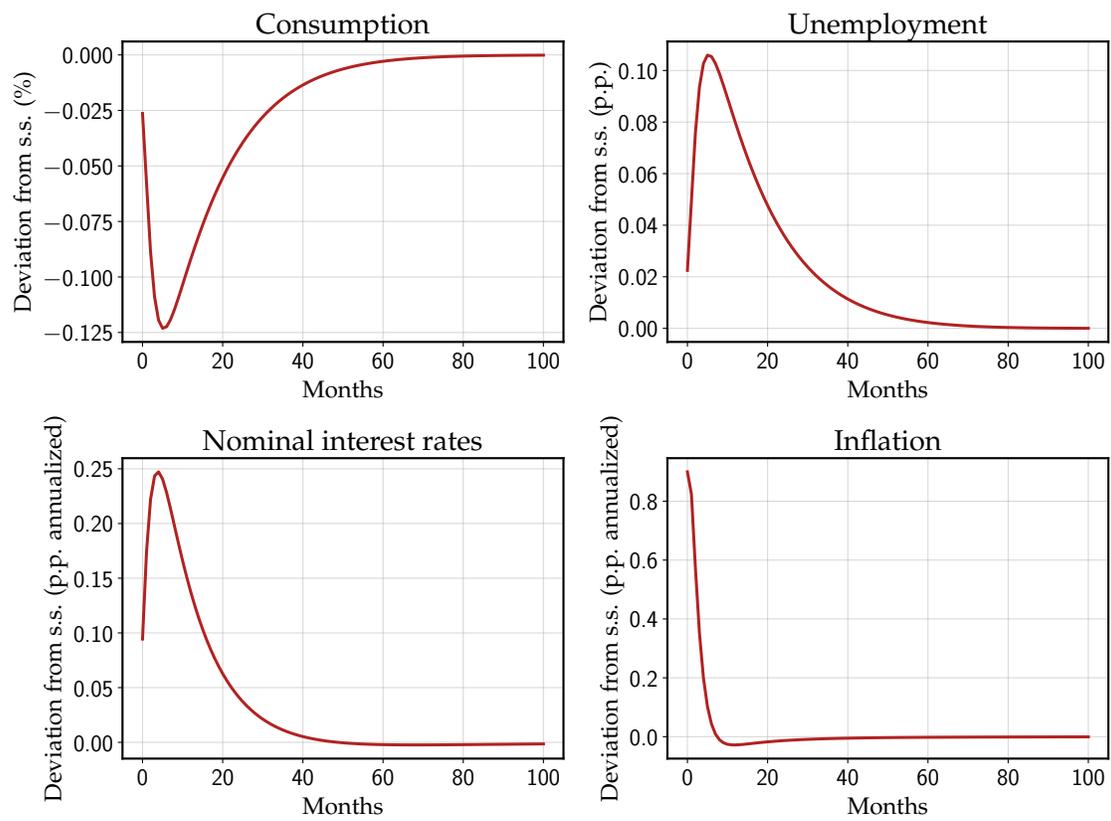


Figure 12: IRF to wage mark-up shock ε_t^w

D Data Appendix

This appendix describes how we construct our data, starting from the raw SIPP files up to the datasets used for regression analyses. The SIPP is a large public dataset. A detailed description of cleaning and manipulating the SIPP data based on its documentation is infeasible, and hence we refer the interested reader to the survey’s documentation²³ for further details.

D.1 Survey of Income and Program Participation (SIPP)

We start with the NBER-provided SIPP files.²⁴²⁵²⁶ The SIPP is a rotating panel that interviews respondents over waves. Each panel is composed of a given number of waves, which may vary by panel. Each wave corresponds roughly to four months, with respondents interviewed at some point over the wave. Some information is collected each wave (e.g. income, marital status, household affiliation, etc), and is recorded in the Core files. Additionally, there are supplements to the core files – the topical files – which contain variables related to specific topics. Asset holdings in the SIPP are not collected in every wave and are instead recorded in topical supplements conducted in a few waves over the life of the panel. As a rough rule of thumb, every (calendar) year or so in the panel, a wealth topical module is available. The table below provides which waves have asset holdings data for each of the panels we use. We download core files for all panels starting from 1984 up to 2008, excluding 1987 and 1988, as these panels have no asset data. We also download all topical files with asset variables for each of these panels. As the survey went through a substantial redesign following the 2008 panel, we do not use subsequent panels in our analyses.

Unfortunately, variable names and codings are not harmonized across SIPP panels (crosswalks have become available only following the 2008 panel). We select variables of interest and harmonize their names across all core and topical files we use. To do so, we made extensive use of the documentation provided in the website of the US Census Bureau. For example, if a variable changes names between different panels, we ensure it has a single name across panels in our data.

²³<https://www.census.gov/programs-surveys/sipp/tech-documentation.html>

²⁴These are available with Stata dictionaries from <https://www.nber.org/research/data/survey-income-and-program-participation-sipp>

²⁵To our knowledge, the SIPP is the only public dataset that allows econometricians to observe both household balance sheets and state of residence. The restricted files for the Survey of Consumer Finances (housed in the Federal Reserve Board) allow researchers to jointly observe wealth and state as well. However, access to the restricted data is extremely limited.

²⁶An alternative commonly employed in the Household Finance literature is to employ wealth registries from countries that at some point (or to this day) administer wealth taxes – Denmark, Norway, and Sweden. We encountered two obstacles that pushed us away from pursuing such data. First, there is an external validity concern. Households in these countries are wealthier, face markedly different levels of unemployment risk, and participate in a more generous welfare system. Second, though policy variation in these countries has been used to detect effects of UI generosity on search behavior among unemployed individuals (Falch 2015, Kolsrud et al 2018) or MPCs out of transfers (Landais and Spinnewijn 2020), we do not believe that there would be enough variation to detect effects on liquid asset buffers. Either geographical variation was meager, or any reforms affected elements of the schedule that had limited impact in the overall generosity of UI policy (eg. duration-dependent caps, minimum benefits).

SIPP Panels Used		
SIPP Panel	# Waves	Wealth Waves
1984	9	4, 7
1985	8	3, 7
1986	7	4, 7
1987	7	4
1990	8	4
1991	8	7
1992	9	4
1993	9	7
1996	12	3, 6, 9, 12
2001	9	3, 6, 9
2004	12	3, 6
2008	16	4, 7, 10

Table 9: SIPP Panels

Relatedly, if a categorical variable changes values across consecutive years, we recode it (using our best judgment in selected cases) to provide harmonized categories across panels.

Our ultimate goal is to have a repeated cross-sectional²⁷ dataset in which the unit of observation is a household \times panel \times wave. This is dictated by asset data availability: such information is *only available at the household-level* and the respondent is asked about his asset holdings as of the interview month, so any other months within the wave are populated with the same information. We focus on household heads following Engen and Gruber (2001), for whom we will collect individual-level covariates that will be used as controls. We first clean the asset variables, which constitute our outcomes. Then, we clean core files up to the last available wealth wave for each panel. We use the core waves to select our sample and construct covariates (e.g. income, replacement rates, sex, etc). Then, we merge household-level core files with the topical files. Finally, we only keep waves for which we observe assets. We detail the procedure more carefully below. Replication files with extensive comments are available upon request.

First, we clean the wealth topical modules for each panel. Asset holdings data is available at the household level. We use data constructed by the SIPP staff whenever possible (variables starting with “THH” in the technical documentation). Following recommendations from the technical documentation, we disconsider asset data from respondents with improper interview statuses (e.g. non-interviewed respondents, for which any information has been imputed). More precisely, we

²⁷Many researchers studying unemployment spells using SIPP data have conducted longitudinal analyses. Unfortunately, given we focus on wealth, any longitudinal analyses quickly stumble on data limitations. Wealth is only surveyed in a few waves, with the 1987-1993 panels not even having more than one wealth wave. For the earlier panels, the amount of households for which we observe wealth in both wealth waves is small. The longitudinal component of the survey improves after the 1996 redesign, but, even then, it’s not so common to observe households for which we sample wealth over two consecutive waves. Restricting ourselves to a longitudinal analysis would substantially reduce our sample size and thus our power to detect any reasonable effects.

turn all asset data we use to missing if the respondent is marked as having an improper interview status (e.g. EPPINTVW of 3 and 4 in post-1996 panels) or is marked as having all its asset data imputed, despite having valid interview status.

Additionally, we avoid imputed asset values whenever possible to avoid introducing additional measurement error arising from the imputation procedure. The SIPP-provided household-level asset totals are not imputation free, but the SIPP provides imputation markers, which we use to recover imputation-free asset holdings. This often requires some judgement, as the documentation does not explicitly describe formulas used for aggregation. However, the descriptions are informative, so that one can reasonably back out which more granular variables are summed to arrive at a particular aggregate. We do not find much disparity in the distributions of these aggregated variables across survey panels after our cleaning procedure, bolstering our belief that our cleaning procedure is reasonable.

If a variable is an aggregate of other more granular variables (e.g. net financial assets), we set missing values or imputed values to zero when taking sums, except when all sum constituents are missing or imputed, in which case the sum is set to missing. There is a trade-off in increasing sample size and thus power versus mitigating measurement error. We opt to privilege the former, especially as the measurement error would be affecting the dependent variable and therefore would contribute to inflated standard errors rather than biased coefficients (of course, under classical measurement error assumptions). In unreported results, we repeat our empirical exercises, but setting sums to missing if any of the underlying components is missing. Our results remain very similar. See Curtis et al (1989) and Hoynes et al (1995) for early criticisms of the SIPP imputation methodology. To allow for cross-year comparisons of asset totals, we transform all monetary variables are transformed into real 2011:6 USD using the PCE price index (from FRED).

Second, we clean the core files. As we will only use information up to the last wealth wave in a panel, we only clean core waves up to then (e.g. for 1992, we only clean core waves 1-4, but, for 1996, we clean all core waves). Following the technical documentation, we do not consider data from respondents with improper interview statuses (we treat such data as missing). We use the core waves to perform sample restrictions and construct our main regressors, namely UI generosity variables using data on earnings histories.

We impose the following sample restrictions, following Engen and Gruber (2001). First, we only consider single-family households. Second, we restrict attention to household heads (assets are observed at the household level) aged 25-64, at least for two reasons: (i) outside this range, households are in different stages of wealth accumulation over their life cycles (education or retirement); (ii) this demographic group is predominantly in the labor force and exposed to unemployment risk. Third, we require no changes in marital status in the waves preceding (and including) the wave of the asset observation.²⁸ This mitigates changes in the availability of spousal insurance or in household wealth. Fourth, we require some wage or salary earnings in the waves preceding (and including) the wave of the asset observation. This ensures we observe households with stable

²⁸For example, if we observe asset data in wave 9 of the 1996 panel for a household head, we require that his or her marital status has not changed since we first observe such individual in the panel.

(labor) earnings histories.

The last sample restriction merits discussion. A large literature in Public Finance studies behavior over unemployment spells using SIPP data (e.g. Chetty 2008, Kuka 2020). However, our theoretical model makes predictions for the behavior of *employed households* who might be at risk of losing their jobs. We ask how much would these households adjust their liquid asset holdings following changes in the consumption risk they face were they to become unemployed. Therefore, isolating a sample of households that are (likely to be) employed is of the essence. Moreover, we condition on earnings rather than formal employment since earnings ultimately matter for wealth accumulation. Fifth and last, we exclude observations for which we do not observe state codes, since we will simulate individual-level replacement rates using information state-level information on UI schedules.²⁹

Beyond sample selection, we collect information on age, sex, race, marital status, years of education, number of children, and weekly wages for the household head. We also collect detailed information on social insurance reciprocity. As with wealth variables, the SIPP has rich information on social insurance payments, with over 40 different kinds of transfer payments in some panels. We ignore imputed values and group different transfers into broad categories: UI, non-UI social security, pension retirement (from government or private institutions), private insurance. Consistent with us selecting individuals with stable employment histories and of prime-age demographic groups, we see very low levels of social security reciprocity in our sample.

We then use the information on previous earnings to simulate UI payments that households would be eligible for if they were to become unemployed. For this exercise, we leverage the UI calculator as well as replication codes available from Kuka (2020). Her UI calculator itself also builds upon previous calculators that have been routinely used in the Public Finance literature. We provide a more comprehensive description of the UI calculator below. For now, it suffices to understand that the calculator inputs measures of previous earnings, number of children, and state of residence, and outputs the weekly benefit allowance (WBA) that the household head would receive if he or she became unemployed, regardless of eligibility. A different calculator takes in the same inputs and outputs whether the household head would be eligible to receive UI. To obtain replacement rates, we divide the WBA by a weekly measure of previous earnings frequently used in state UI schedules to determine WBAs. Again, all monetary variables are transformed into real 2011:6 USD using the PCE price index.

Having cleaned both core and topical files, we merge them and keep only the waves in which we observe assets. We additionally merge supporting datasets containing state-time level controls from Kuka (2020),³⁰ as well as state-level unemployment rates, labor force participation rates, and resident population from GeofRED. We then append all panels together, and collapse the data to

²⁹Fortunately, there are few such cases, as the primary motivation for masking individual state markers are small cell concerns. Moreover, starting 2004, there are no instances of masking.

³⁰These data overall contain information on state-level generosity of other social insurance programs, such as total welfare spending; indicators for adoption of the large welfare reforms in the 1990s; as well as state-level UI expenditure data, such as benefits paid (broken into state, federal and emergency payments) and net reserves. All monetary values are converted to 2011:6 USD using the PCE index.

Selection criteria	Number of obs.	Level of obs.
All wealth waves	5281999	Person-wave-month
Single family HH	4552157	Person-wave-month
Aged 25-64	2318083	Person-wave-month
Ref. person	1344691	Person/HH-wave-month
No change in marital status	1298346	Person/HH-wave-month
Some wages	744897	Person/HH-wave-month
Collapsing	189012	Person/HH-wave
Valid wealth	188930	Person/HH-wave
Grouped + survey adj.	188929	Person/HH-wave
Baseline regression	178121	Person/HH-wave

Table 10: Sample Accounting

Notes: This table performs sample accounting. The first row counts all observations in our merged core-topical files for which wealth is observed. The second row restricts to single-family households. The third row imposes our age restriction. The fourth row restricts to household heads (reference person, in SIPP technical language). The fifth row restricts to heads with no changes in marital status up to the wealth observation, as described in the main text. The sixth row imposes the restriction of having some labor earnings (wages) in each wave up to the wave of the wealth observation. The seventh row collapses from monthly to wave-level observations for each household head. The eighth row restricts to households heads with valid wealth observations, as flagged by the SIPP technical documentation. The ninth row drops any remaining individuals flagged with improper survey codes. The last row indicates our baseline regression sample with all non-missing variables included in the TWFE regression Column 6 of Table 2.

the household-panel-wave level.³¹ Table 10 performs a sample accounting exercise, illustrating how many observations we lose with each of our sample restrictions and data cleaning. Our final dataset contains around 180,000 household heads spread over 51 states and 22 years.

D.2 UI Calculators and Simulated Policy Instruments

Our key regressor is a measure of UI generosity. Although the theoretical public finance and macroeconomics literature has normally modeled UI conveniently as a replacement rate b over previous employed earnings, reality is tremendously more complex. On one hand, this generates substantial cross-state and cross-time variation that we can leverage for our empirical exercises. On the other hand, mapping data to theoretical objects necessitates judgment calls. This is also true for most social insurance programs, for which the gap between theoretical modeling and actual schedules can be enormous.

The solution in the Public Finance literature since the seminal work of Currie and Gruber (1996) has been to construct calculators that simulate policies given the information that states would have used to determine transfer payments, and then compute a summary measure of the policy’s generosity. In our context, we employ UI calculators that have been curated by previous UI researchers over the past 25 years, and consider, for the most part, the replacement rate to summarize the generosity of the policy. Our UI calculator comes from Kuka (2020), who builds on previous

³¹Recall each wave has four months, and the wealth variables are available at the wave level. For any categorical variables, we take the value observed in the last month of the wave, a standard approach to deal with seam bias in the SIPP. For monetary values (e.g. income), we take averages over the wave.

calculators in the literature. The calculator is simply a script that takes in information about previous earnings, state of residence, time period, etc, and outputs either an eligibility dummy or the WBA that an individual would receive if they were to lose their jobs in that particular state-time, given previous earnings and other covariates that may enter the policy schedule.

The calculator contains information on the particular formulae each state uses at a given point in time to compute WBAs or eligibility. These formulas were hand-collected by previous researchers from statutory files. A quick glance reveals tremendous heterogeneity in what enters schedules across states and over time. For example, some states consider the highest quarterly earnings over a given time period as their measure of previous earnings. Some consider a measure of annual wages. Some have slopes, kinks, or intercepts that are dependent on the number of children. Moreover, for WBA, there are statutory minimum and maximum payments, which also generate kinks in the schedule. Within some states, the parameters of the schedule frequently change over time.

Our calculator is rich enough that we observe the schedule for all US states at a semesterly frequency. The calculator does not make any simplifying assumptions, and constructs the information necessary to compute WBAs and eligibility given the schedule in each state-year pair. We transform WBA into replacement rates by dividing it by the highest quarterly earnings prior to the assumed time of job loss (converted to a weekly amount). This is the measure of previous earnings most commonly used across states and over time. Constructing the replacement rates with the second most frequently used measure of pre-job loss earnings delivers replacement rates that are extremely correlated with our main measure.

When constructing our main dataset, we simulate eligibility and replacement rates at the individual-time level. Namely, we ask “how much would this household head make in UI payments if he or she were to lose his job in this month, given the state of residence and previous earnings history?”. This is what our calculator computes for each household head in our sample. Although this is a very granular measure of policy generosity, since at least Chetty (2008), UI researchers have recognized that individual-level simulated policies can be noisy, primarily due to measurement error in inputs. Moreover, from an identification perspective, there are several concerns about using individual replacement rates as a measure of policy generosity.

First, individuals may select into a particular state because of the policy generosity. If there are unobserved factors that influence the outcome of interest and the decision to live in a state due to policy generosity (“spatial selection”), a regression of the outcome on individual-level simulated policy variables will provide confounded estimates. More broadly, borrowing from DAG logic, if there are any unobserved covariates that influence an individual-level input of the policy schedule and the outcome of interest, our estimates will also be confounded. At the core of the issue is that individual-level simulated policy measures leverage individual-level information that may contaminate the “true” generosity of the policy.

We follow Currie and Gruber (1996) and a large ensuing literature that overcomes this issue by constructing *simulated policy instruments*. To obtain a measure of policy generosity that is un-

confounded, we simulate the policy using a *counterfactual sample* of individuals. To be concrete, if we want to obtain a measure of how generous UI is in Alabama in 1987, we can construct a set of individuals sampled from different state-year pairs, assume that they lose their jobs, compute how much UI they would make if they were in Alabama in 1987, and then take a moment of the distribution of simulated policies as a clean measure of policy generosity, e.g. the mean or median replacement rate in this counterfactual sample.

This is precisely what we do with our data. We start from our sample of household heads, but do not condition on having valid wealth observations. We then strip these observations from state and time variables. This is a national, cross-time sample of individuals. We assume that all such individuals lost their jobs in a particular state-time cell (e.g. Alabama in 1987), and compute their eligibility and their WBAs using the UI schedule from that particular state-time cell. We then compute the mean and median replacement rates. These moments serve as our simulated policy instruments for that particular state-time cell. We repeat this for all states and time periods (semester-years) from 1984:H1 to 2011:H2.

The simulated policy instrument solves each concern raised by using individual-level simulated policy variables. First, we use a national counterfactual sample, thus circumventing the issue of spatial selection. Second, by pooling multiple individuals and considering sample moments, unobservables that might bias the measure of policy generosity wash out. Third, by taking moments or quantiles of the distribution of simulated replacement rates, we considerably reduce noise, so that our simulated policy instrument aids with measurement error concerns in the spirit of Griliches (1986).

In practice, one can compute mean and median replacement rates over more narrowly defined cells than state-time cells. For instance, many states condition UI generosity on the number of children. This is actual variation in the generosity of the schedule that the researcher may want to explore. In that case, one can take means or medians by state-time-number of kids cells. Relatedly, UI is by design a progressive policy. One can capture this progressivity by taking means or median over cells defined by state, time, and quantiles of previous earnings. The only care is that there should be sufficiently many observations in each cell to ensure that individual-level unobservables do not dominate the moments or percentiles of interest.

Following this logic, we compute simulated instruments at different levels of granularity. Throughout our regression exercises, unless explicitly noted, we take median replacement rates at the state-time-number of kids level as our measure of policy generosity. First, it strikes a good balance between the granularity of the instrument (there is substantial variation by number of kids)³² and average cell size. Second, taking medians generates a measure of centrality of generosity without risking having outliers unduly influence the measure, as it would be the case if we had taken means. Having obtained our simulated instruments, we merge them with our cleaned SIPP files, which concludes the construction of the main dataset.

Appendix E presents the results of the policy simulation. In particular, for each state, we plot

³²For number of kids, we consider categories, 0, 1, 2, 3 and 4+.

median replacement rates (by the number of children) for the 1984:H1-2011:H2 period. A number of broad patterns emerge. First, there is substantial cross-sectional variation in average replacement rates. Second, replacement rates have been either stable or slightly increasing over time. Third, most changes in UI generosity are incremental, with states tweaking schedules frequently, mostly by changing minimum or maximum weekly benefits. Nonetheless, there have been episodes with substantial changes in replacement rates within-state (e.g. California in the early 2000s, or Louisiana in the early 1980s). Fourth, there is considerable variation in generosity by the number of children, with some states (e.g. Connecticut) having jumps of up to 5p.p. per dependent child. Fifth, in certain states, the median replacement rate is flat over time due to a binding maximum (e.g. California over the 2000s). Sixth, absent reforms in the schedule, replacement rates fall over time, as income often grows faster than weekly benefits.

D.3 Summary Statistics

Appendix table 11 collects summary statistics for our sample. All nominal variables are expressed in 2011 USD. Panel A collects individual-level characteristics of household heads. On average, 40% are female, 84% are caucasian, and 64% are married. The average head is 43 years old, has 14 years of education, one child, and earns a weekly wage of \$948. Annual household earnings average \$70,000, with considerable dispersion. Panel B reports balance sheet information. Our sample's balance sheets are qualitatively consistent with the evidence presented in Campbell (2006) based on broader data from the SCF.

Our liquid assets measure is Net Financial Assets (NFA), which sums assets in banking institutions (interest-bearing or not), assets in other financial institutions (interest-bearing or not), equity in stocks or mutual funds, and other financial assets, and then subtracts unsecured debts. This is the clearest mapping between liquid assets in the model and in the data. As we seek to capture liquid assets, NFA excludes retirement accounts and illiquid assets (eg. real estate or vehicles).

The average household holds about \$33,000 in NFA, the bulk of which comes from interest-bearing deposits. Other components of NFA only add mass to the tails of the distribution. Up to the 25th percentile, households in our sample are net borrowers in liquid assets. The median household only has \$0 in NFA. Only at the 75th percentile, we see NFA reaching about \$4,000, about 6% of median yearly earnings. A considerable share of wealth is stored in illiquid assets: vehicles in the lower tails of the distribution, and housing (often associated with sizable mortgages), as we move along the distribution. This apparent "illiquidity" of US households has been consistently documented in the Household Finance literature.

Panel C collects individual-level measures of UI generosity and social insurance reciprocity. Most heads (96%) are eligible for UI and, conditional on being eligible, would earn about \$1000 a month in benefits. The average replacement rate is 41%, ranging from 20% at the 10th percentile to 55% at the 90th percentile. Consistent with our sample selection efforts, only 5% of sample heads received UI at some point in the data, and only 7% received some non-UI social insurance transfer. Amounts received in social or private insurance transfers are negligible. Last, Panel D collects our

simulated policy instruments. The median head is eligible and enjoys a replacement rate of 42%, ranging from 32% at the 10th percentile to 52% at the 90th percentile.

Note that eligibility-adjusted measures (i.e. setting benefits to zero for individuals who are not eligible) are smaller on average, reflecting that many low-income heads are not eligible for UI. Ineligibility creates a right tail in the distribution of unadjusted replacement rates. This exemplifies the potential concerns with using individual-level replacement rates: the policy appears very generous, but this is driven by measurement error in income or an “unobservable” (in this case, eligibility) that confounds the generosity of the policy. Similarly, state-level instruments dramatically reduce the dispersion of simulated policy variables, better capturing statutory variations in UI schedules.

	Mean	S.D.	10th	25th	50th	75th	90th
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A: head characteristics</i>							
Age	42.69	10.12	29	34	42	51	57
Female	0.4						
Years of education	13.69	2.77	12	12	13	16	18
White	0.84						
Married	0.64						
Number of kids	1.06	1.13	0	0	1	2	3
Weekly wage	948.23	803.88	316	505	787	1172	1686
Annual HH earnings	71510	51480	24390	38758	60817	90659	126645
<i>Panel B: households balance sheets</i>							
NFA/Earnings	0.06	0.32	-0.19	-0.04	0.00	0.06	0.35
Net financial assets	6647	33490	-10423	-1942	0	3960	26877
Interest bearing	7795	19777	0	0	757	6073	20933
Stocks	3934	20876	0	0	0	0	2993
Retirement	7979	32249	0	0	0	0	17052
Vehicles	7974	10060	0	1324	5587	12487	20666
Home equity	61279	98456	0	0	22316	86140	178940
Total wealth	110934	183392	781	7172	44121	138977	298781
Unsecured debts	5596	16654	0	0	1018	6238	16199
Secured debts	68897	108170	0	0	23294	108071	191580
Net worth	106578	183954	-1324	3912	39160	134849	2995978
<i>Panel C: social insurance</i>							
Eligibility	0.96						
Weekly benefits (WBA)	261.29	106.54	133	186	250	324	405
Eligibility-adj. WBA	251.14	113.78	118	179	242	316	405
Replacement rate	0.41	1.98	0.2	0.29	0.41	0.50	0.55
Eligibility-adj. RR	0.38	0.15	0.17	0.27	0.39	0.50	0.54
Any UI	0.05						
UI	5.33	61.01	0	0	0	0	0
Any SI	0.07						
Non-UI transfers	12.10	93.07	0	0	0	0	0
Private insurance	18.39	159.98	0	0	0	0	0
<i>Panel D: simulated instruments</i>							
Mean eligibility	0.97						
Median RR	0.42	0.08	0.32	0.37	0.42	0.49	0.52
Median RR (elig. adj.)	0.41	0.08	0.31	0.36	0.41	0.48	0.51
Number of obs.	178121						

Table 11: Summary Statistics

Notes: This table reports summary statistics for our sample. Panel A contains individual household head characteristics such as demographics, household composition, and income. Panel B provides a snapshot of household balance sheets. Panel C reports individual household head information on social insurance reciprocity and (potential) unemployment insurance generosity. Panel D reports moments of our simulated policy measures. All dollar values are expressed in 2011 USD using the PCE series.

E Additional Empirical Results

Net Financial Assets			
	(1)	(2)	(3)
Panel A: two-stage least squares			
Replacement rate	-0.0035 (0.0456)	-0.0635 (0.0468)	0.0202 (0.0470)
Panel B: first stage			
Median RR	0.7041*** (0.0651)	0.5794*** (0.0536)	0.6853*** (0.0637)
<i>KP F-stat</i>	116.95	116.87	115.67
Outcome mean	0.06	0.06	0.06
<i>N</i>	177813	177813	177813
Individual controls	✓	✓	✓
State controls		✓	✓
Time FE	✓		✓
State FE	✓		✓

Table 12: TWFE IV Regressions

Notes: this table reports estimates of IV regressions of net financial assets (normalized by annual household earnings) on simulated unemployment insurance replacement rates. Panel A reports second-stage coefficients estimates by two-stage least squares. Panel B reports the corresponding first-stage estimates. Kleinbergen-Paap (2006) robust F-statistics are displayed below first-stage standard errors. Column (1) reports estimates of with individual-level controls for the household head, as well as state and time fixed effects. Column (2) removes the fixed effects and adds state-level controls. Column (3) adds back the fixed effects. All dollar values are expressed in 2011 dollars. Standard errors are clustered at the state level. *, **, *** denote significance at the 10, 5 and 1% levels.

	Baseline	No Small	No recessions	< Granular	GFA	Excl. risky	No trim	Trim 10%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Median RR	0.0566 (0.0535)	0.0593 (0.0546)	0.0690 (0.0525)	-0.0275 (0.0811)	0.0597** (0.0244)	0.0308 (0.0433)	0.5014 (0.3991)	0.0685** (0.0334)
Outc. mean	0.07	0.07	0.06	0.07	0.10	0.01	0.09	0.04
<i>N</i>	1180138	1168592	650835	1180138	1177715	1165945	1242233	1117999

Table 13: Stacked Regressions: Robustness

Notes: this table assesses robustness of our estimates in our stacked differences design. There are 47 reforms in our sample. All columns report estimates controlling for state-by-dataset and time-by-dataset fixed effects. Column (1) reports our baseline estimates. Column (2) drops any state-year cells with fewer than 25 observations. Column (3) estimates in the subset of reforms that do not occur during NBER recessions. Column (4) instead uses a less granular, state-time level measure of UI generosity rather than a state-time-number of kids measure. Columns (5) considers gross instead of net financial assets. Column (6) excludes stock and mutual funds from the computation of NFA. Column (7) repeats column (1) without trimming the outcome distribution. Column (8) trims the sample by discarding observations below the 5th above the 95th percentile of NFA over earnings. All dollar values are expressed in 2011 dollars. Standard errors are clustered at the state level. *, **, *** denote significance at the 10, 5 and 1% levels.

