Precautionary Savings and Stabilization Policy in a Present-Biased Economy

Adriano Fernandes* IMF

Rodolfo Rigato † ECB

First Draft: October 2022

This Draft: June 2026

Abstract

We study the sensitivity of precautionary savings to consumption risk in standard heterogeneous-agent (HA) models and find it is strong relative to the data. Using variation in unemployment insurance (UI) generosity across U.S. states, we rule out savings responses larger than about 0.75% of annual income per 10 percentage-point increase in replacement rates. This is at odds with standard HA models but consistent with a model with present-biased households. Disciplining precautionary savings is policy-relevant: an estimated HANK model predicts a 45% smaller fiscal multiplier of temporary UI extensions; and a 40% weaker effect of UI in reducing aggregate consumption volatility.

JEL classification codes: E32, E62, H53.

Keywords: Precautionary Saving, Unemployment Insurance, Macroeconomic Stabilization.

*International Monetary Fund, Research Department. Email: afernandes@imf.org.

†European Central Bank. Email: rodolfo.dinis_rigato@ecb.europa.eu.

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, Damien Capelle, Veronica De Falco, Xavier Gabaix, Benny Goldman, Pierre-Olivier Gourinchas, Omeed Maghzian, Peter Maxted, Ben Moll, Guilherme Neves, and numerous seminar participants for their helpful insights. The views expressed here are solely those of the authors and do not represent those of the European Central Bank or the International Monetary Fund, their Executive Boards, or their management. Any errors are our own.

1 Introduction

Heterogeneous-agent (HA) models have become central to the study of business cycles and stabilization policy, in large part because they generate large and dispersed marginal propensities to consume (McKay et al. 2016,Kaplan et al. 2018,Auclert et al. 2018). However, uninsurable consumption risk also leads agents to engage in precautionary savings — a channel that is quantitatively important for policy transmission but whose strength has received less scrutiny. 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?

Our first contribution is to document that in standard HA models, the precautionary savings response to changes in consumption risk is larger than what the data supports; but that the introduction of a single ingredient – present bias – substantially narrows this gap. Throughout, we focus on consumption risk due to *unemployment risk*, consistent with an extensive literature documenting its key role for income fluctuations.[1] We establish the gap between model and data in three steps. First, we set up two different HA models: one with standard (i.e., exponential) time preferences and another with present bias. We calibrate both models to successfully match a host of moments emphasized in previous work. 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 directly affects the consumption risk households face. In doing so, we derive results for calibrating HA models with unemployment risk; and show that present-biased models help match such moments, while the standard model struggles to do so. Second, we turn to the data and revisit empirical estimates of precautionary savings out of unemployment risk. Using quasi-random variation in UI schedules across U.S. states over time, we find small, statistically non-significant effects.

We then compare the sensitivity of precautionary savings implied by the two models to our empirical estimates. When running the same regression in the standard model, we obtain a large coefficient that falls outside the estimated confidence interval. The standard model's precautionary savings prediction is thus at odds with the data. The present-biased model, on the other hand, generates coefficients that lie within the empirical confidence intervals and is thus consistent with the data. This comparison is anchored by a calibration strategy that disciplines both models with two independent sets of micro moments — intertemporal MPCs and unemployment spending paths — before testing them on a third moment they were not calibrated to match. This result survives even when we allow the standard model to feature heterogeneous discount factors. Our empirical finding — that precautionary savings respond weakly to UI generosity — is consistent with other behavioral mechanisms, including inattention to policy changes or overly optimistic beliefs about reemployment prospects. We emphasize present bias because it is parsimonious: a single parameter modification simultaneously matches the weak precautionary response documented here and the sharp consumption drop at UI benefit expiration — a feature that standard models cannot

[1] For example, Guvenen et al. (2014) show that the distribution of yearly log income changes features fat tails. The left-tail can be explained by models featuring non-employment shocks, as recently shown by Guvenen et al. (2025).

replicate.

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. Our focus on UI is deliberate: it has become a staple policy across developed economies, with policymakers frequently alluding to its stabilization properties Blanchard et al. (2013). As such, we study the role of UI both as a discretionary policy tool and as an automatic stabilizer. As a discretionary stabilization tool, UI affects consumption not only by sustaining the consumption of the unemployed, but also largely through precautionary saving of the employed. 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 45% smaller when agents are present-biased. As an automatic stabilizer, the effect of UI in reducing aggregate consumption volatility is 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 a fraction of the unemployed do not collect their benefits, in accordance with empirical evidence. 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 time preferences and one with present bias. Present-biased agents strongly discount future utility relative to the present, more than the additional discount between two consecutive future dates implied by exponential discounting. Ericson and Laibson (2019) reviews a voluminous literature in psychology and behavioral economics providing empirical support to present-biased models in field and laboratory settings.[2] We focus on present bias of the naïve type: agents are not aware of their present bias, leading them to make dynamically inconsistent choices. In our model, these agents wish to self-insure against future income risk, but procrastinate. They do not start saving today because they wrongly believe they can commit to start saving tomorrow.

Both versions of the model – exponential and present-biased – are calibrated to match the same moments. Two sets of empirical moments are crucial. The first are the now usual intertemporal marginal propensities to consume (iMPC), as in Auclert et al. (2018). Matching iMPCs disciplines the dynamic response of consumption to marginal changes in disposable income, which could be generated by an increase in labor demand or in UI generosity. Both standard and present-biased models can match iMPCs.

The second set of moments is the average spending along an unemployment spell documented

[2]DellaVigna (2018) additionally 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.

by Ganong and Noel (2019). Matching these spending paths is important for three reasons. First, it directly measures consumption risk posed by unemployment. Second, we prove that a similar set of moments characterizes the partial equilibrium response of consumption to unexpected changes in job finding and separation rates. Therefore, the spending paths play a role similar to marginal propensities to consume to *extensive margin* variations in labor demand, i.e., changes in employment rather than in the hours worked by the employed. Third, the spending paths give us a way to calibrate the degree of present bias in the economy. Confirming the quantitative results in Ganong and Noel (2019), we show that the consumption drop upon UI expiration – a perfectly predictable event – is inconsistent with the consumption smoothing behavior of standard preferences, but is consistent with present bias.

With both calibrated models, we 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. 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. 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. To shed some light on these results, we analytically characterize the sensitivity of the equilibrium real interest rate to changes in social insurance in a zero liquidity limit, and find that time inconsistency accounts for the reduction in precautionary behavior.

Is the reduced sensitivity of precautionary saving for unemployment under present bias consistent with the data? To address this empirical issue, we revisit the empirical precautionary savings literature by exploiting variation in UI generosity across US states over time to estimate how strongly savings respond to UI reforms. Using Survey of Income and Program Participation (SIPP) data, we first expand on Engen and Gruber (2001) by leveraging a simulated instrument strategy 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 using a canonical TWFE design. This finding survives across a battery of robustness tests, including controlling for the level effect of UI changes on consumption risk. In our baseline 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. A formal power analysis confirms that this near-zero estimate is informative: our preferred design has sufficient power to detect precautionary responses of the magnitude predicted by the standard exponential model.

To tighten identification, we leverage a stacked differences-in-differences design (Cengiz et al., 2019). We first isolate states that experienced large changes in UI generosity using our simulated

replacement rates. Using these events, we estimate the dynamics of household wealth-to-income ratios in states that experienced a reform relative to those who did not. We confirm the results from our TWFE design: the effects are small and statistically non-significant.

To bridge the reduced form empirics and the model, we treat the regression estimates as external moments and reproduce the exact same regressions using model-simulated data. In the exponential model, 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 empirical confidence interval. The data thus rejects this prediction of the standard HA model. The present-biased model predicts an average increase of 0.21%, within the empirical confidence intervals.

We conclude the standard model generates precautionary savings out of unemployment consumption risk that are large in comparison to the data, and move to the implications for stabilization policy. To do so, we augment our model with wage rigidity, monetary and fiscal rules, and aggregate shocks; and estimate the additional parameters using Bayesian methods. First, we analyze the effects of discretionary UI extensions on aggregate consumption. Despite not drawing as much academic scrutiny as monetary policy or untargeted fiscal transfers, UI extensions have been routinely deployed in the US during recessions, and have therefore become a crucial stabilization tool. As in Kekre (2021), we find that 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. Our novel result is that most of the stabilization ammunition of UI in the model is driven by this precautionary response of employed workers. Quantitatively, in the standard HA model, a 12-month increase in UI duration has a fiscal multiplier of 1.13. When precautionary behavior is attenuated by present bias, the multiplier is 0.62 – a substantially smaller value.

We then turn to the role of UI as an automatic stabilizer, which have received renewed interest and support in policy circles. 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 9.7%, mostly by increasing anticipation effects. In a present-biased economy, this effect is dampened, and the same policy leads to a 5.9% 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 studies the effects of household heterogeneity for the transmission of macroeconomic shocks and stabilization policy (Kaplan and Violante, 2014; Werning, 2015; McKay et al., 2016; 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. 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 time preferences. We extend their analyses by first showing that standard models generate predictions about precautionary savings that exceed what the data supports, looking at aggregate predictions of these models only after narrowing 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.

A few papers in the aforementioned literature deserve special attention for focusing on UI or other automatic stabilizers. McKay and Reis (2016) 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. Similarly, Ferrière and Navarro (2025) compare the efficacy of a set of stabilization tools, including unemployment insurance. 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 may overstate their efficacy by featuring too strong precautionary behavior. In related work, McKay and Reis (2021) characterize the optimal design of automatic stabilizers in a zero liquidity setting. We abstain from normative analyses, focusing only on the positive effects of UI in a fully-fledged quantitative model.

Another closely related paper is Kekre (2021), who finds fiscal multipliers of discretionary UI extensions close to 1 in a quantitative model with exponential preferences. We confirm his predictions in our model with standard preferences. However, the present-biased specification delivers a fiscal multiplier around 45% lower, showing that matching the sensitivity of precautionary savings in the data is quantitatively meaningful . In a setting complementary to ours Bardóczy and Guerreiro (2023) study the role of expectation formation mechanisms in affecting the fiscal multiplier of unemployment insurance extensions. The main difference between both approaches is that we use microdata on household asset holdings to discipline the strength of precautionary saving, while Bardóczy and Guerreiro (2023) use forecast errors of professional forecasters.

Another related strand of literature studies how behavioral frictions interact with stabilization policy (Gabaix, 2020; Erceg et al., 2021; Vimercati et al., 2021). Closest to us is Laibson et al. (2021), who show how procrastination generated by present bias can slow the mortgage refinancing channel of monetary policy in partial equilibrium. We differ from them in three ways. First, we focus

on a different implication of present bias, namely reduced precautionary behavior in liquid asset accumulation. Second, we study fiscal policy and automatic stabilizers. Third, we entertain our counterfactual policy exercises in general equilibrium. Another related work is by Maxted (2020), who studies the implications of present bias for properties of heterogeneous-agent models (e.g., wealth distributions), but not for business cycles.

Last, a vast body of empirical work has studied the effects on income risk on precautionary savings. We discuss how our estimates fit into this literature in section 3.2. Japelli and Pistaferri (2017) summarize theoretical foundations and the ensuing empirical tests. . Our contribution to this literature is two-fold. First, our stacked DiD design is of independent methodological interest: we contribute a stacked differences-in-differences design that isolates large, discrete reforms in state-level UI schedules for identification. Second, we align reduced form econometric estimates and quantitative HA macro models.[3] Our starting point is Engen and Gruber (2001). Empirically, we (i) incorporate more data and (ii) tighten identification with stacked regressions. Theoretically, we abstract from life-cycle considerations, but incorporate general equilibrium forces in the HA macro tradition and by match more recent, frontier estimates of other moments relevant for quantitative predictions.

Outline. The paper proceeds as follows. In section 2, we introduce a standard model featuring 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. Section 4 introduces business cycles and shows the implications for stabilization policy. Section 5 then presents our concluding remarks.

2 A Standard Model of Precautionary Savings

We introduce a standard model of consumption 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 savings, leaving features such as monetary policy and nominal rigidity for section 4.1. Until then, all variables are expressed in real terms.

[3] A notable exception is Carroll and Samwick (1997), who find that "wealth is estimated to be far less sensitive to uncertainty than implied by that [life cycle] model.". Effectively, their result is an early version of ours, but in a simpler life-cycle model: the liquidity constraint plus income risk chassis of many consumption models predict too strong precautionary savings relative to the data. Other papers motivate the reduced form analysis with some quantitative simulations (as we do), but do not "close the loop" and attempt to make the quantitative model match the regression estimates.

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

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

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. A recursive formulation of this problem is shown in the appendix.

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.[4] 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:

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

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

[4] Alternatively, one could assume that agents are *sophisticates*: they understand their future selves will deviate from their current plans. Harris and Laibson (2001) derive a generalized Euler equation allowing for sophistication. 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.

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 unemployed 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 4.1, 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 Coglianese 2019). To account for this fact, we follow Kekre (2021) in assuming that only a randomly selected fraction ξ 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

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

The government budget constraint is, therefore,

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

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

$$Y_t^{firms} = Z \cdot N_t, \tag{5}$$

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

$$N_t = E_t \cdot n_t. \tag{6}$$

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

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

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

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

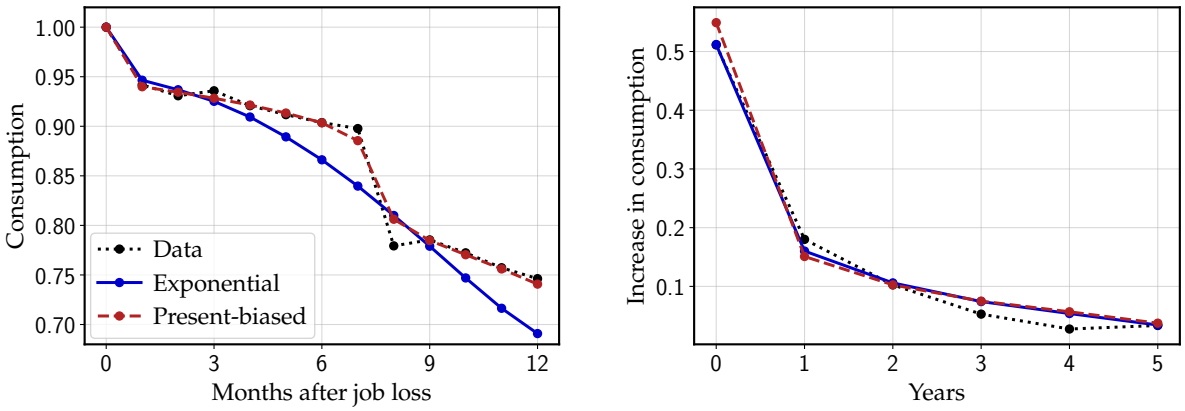
2.2 Calibration

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. Importantly, here we display benchmark calibrations of both models. Section 3.3 considers numerous extensions when comparing model properties with the data, such preference heterogeneity, different assumptions on borrowing and saving, and reduced-form search behavior.

We calibrate some 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 together with the values generated by each of our models, normalizing pre-unemployment consumption to 1. Ganong and Noel (2019) compute these moments using the universe of Chase consumer checking and credit card accounts. They 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, typically after 6 months. They also show that present-biased preferences are consistent with this feature of the data. Figure 1 thus replicates their findings.[5]

These moments are important beyond their ability to compare present biased and standard time preferences. First, we are interested in precautionary savings, which depend crucially on the consumption risk generated by unemployment risk. This is precisely what the spending paths measure: how much consumption is at risk for employed households were they to become unemployed. They are thus directly informative about a key driver of precautionary savings in many models. Second, moving from micro to macro, we show that consumption dynamics during unemployment are tightly linked to the response of aggregate consumption to employment fluctuations.

[5]We also confirm that the implied path of income during unemployment aligns with the empirical counterparts in Ganong and Noel (2019) across both models. A competing behavioral mechanism is overly optimistic beliefs about job finding rates, as empirically documented by Mueller et al. (2021). Ganong and Noel (2019) explore this alternative, but find that the degree of optimism that would be required to match the drop in consumption upon UI exhaustion far exceeds what is in the data.



(a) Consumption during unemployment

(b) iMPCs

Figure 1: Calibrated consumption during unemployment and iMPCs.

In the appendix, we derive an expression relating a more general set of spending paths during unemployment to the response of aggregate consumption to changes in job finding rates – akin to how MPCs determine the response of aggregate consumption to marginal variations in disposable income.

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). This implies a coefficient of relative prudence of 3, in line with frontier empirical estimates (e.g., Fagereng et al., 2017).[6] The large precautionary savings produced by the standard model are thus not driven by implausible estimates of relative prudence, and we later show our results remain if we instead set $\sigma = 1$. The AR(1) parameters of the idiosyncratic labor productivity process are taken from Floden and Lindé (2001), adjusted for our monthly setting.[7] The public debt B is set at 26% of aggregate income to match the aggregate liquid asset holdings of households

[6]Our choice of relative prudence also falls in the range of other notable estimates (Gourinchas and Parker, 2002; Bertola et al., 2005). Our parameter values also align with Laibson et al. (2024), who estimate a large-scale life-cycle model a la Gourinchas and Parker (2002) with present-bias and CRRA utility. They find a coefficient of relative prudence of 2.9, a discount rate of 0.99 and a present-bias coefficient of 0.53.

[7]Floden 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.

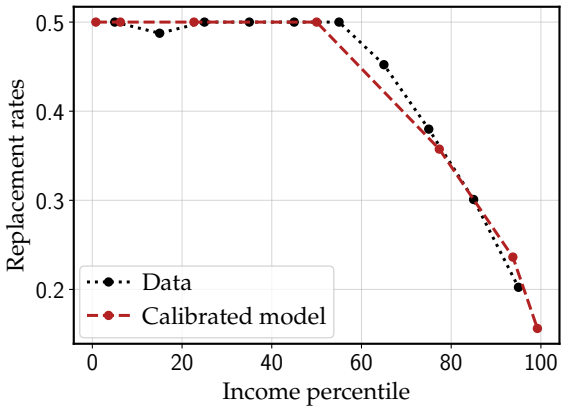


Figure 2: Replacement rates along the income distribution.

(Kaplan et al., 2018).

We set $b = 0.5$ and $\bar{b} = 0.48$ to calibrate the UI system. As we normalize $n = w = 1$ in steady state, the latter value corresponds to 48% of the average pre-tax labor income. The UI parameters are set to match the distribution of replacement rates in the SIPP data (see section 3.2) plotted 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. To our knowledge, no other quantitative model with HA and UI has targeted and matched this crucial policy dimension. Last, UI duration is set $N^{ui} = 7$, as in Ganong and Noel (2019) and most UI laws.

We choose the take-up rate $\xi = 0.77$ to match the estimated UI take-up among eligible workers from Auray et al. (2019).[8]. 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 $\mu^w = 1.1$. 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 computed by Laibson et al. (2021). The monthly 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 an average consumption of the unemployed equal to 80.4% of the average consumption of the employed, in line with empirical estimates for the United

[8]In section 3.2, some sample restriction assumptions are required to better capture the effects of UI on asset accumulation of employed households. They leave us almost exclusively with a sample of eligible households. Our chosen take-up rate further allows us to better compare our model with the empirical evidence.

Parameter	Meaning	Value		Source/Target
Externally calibrated parameters				
σ	CRRA	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	$\{0.5, 0.9\}$	iMPC and consumption during unemployment
ω_H	Share of high β	–	0.74	
y^{priv}	Non-UI income	0.320	0.320	
T^{sn}	Safety net transfers	0.087	0.087	
$\underline{a}$	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.

States (McKay and Reis, 2021). As the drop in consumption upon unemployment is an important determinant of precautionary saving (Schaab, 2020), we calibrate the exponential model to generate this ratio. Crucially, the value $\underline{a} = 0$ for the exponential model is not a restriction.– It is necessary for matching iMPCs and the 80.4% ratio. Even with $\underline{a} = 0$, the exponential model is not able to generate such a large consumption drop , delivering a ratio of 85.7%. Therefore, there is a trade-off between imposing the same parameter values in both models or having them match the same set of observable moments. We choose the latter approach, allowing the borrowing constraints to differ in the baseline calibration. However, when contrasting models and data in section 3, we relax the borrowing constraint of the exponential model as a robustness check.

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.

Our calibration strategy disciplines both models with two independent sets of micro moments before testing them against a third. Both economies target the same iMPC schedule and the same unemployment spending paths documented byGanong and Noel (2019) — though the present-biased model matches the sharp drop in consumption at UI benefit expiration substantially better than the exponential model, which smooths through expiration more than the data suggest. The question we then ask is whether a model that has been disciplined by these moments also matches

the precautionary savings response to UI generosity — a moment neither model was calibrated to hit.

The present-biased model features heterogeneous β, while the baseline exponential model has a single discount factor. To put both on equal footing, we also calibrate an exponential model with heterogeneous discount factors (δ_L,δ_H), targeting the same moments. Even with extreme heterogeneity — an annualized δ_L as low as 0.78 — the model still cannot replicate the sharp consumption drop at benefit expiration: impatient types run down assets quickly but do not generate the distinctive spending cliff that present bias produces. Full details appear in the Online Appendix.

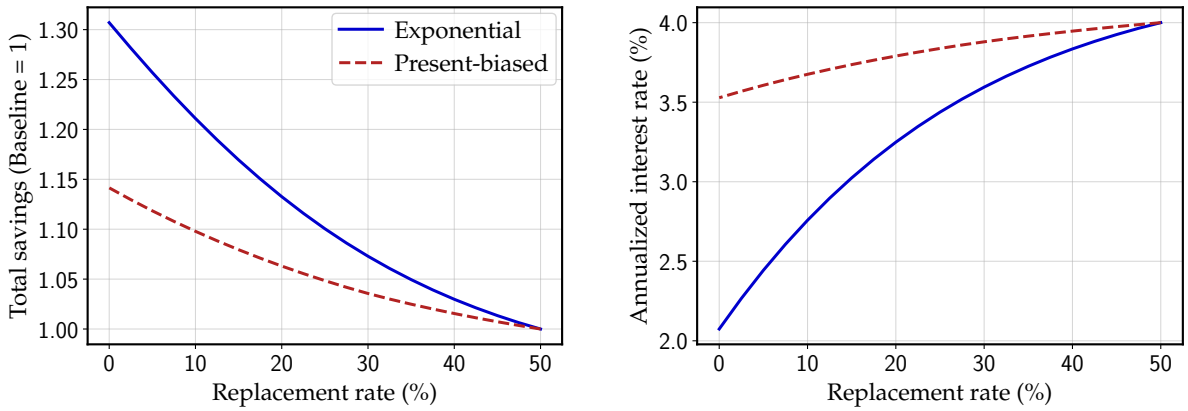
3 UI and Precautionary Savings: Models and Data

We now turn to the following question: if there is decline in consumption insurance (e.g., a reduction in UI), by how much do precautionary savings increase? We start by characterizing this sensitivity of savings to changes in social insurance in our models, both numerically and analytically for a tractable special case. Then, we provide estimates of similar objects using SIPP data and show that the exponential 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 Precautionary Savings in the Models

Before studying the strength of precautionary savings in our models, it is helpful to first define what we mean by precautionary saving and how this relates to MPCs.The MPC is the slope of the consumption function with respect to cash-on-hand at this agent's income and wealth states. In HA models, consumption functions are typically concave (equivalently, MPCs are decreasing) on wealth due to precautionary saving: agents in good states (with high income or large asset holdings) would save a larger fraction of a marginal dollar than agents in bad states, where they are close to the borrowing constraint and have near-unit MPCs. Precautionary savings and MPCs are, therefore, tightly connected (Kimball, 1990; Achdou et al., 2021).

In what follows, we define the strength of precautionary saving as the change in asset holdings or interest rates in response to a change in a parameter that affects consumption risk (in our case, the UI schedule). Therefore, our definition of precautionary saving is related to the *shift* in the consumption function following a policy change, as opposed to its slope, although MPCs may also change. Interestingly, this shift happens even for agents that are not immediately affected by the policy – for instance, the consumption behavior of employed agents changes even if their income remains the same after a UI reform (except, of course, by potentially small changes in the taxes that are used by the government to finance UI). As we show below, the exponential and present-biased models display very different precautionary saving behavior despite the fact that they are initially calibrated to match the same average MPCs.



(a) Partial equilibrium

(b) General equilibrium

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

We start by computing the response of precautionary savings to changes in UI. 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.[9] 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 further understand the role of present bias, in the appendix we describe a limiting case where we can derive closed-form expressions for the GE case. In a zero liquidity limit, we show that most of the attenuation in precautionary savings is driven by a *time inconsistency* effect: households naively believe that their future selves will start saving, but, in equilibrium, never do so. This "procrastination" effect has received extensive empirical support (Beshears et al., 2018).[10]

[9]MPCs and precautionary savings are connected: both large MPCs at low wealth levels and precautionary savings are properties of concave consumption functions (Kimball, 1990). It is convenient to reiterate the proper definition of these concepts in our context. MPCs are related to the average slope of the consumption function (as a function of liquid asset holdings), while precautionary savings refer to shifts in these functions in response to changes in idiosyncratic income risk, here proxied by the UI replacement rate. Analogously, matching iMPCs is not sufficient for matching precautionary savings.

[10]Carroll et al. (2021) shows that reducing the risk of hitting liquidity constraints will tend to make the consumption

In sum, our theoretical analyses highlighted two points. First, exponential agents in standard models exhibit sizable precautionary behavior. Second, introducing present bias attenuates precautionary behavior. Towards better discriminating across models, we measure the sensitivity of precautionary savings to UI generosity in the data, and then show that the implications of the present-biased model are more aligned with our empirical estimates.

3.2 Empirical Evidence

We revisit the empirical precautionary savings literature by measuring the effect of varying the generosity of UI on the stock of liquid assets of households. We leverage an institutional feature of the US UI system: each 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. We obtain small, statistically non-significant effects — informative bounds that allow us to compare across models. We then run the same regressions in model-generated data, both for our exponential and present-biased calibrations. The exponential model's predictions fall outside our confidence bands, while the present-biased model's predictions lie within them.

Data. We use the 1984–2008 panels from the Survey of Income and Program Participation (SIPP), focusing on household heads with stable earnings histories for whom we observe both UI generosity and liquid asset holdings. Following Engen and Gruber (2001), our final sample covers all 50 states plus DC over the 1984–2011 period. We measure precautionary savings using the ratio of net financial assets to annual income. Individual-level UI replacement rates are constructed using the calculator from Kuka (2020), which codifies state-specific benefit formulae. Detailed descriptions of the SIPP sample construction, variable definitions, and UI institutional features are in the online data appendix.

We use simulated policy instruments following Currie and Gruber (1996) rather than individual-level replacement rates. Individual-level rates raise endogeneity concerns from spatial selection and measurement error from noisy earnings histories. To construct the instruments, we strip state-time information from our sample, apply each state-time UI schedule to this national counterfactual sample, and compute median replacement rates for each state-time-number-of-children cell. This approach eliminates confounds from individual-level selection while aggregating away idiosyncratic measurement error. See the online data appendix for details.

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:

function less concave. Present bias is operative through this mechanism as well: by erroneously believing their future selves will self-insure against unemployment risk, present-biased agents mistakenly understate their current and future risk of hitting the borrowing constraint, therefore lowering their precautionary savings.

$$\left(\frac{Liquid\ Assets}{Earnings}\right)_{ist} = \alpha_s + \tau_t + \gamma \times RR_{st} + X'_{ist}\beta + \varepsilon_{ist} \tag{8}$$

The dependent variable is 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. 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. 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), will confound estimates of simple cross-sectional regressions of liquid assets on UI generosity. Inclusion of state effects controls for these time-invariant state-level confounders.

Second, suppose all states have the same UI schedule. Omitted trends or cyclical factors affecting the time series of UI and liquid assets simultaneously will confound estimates of time series regressions of liquid assets on UI generosity. For instance, households' asset values fall in recessions, while generosity tends to rise due to federal-level UI extensions. Inclusion of time effects controls for these state-invariant time-level confounders.[11] Controlling for time and state effects is essential to isolate unobservable confounders that are present in the data but not in the model.

As standard in TWFE designs, identification relies on within-state time series variation. In plain English, we rely neither on comparing California and Missouri in 2001 nor on comparing trends or fluctuations in aggregate liquid asset buffers and UI generosities. 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: a quadratic polynomial 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.[12] All specifications control for simulated eligibility.

Another possible concern is that unobserved factors at the state-time level could simultane-

[11] Time 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.

[12] We 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.

ously influence the generosity of UI and households' precautionary savings. For instance, UI changes could be driven by local business cycles or correlated with changes in other social insurance programs. To mitigate this concern, we include time-varying state-level controls: the unemployment rate to absorb time-varying local economic slack; measures of welfare spending to control for the generosity of other policies households may have access to;[13] and measures of UI spending (regular state payments, regular federal payments, and emergency payments) and net UI reserves to control for budgetary and political considerations. Last, we include dummies for discretionary state-level UI extensions not captured in our model.

Summary Statistics. The appendix collects summary statistics for our sample and provides an extensive discussion. We summarize 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).[14]

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. Once we normalize by household earnings,[15] 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.

[13]For 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.

[14]Although 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.

[15]More 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.

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.[16] The appendix shows the policy variation we leverage. 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). Second, most variation in generosity is cross-sectional: state (time) effects absorb 62% (15%) of the variation in median replacement rates. As a reminder, our TWFE specification leverages neither pure cross-sectional nor pure time-series variation. 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 within-state changes in replacement rates, a source of variation we later explore.

TWFE Results. Table 2 reports OLS estimates of specification (8) 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 for comparing the two models.

Column 1 shows 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, , median replacement rates change on average by 0.26 p.p. year-over-year within a state, with a standard deviation of 2.5 p.p.[17] For this average change, 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.

We remove fixed effects in column 2, adding instead state-time controls. 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.

[16] Consistent 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 recipiency in our sample might seem low compared to our calibration. They are not. Social or private insurance recipiency 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.

[17] To be precise, 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. Without residualizing by the fixed effects, the standard deviation rises to 3.5p.p. However, residualized changes better reflect the variation we leverage for estimation.

	Net Financial Assets		
	(1)	(2)	(3)
Median RR	−0.0051	−0.0391	0.0109
	(0.0299)	(0.0261)	(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 bearings assets on banking institutions, interest-bearing or non-interest bearings 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.

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 the appendix 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.

A final concern is that UI variation may lead to changes in both levels and risk of consumption. As replacement rates rise, income and consumption during unemployment rise, so the (expected) ratio of unemployed to employed consumption rises. This may lead to an income effect, encouraging households to increase consumption and lower savings. Hence, the level effect would

confound the effects of UI on consumption risk. Column 9 addresses this concern. It controls for a measure of the consumption drop upon unemployment that a household would face if they lived in a state-time cell (s, t), estimated by extending the methodology of Chodorow-Reich and Karabarbounis (2016). Details can be found in the appendix. The point estimate barely moves relative to our baseline. The standard errors are slightly larger given the sample reduction, but the two-standard error lower band still rules out decreases in asset-income ratios smaller than -0.006 per 10 p.p. increase in median replacement rates, as in our main result.

	Baseline	GFA	Excl. risky	Elig. adj.	< Granular	No trim	Trim 10%	Time cluster	Levels Ctrl
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(8)
Median RR	0.0109	0.0229	−0.0066	0.0100	0.0010	−0.2432	0.0132	0.0109	0.0155
	(0.0298)	(0.0238)	(0.0233)	(0.0293)	(0.0364)	(0.2031)	(0.0241)	(0.0294)	(0.0391)
Outc. mean	0.06	0.10	0.01	0.06	0.06	0.17	0.04	0.06	0.05
N	178119	177812	176024	178119	178119	187504	168731	178119	113785

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. Column (9) includes the state-time level controls for consumption drops constructed based on the methods by Chodorow-Reich and Karabarbounis (2016). Except for column (8), standard errors are clustered at the state level.

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. When we observe variation in replacement rates, there are two possibilities. First, the parameters of the UI schedule have changed. This is a *statutory policy* variation. Second, parameters may have remained fixed, but inputs into benefits calculations may have changed. This is *not* policy variation. Our model makes statements about *policy* variation. When confronting our model with the data, the relevant identified moment is γ, the OLS estimate.[18]

The small and statistically insignificant coefficients in our baseline specifications are not uninformative zeros. A power analysis shows that our preferred TWFE design can rule out declines in net financial assets larger than about 0.75% of annual income for a 10 percentage-point increase in the replacement rate. Put differently, the empirical design is sufficiently informative to rule out economically meaningful precautionary-savings responses. As we will see, the lower bound of our confidence bands rule out the effects predicted by exponential models but are consistent with

[18] Appendix table **??** presents an exactly-identified IV model for completeness. Panel B **??** reports first-stage results. The effects are all significant at the 1% level. Panel A reports two-stage least squares estimates. Expectedly, our lower than 1-1 passthrough estimates inflate our IV estimates. But the effects remain a statistical zero, and the sensitivity to controls parallels exactly that observed in Table 2.

the effects predicted by the behavioral model.[19]

Stacked Regressions A recent econometrics literature has raised concerns about interpreting the canonical TWFE specification in some settings (see De Chaisemartin and D'Haultfoeuille, 2022 for a review). First, we discuss how these concerns apply to our case. Second, we implement a complementary identification strategy that addresses these concerns by exploiting a different source of variation: large, discrete UI reform events in a stacked differences-in-differences design (Gormley and Matsa, 2011,Cengiz et al., 2019). Both identification strategies yield the same substantive conclusion.

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 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, the estimator includes comparisons between states whose replacement rates have not changed much and states whose replacement rates have changed a bit. 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. 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, the appendix shows that states often change UI generosity, with a much more complex timing than simple one-

[19] We focus on static estimates rather than a distributed lag specification with leads and lags of the replacement rate. Given the high persistence of simulated replacement rates within states, individual lead and lag coefficients are poorly identified due to multicollinearity in the dynamic regressor matrix (**?**;**?**). The static coefficient — a persistence-weighted average of dynamic effects — remains well-identified. In the dynamic specification, pre-treatment coefficients are jointly insignificant (p = 0.19), supporting the parallel trends assumption. Post-treatment coefficients are jointly significant (p = 0.004), though individually imprecise. The cumulative post-treatment effect (0.036, s.e. 0.052) is statistically equivalent to the static estimate.

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 quasi-exogenous conditional on controls, one can interpret the exercise as 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 4 years prior to 8 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.[20] 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.[21] 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, 42 are clean events that can be used for estimation. We stack all datasets, assigning to each an identifier. We then estimate (8) 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

$$\left(\frac{Liquid\,Assets}{Earnings}\right)_{isth} = \alpha_{sh} + \tau_{th} + \gamma \times RR_{sth} + X'_{isth}\beta + \varepsilon_{isth} \tag{9}$$

[20] A 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.

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

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 -0.0068 per 10 p.p. increase in median replacement rates.[22]

In the appendix, we show our results are robust to a battery of robustness tests: dropping all state-year observations with less than 25 observations; discarding all events that happen during NBER recessions; controlling for levels effect of UI using our state-time consumption drops upon unemployment extended from Chodorow-Reich and Karabarbounis (2016); excluding stocks or mutual funds from NFA; or trimming 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 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 inflate standard errors too much for us to be able to detect any reasonable effects. 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 $\approx$2.5, more in line with model-based wealth distributions, and also imply estimates quantitatively consistent with our other results.[23]

In the online appendix, we validate the stacked design using a event study. Reforms produce a persistent increase in replacement rates of approximately 3 percentage points. The effects on wealth are uniformly small and statistically insignificant across all event-time horizons.

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 differences reinforces our belief that our estimated estimates are not artificially lower due to some of the mechanics of the canonical TWFE design.

Previous Literature. The literature attempting to estimate precautionary savings is extensive. We nevertheless try to situate our strategy relative to this vast literature.

Dynan (1993) provides an early, direct estimate of precautionary savings. Using GMM and alternative sets of instruments to estimate a quadratic approximation to the consumption Euler equation, she finds very small coefficients of relative prudence. Other works would pursue less

[22] In the Online Appendix, we validate the stacked design using a binary event study. Reforms produce a persistent increase in replacement rates of approximately 3 percentage points , with no evidence of pre-reform differential trends. The reduced-form effect on net financial assets is near-zero across all event-time horizons, confirming that the null result is not driven by"forbidden comparisons" or dynamic misspecification.

[23] Results are also robust to replacing household-level controls with demographic-cell fixed effects defined by 10-year age bins, four education categories, sex, and race. Results are available upon request.

	Net Financial Assets		
	(1)	(2)	(3)
Median RR	0.1014*	−0.0238	0.0616
	(0.0512)	(0.0377)	(0.0665)
Outcome mean	0.07	0.07	0.07
N	1127940	1127940	1127940
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.

direct tests, often testing for associations between wealth holding and income risk measures.[24] Table 5.2 in Browning and Lusardi (1996) summarized the state of the literature, finding tremendous variation in empirical contexts, datasets, econometric methods, measures of income or consumption risk, and main findings. At the time, they concluded that "while the precautionary motive is important for some people at some times, it is unlikely to be so for most people (...) [and] unlikely to be as important as some studies suggest." (p. 1838).

More closely to us, Carroll et al. (2003) focus on unemployment risk. They estimates the relationship between household total net worth and one-year-ahead job loss risk, relying on household region as an instrument.[25] Their results are mixed: a pooled model finds no evidence of precautionary savings, but heterogeneity analyses detect positive effects for households above the 30th percentile of permanent income. Moreover, they find no statistically significant relationship when looking at net financial assets, consistent with our findings.

A consistent challenge across these studies was that realized or subjective measures of income

[24]Carroll and Samwick (1997) estimate a significant association between wealth and the variance of income shocks using data from the PSID. Carroll and Samwick (1998) establish a similar empirical result, but using an empirical proxy of Kimball's equivalent prudence premium. Both studies, however, rely on small samples, non-clustered standard errors, samples of positive wealth, and measures of income risk or prudence premium that may suffer from model misspecification. Other influential studies leveraged surveys where households were elicited their subjective beliefs of income risk. Guiso et al. (1992) and Lusardi (1998) implemented similar strategies in Italy and US data, respectively, concluding that precautionary savings played a little role in wealth accumulation.

[25]There are some econometric concerns. First, the exogeneity restriction may be violated, as households may spatially select into locations with lower job loss risk. Second, the use of inverse hyperbolic signs to deal with non-positive skewed outcomes has been criticized by Chen and Roth (2023), who show that estimates of arbitrary sizes can be obtained with such transforms depending on scaling and instead recommend the use of scaled ratios, as we do. Third, results often become insignificant under alternative instrument sets. Last, and perhaps most importantly, the standard errors fail to account for clustering in treatment (i.e. correlated job loss risk) and thus likely overstate the statistical significance of their findings.

or unemployment risk suffer from confounding or selection concerns. Early IV strategies – education, occupation, or industry instruments – did not manage to circumvent these issues, as they are choice variables possibly correlated with risk preferences.[26] We instead prefer to leverage statutory policy variation, as Engen and Gruber (2001). Further, we tightened their original identification strategy by (i) using a simulated policy approach to reduce concerns of spatial selection, (ii) expanding the set of state-time controls, and (iii) deploying a novel stacked differences specification to circumvent recent econometric concerns related to TWFE designs.[27]

The most recent estimates leverage Scandinavian administrative registries. Juelsrud and Wold (2024) revisit the pass-through of unemployment risk to household savings studied by Carroll et al. (2003). They compare oil engineers versus other high-skilled workers in two Norwegian oil-producing regions following a 2014 collapse in oil prices, finding that a 1 p.p. increase in separation rates leads to a 1.3% increase in liquid assets, but null effects in illiquid assets. The effect is strong evidence in favor of precautionary savings out of unemployment risk. However, the comparison between oil engineers and other high-skilled workers is too specific to be portable for a business cycles macro model.

Fagereng et al. (2017) provides a more portable moment by estimating Euler equations in Norwegian data. Their approach isolates shocks to firm value-added, matches them to households via employer-employee registries, and then uses them as instruments for residual earnings or consumption volatility, striking a better balance between internal and external validity. They find a coefficient of relative prudence of 2, which is close to our calibration's value. Christelis et al. (2020) and Guiso and Jappelli (2024) also find coefficients of relative prudence in the range of 2 to 3 in Dutch and Italian survey data, respectively. Last, Laibson et al. (2024) extend the seminal model of Gourinchas and Parker (2002) to jointly estimates discount rates, present bias, and relative prudence under CRRA preferences using simulated method of moments. They show that allowing for present bias is crucial to match credit card indebtedness patterns in the US, and find a coefficient of relative prudence of 3. Our calibration of relative prudence is in line with the literature.

Taking stock, isolating precautionary savings is empirically challenging, with no clear preferred empirical strategy. We leveraged statutory, quasi-random policy variation, which trades-off well internal and external validity. Ours is not the only approach. Other applications may find other moments more useful, and further research is needed. But all in all, we offer credible and portable micro elasticities for HA business cycle models relative to previous estimates, allowing us to align macro models to such micro moments. Moreover, our calibration is consistent with

[26] Fuchs-Schündeln and Schündeln (2005) provided evidence of a particular mechanism: risk averse individuals select into less risky occupations. Using the German reunification experience, they show that measures frequently used by the previous literature could understate the degree of precautionary motives. Their point estimates, however, are hard to translate into quantitative magnitudes that are relevant for calibrating HA models. In turn, our approach leveraging policy variation does not suffer from self-selection concerns (especially as we use a simulated policy strategy) and allow for a direct replication of the empirical strategy in model-simulated data.

[27] Bornstein and Indarte (2023) also estimated the effects of social insurance on household savings by leveraging plausibly exogenous policy variation across geographies over time. They find that a 1 p.p. increase in Medicaid eligibility increases credit card borrowing by 1.17%. Our model cannot speak to their point estimates, as we focus on unemployment rather than health shocks as a source of consumption risk.

increasing consensus among frontier structural estimates of relative prudence.

3.3 Bridging Model and Data

We now compare the properties of the two models to our empirical estimates of the sensitivity of precautionary savings to UI generosity. To do so, we run the same reduced-form regressions in model-generated data using both the exponential and present-biased calibrations. We then compare the model-based regression coefficients to the confidence bands estimated in the data. Effectively, we are testing whether each model can pass an external validation test, where the external moment is the reduced-form regression coefficient we estimated from the data. This approach is analogous to the common practice in macroeconomics of checking whether model-implied impulse response functions fall within the confidence bands of empirically estimated IRFs. We apply the same logic to a micro-empirical regression coefficient rather than a time-series IRF.

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.

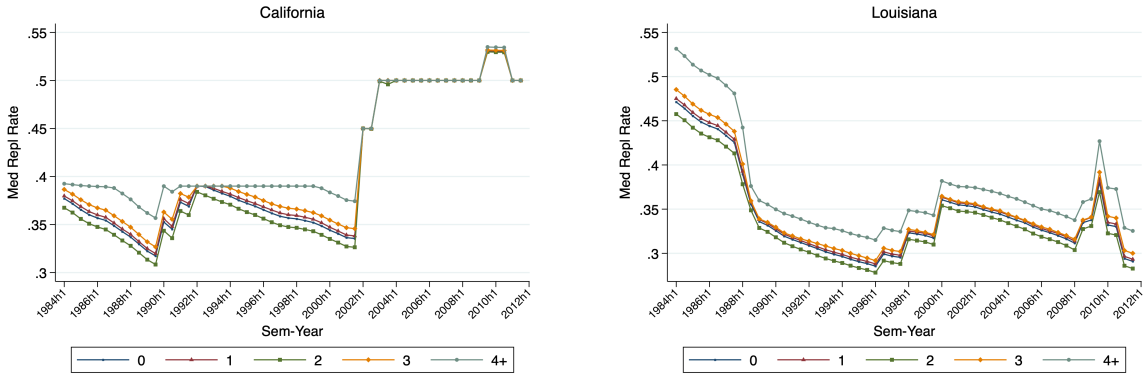


Figure 4: Median replacement rates in California and Louisiana by number of children.

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 the appendix). 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.[28] 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. To ensure that the results from our simulated regressions are not influenced by any type of sampling variation, we repeat this procedure 300 times and report the average estimated coefficient across simulations. Table 5 shows the results. The exponential model delivers coefficients that are substantially larger than our empirical estimates, and always outside the estimated confidence intervals. On the other hand, the model with present-biased households delivers much lower coefficients that lie within confidence intervals.

A formal power analysis confirms that this comparison is not an artifact of imprecise estimation. In the main TWFE specification, the exponential benchmark implies a decline of roughly 0.64% of annual income per 10 percentage-point increase in the replacement rate — a magnitude that falls near the design's rejection threshold ($\beta^* = 0.64\%$), at which the test achieves conventional power. The present-bias benchmark, by contrast, implies a decline of about 0.21%, well inside the region where the design cannot reject. The near-coincidence of the exponential prediction and the rejection threshold means the data are just informative enough to distinguish the two models: a substantially weaker design would not have been able to do so. The online appendix reports rejection thresholds across all robustness specifications.

Note that we are *not* estimating any model parameters based on this exercise.[29] As discussed above, this mirrors the standard IRF validation exercise in macroeconomics. We are asking the competing models to play a fair game: at first, we give them ample room to match a host of moments that the previous literature has shown to be important; then, we ask which calibrated model can best match a new external moment.

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

[29] This is also why we do not need to consider model-based "standard errors" when comparing the model estimates with the data estimates. This is because the only source of uncertainty in the model is simulation uncertainty. The standard error of the model-based regression coefficient goes to zero at a rate $o(\sqrt{N_{sims}})$, where N_{sims} is the number of simulations. Hence, we can make the standard error arbitrarily small by just increasing the number of simulations. The only variability left is sampling error from the data, which the standard errors of the data estimates already capture.

	Data	Exponential	Present-biased
Median RR	0.0109	-0.0636	-0.0207
	(0.0298)	–	–
Lower 95% confidence band	-0.0475	–	–
Power vs model (%)		79.2	27.0
β_{model} in SEs below est.		2.50	1.06

Notes: The 'Data' column reports our preferred TWFE estimate (Table 3, Column 4). The 'Exponential' and 'Present-biased' columns report coefficients from regressions of the same specification run on model-simulated data. Power vs model reports the probability of rejecting the null hypothesis of no effect in a one-sided test at the 5% significance level, assuming the true coefficient equals the model-implied value. Formally, Power =$\Pr(t > t_{0.05} \mid \beta = \beta_{\text{model}}) = \Phi(|\beta_{\text{model}} - \hat{\beta}|/\text{SE} - t_{0.05})$, where$\hat{\beta}$ and SE are the data estimate and its standard error, and$t_{0.05}$ is the critical value for the regression's residual degrees of freedom. We have 79.2% power against the exponential model, meaning that if the true effect were as large as the exponential model predicts, we would reject the null in about 4 out of 5 replications of the study. Power against the present-bias model is 27.0%, consistent with our inability to distinguish it from a zero effect.

Table 5: Regressions in model-simulated data.

Robustness. We now turn to robustness exercises that show that our result is indeed a feature of many HA models with exponential discounting, not only our specific one. 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. For robustnesss, 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. In all variants, the predicted coefficients remain outside the 95% confidence bands estimated in the data, suggesting that our findings are not an artifact of our specific calibration choices.

The first variant considers log utility, corresponding to a coefficient of relative risk aversion$\sigma = 1$. Next, we generalize our baseline model by allowing the job finding rates of unemployed agents that receive UI to be functions of the replacement rate. This extension parsimoniously allows us to model the disincentive effects of UI on search effort, a supply-side behavioral response emphasized in previous models. 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 the elasticity of job finding rates to UI generosity to be -0.02, as estimated by Ganong et al. (2021). We also consider reducing the supply of liquid assets in the economy by half, as well as increasing it by 50%. Another variant follows a different calibration strategy for the exponential model. We set $\underline{a} = -1.48$ and $\chi = 0.0098$, as in the

Variant	Coefficient
Log utility ($\sigma = 1$)	-0.0551
Varying job finding rate (f)	-0.0589
Low liquidity	-0.0611
High liquidity	-0.0607
Borrowing	-0.0802
Het. discount factors	-0.0578
No private insurance ($y^{priv} = 0$)	-0.1885
No UI expiration ($N^{ui} = \infty$)	-0.0842
Full take-up ($\xi = 1$)	-0.0903

Table 6: Regressions for variants of the exponential model.

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 calibration.

We also show results from a model with discount factor heterogeneity, described in more detail in the appendix. One possible concern regarding our calibration strategy is that the present-biased model features more degrees of freedom due to β heterogeneity, which could potentially lead to an unfair comparison against the standard model. This specification addresses this issue by considering a model with two discount factors $\delta \in \{\delta_L, \delta_H\}$.

So far, most variations of the model delivered slightly smaller coefficients, but still outside the confidence intervals estimated in the data. The last three rows of the table 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). The first of these removes private insurance, so the only source of insurance is the statutory UI schedule. The remaining two 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 predicts precautionary savings responses to changes in social insurance that exceed the empirical bound established in Section3. The present-biased model, on the other hand, delivers predictions that lie within the estimated confidence intervals. The empirical evidence implies that the precautionary channel is weaker than standard models assume — with direct consequences for the stabilization properties of UI, which we now explore.

4 Implications for Stabilization Policy

We now turn to the implications of dampened precautionary saving for stabilization policy. As a stabilization tool, UI 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 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).

4.1 Closing the Model

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. But to understand the role of UI for stabilization policy, we need to extend the model. We now specify nominal rigidities, cyclical properties of the labor market, and the behavior of fiscal and monetary authorities. Last, we estimate the new parameters with Bayesian methods.

Wage Stickiness. Stabilization policy stimulates 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, extending the results in Auclert et al. (2018) to a setting with unemployment. Derivations can be found in the appendix.

First, define the aggregate labor wedge

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

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:

$$\pi_t = \frac{\delta}{1+\delta\iota}\mathbb{E}_t\pi_{t+1} + \frac{\iota}{1+\delta\iota}\pi_{t-1} + \frac{(1-\theta)(1-\delta\theta)}{\theta(1+\delta\iota)} d\log\mu_t + \varepsilon_t^w, \tag{10}$$

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 how 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 cyclicality 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 λ% 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

$$ds_t = -\eta \cdot df_t, \tag{11}$$

where the parameter η captures the relative importance of separation versus job finding rates in driving employment fluctuations.[30]

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

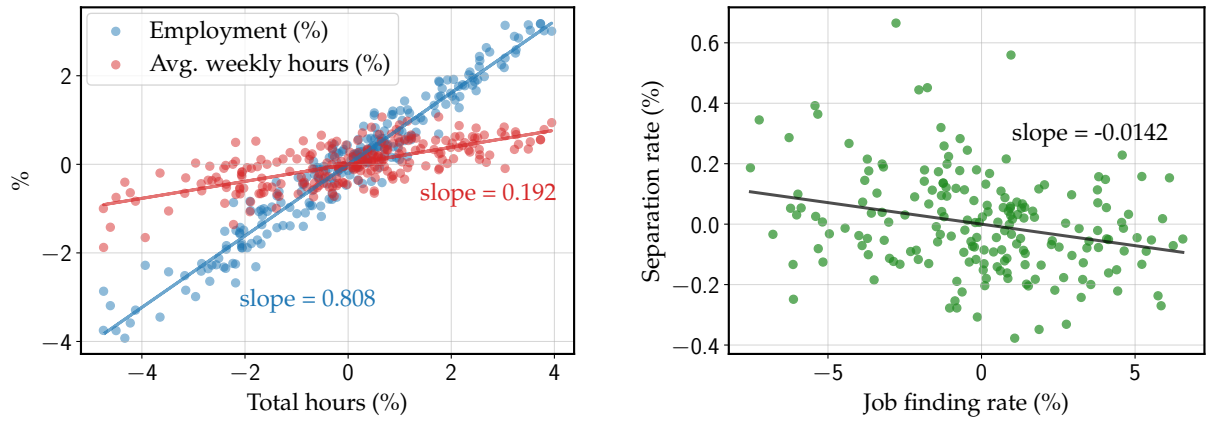
$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t}.$$

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.

[30] Ravn and Sterk (2021) studies how feedback between unemployment risk, separation and finding rates, and aggregate consumption arise endogenously in a model with search and matching.



(a) Hours worked and employment

(b) Separation and job finding rates

Figure 5: Parameters governing labor demand variation.

Calibration and Bayesian Estimation. In section 2.2, we calibrated parameters affecting the steady state of the model. To determine values for the (new) parameters governing the dynamic response of the economy to aggregate shocks, we mix calibration and Bayesian estimation.

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 $\lambda = 0.817$.

We use data from Shimer (2012) to calibrate η, the degree to which job finding rates drive employment fluctuations as opposed to separation rates. 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 OLS 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).[31]

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). 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. The appendix provides details on the data and the

[31] Auclert 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 $R_t \approx \tau_t$.

estimation results (e.g., posterior distributions and impulse response functions to business cycle shocks). Armed with the estimated model, we can entertain our first policy counterfactual.

4.2 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,[32] 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.[33]

Figure 6 shows how output responds. The effect is substantially larger in the exponential economy. 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:

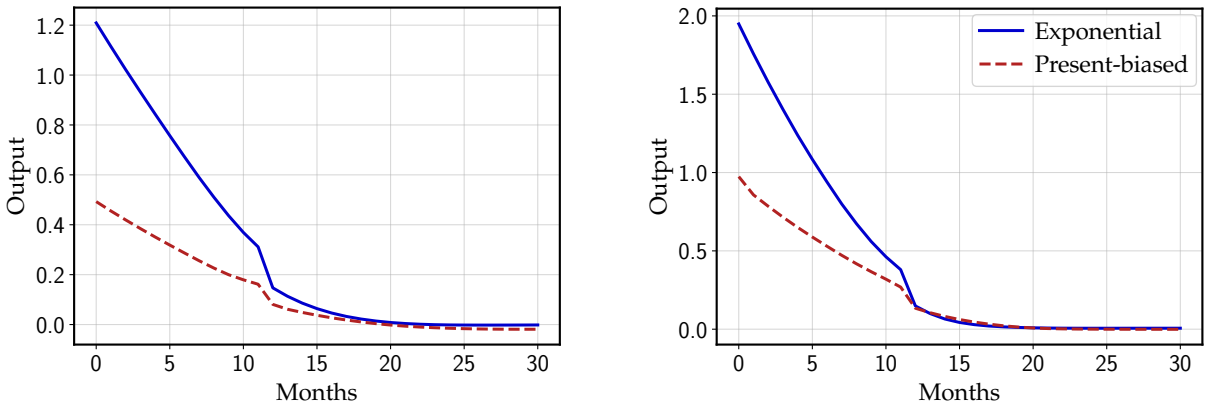
$$\text{multiplier} = \frac{\sum_{t=0}^{\infty}(1+r)^{-t}Y_t}{\sum_{t=0}^{11}(1+r)^{-t}},$$

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 0.78 for the exponential economy and 0.27 with present bias. For the increase in UI duration, they are 1.13 and 0.62, 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

[32] The 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.

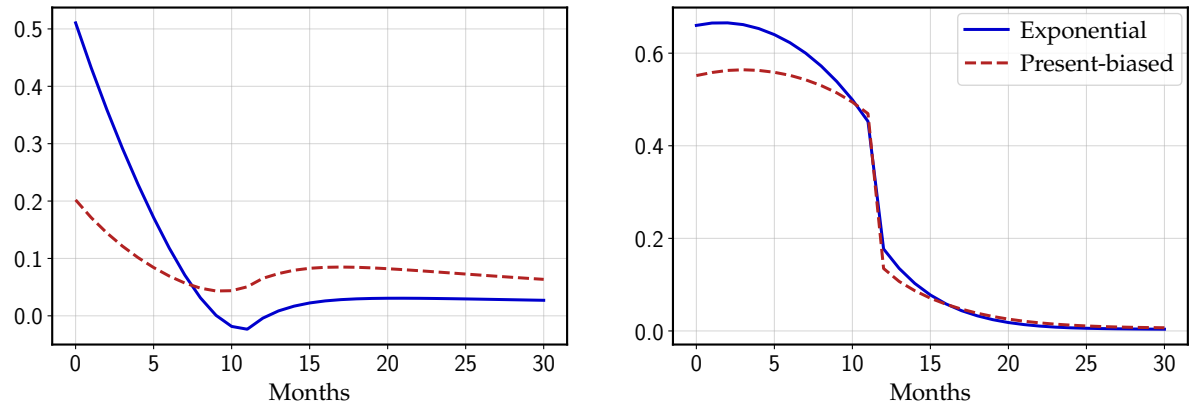
[33] This is without loss of generality, as any increase in spending due to changes in unemployment would be second-order, but we only solve for aggregate dynamics to a first-order approximation.



(a) Transfer to all unemployed

(b) Duration increase

Figure 6: Output response to UI extension



(a) Employed agents

(b) Unemployed agents

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

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.

4.3 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

	Increase in the standard deviation of $\log(C_t)$				Annualized r
	All shocks	Only monetary	Only markup	Only discount	
Exponential	9.7%	8.1%	13.2%	9.8%	3.10%
Present-biased	5.9%	4.7%	10.4%	6.0%	3.72%

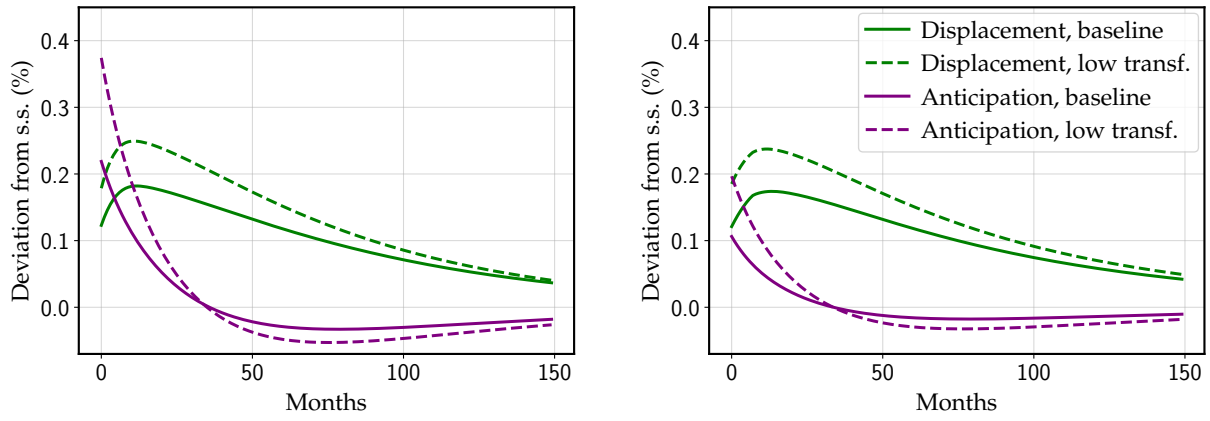
Table 7: The effects of reduced automatic stabilizers on aggregate consumption volatility.

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 7 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.

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 (11).

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



(a) Exponential

(b) Present-biased

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

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.[34]

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 anticipation 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 about 0.1 in the present-biased case and by 0.225 in the exponential case – nearly a two-fold increase, 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 shown in the appendix. 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.

The policy implications of weaker precautionary savings are likely robust to the specific behavioral mechanism (e.g., present bias, inattention, or overly optimistic beliefs about job finding).Ganong and Noel (2019) test these alternative behavioral mechanisms in a structural model of

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

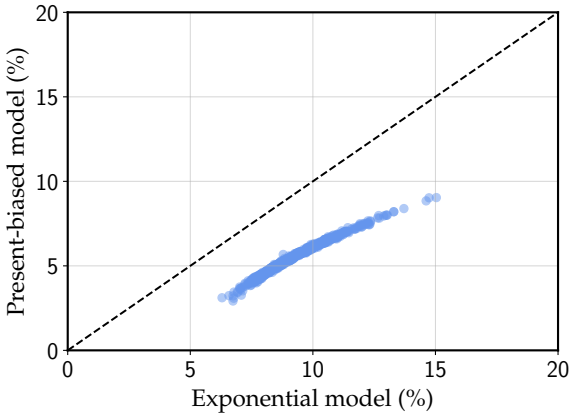


Figure 9: Increase in consumption volatility in the exponential and present-biased models for various parameter values

unemployment spending and find that present bias best accounts for the sharp consumption drop at benefit expiration. Given we show such paths are akin to sufficient statistics, our quantitative results provide a benchmark under the mechanism that best fits this moment. Other moments may be relevant, and future work should continue to explore how other behavioral assumptions may improve the link between model and data.

5 Conclusion

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 — our preferred design rules out precautionary responses larger than about 0.75% of annual income per 10 percentage-point increase in replacement rates, establishing an informative bound on the strength of the precautionary channel.

Next, we contrast our empirical findings to the predictions of two models: one with standard and another with present-biased preferences. By running the same TWFE design in model-simulated data, we find large precautionary savings responses in the exponential model that fall outside the estimated confidence intervals. This model's precautionary savings predictions are therefore at odds with the empirical evidence. 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. Our empirical design is precise enough to detect effects of the magnitude predicted by the standard exponential model, lending credibility to the comparison.

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 45%. 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.

Our work 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? More broadly, while present bias provides a parsimonious account of both the weak precautionary response and the spending drop at UI benefit expiration, other behavioral mechanisms could be informative about precautionary savings. Exploring these mechanisms and quantifying how policy prescriptions differ across them is an important agenda for future work.

References

Acharya, Sushant and Keshav Dogra, "Understanding HANK: Insights from a PRANK," *Econometrica*, 2020, *88* (3), 1113–1158.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 04 2021, *89* (1), 45–86.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton university press, 2009.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ , Matthew Rognlie, and Ludwig Straub, "The intertemporal keynesian cross," Technical Report, National Bureau of Economic Research 2018.

_ , _ , and _ , "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Technical Report, National Bureau of Economic Research 2020.

Auray, Stephane, David L Fuller, and Damba Lkhagvasuren, "Unemployment insurance take-up rates in an equilibrium search model," *European Economic Review*, 2019, *112*, 1–31.

Bardóczy, Bence, "Spousal insurance and the amplification of business cycles," 2022.

_ and Joao Guerreiro, "Unemployment Insurance in Macroeconomic Stabilization with Imperfect Expectations," Technical Report 2023.

Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri, "Uncertainty and Consumer Durables Adjustment," *The Review of Economic Studies*, 10 2005, *72* (4), 973–1007.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, *"Behavioral Household Finance"*, Vol. 1, Elsevier, 2018.

Bilbiie, Florin O and Xavier Ragot, "Optimal monetary policy and liquidity with heterogeneous households," *Review of Economic Dynamics*, 2021, *41*, 71–95.

Bitler, Marianne and Hilary W Hoynes, "The state of the safety net in the post-welfare reform era," Technical Report, National Bureau of Economic Research 2010.

Blanchard, Oliveir, Florence Jaumotte, and Prakash Loungani, "Labor Market Policies and IMF Advice in Advanced Economies During the Great Recession," *International Monetary Fund Staff Discussion Note*, 2013.

Blow, Laura, Martin Browning, and Ian Crawford, "Non-parametric Analysis of Time-Inconsistent Preferences," *The Review of Economic Studies*, 01 2021, *88* (6), 2687–2734.

Bornstein, Gideon and Sasha Indarte, "The Impact of Social Insurance on Household Debt," *Working paper*, 2023.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 1996, *34* (4), 1797–1855.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Campbell, John Y, "Household finance," *The journal of finance*, 2006, *61* (4), 1553–1604.

Carroll, Christopher D, "Buffer-stock saving and the life cycle/permanent income hypothesis," *The Quarterly journal of economics*, 1997, *112* (1), 1–55.

Carroll, Christopher D. and Andrew A. Samwick, "The nature of precautionary wealth," *Journal of Monetary Economics*, 1997, *40* (1), 41–71.

_ and _, "How Important is Precautionary Saving?," *The Review of Economics and Statistics*, 1998, *80* (3), 410–419.

_, Martin B. Holm, and Miles S. Kimball, "Liquidity constraints and precautionary saving," *Journal of Economic Theory*, 2021, *195*, 105276.

Carroll, Christopher, Karen E. Dynan, and Spencer Krane, "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets," *The Review of Economics and Statistics*, 2003, *85* (3), 586–604.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chaisemartin, Clément De and Xavier D'Haultfoeuille, "Difference-in-differences estimators of intertemporal treatment effects," Technical Report, National Bureau of Economic Research 2022.

Challe, Edouard, "Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 241–83.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions*," *The Quarterly Journal of Economics*, 12 2023, *139* (2), 891–936.

Cho, Daeha, "Unemployment risk, MPC heterogeneity, and business cycles," *Manuscript, University of Melbourne*, 2020.

Chodorow-Reich, Gabriel and John M Coglianese, "Unemployment insurance and macroeconomic stabilization," *Unemployment Insurance and Macroeconomic Stabilization." In Recession Ready, ed. Heather Boushey, Ryan Nunn, and Jay Shambaugh*, 2019.

_ **and Loukas Karabarbounis**, "The Cyclicality of the Opportunity Cost of Employment," *Journal of Political Economy*, 2016, *124* (6), 1563–1618.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, and Maarten van Rooij, "Consumption Uncertainty and Precautionary Saving," *The Review of Economics and Statistics*, 03 2020, *102* (1), 148–161.

Currie, Janet and Jonathan Gruber, "Saving babies: The efficacy and cost of recent changes in the Medicaid eligibility of pregnant women," *Journal of political Economy*, 1996, *104* (6), 1263–1296.

DellaVigna, Stefano, *"Structural Behavioral Economics"*, Vol. 1, Elsevier, 2018.

Dynan, Karen E., "How Prudent are Consumers?," *Journal of Political Economy*, 1993, *101* (6), 1104–1113.

Engen, Eric M and Jonathan Gruber, "Unemployment insurance and precautionary saving," *Journal of monetary Economics*, 2001, *47* (3), 545–579.

Erceg, Christopher J, Zoltan Jakab, and Jesper Lindé, "Monetary policy strategies for the European Central Bank," *Journal of Economic Dynamics and Control*, 2021, *132*, 104211.

Ericson, Keith Marzilli and David Laibson, *"Intertemporal Choice"*, Vol. 2, Elsevier, 2019.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri, "Firm-Related Risk and Precautionary Saving Response," *American Economic Review*, May 2017, *107* (5), 393â97.

_ **, Martin B Holm, and Gisle J Natvik**, "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Ferrière, Axelle and Gaston Navarro, "Fiscal management of aggregate demand: the effectiveness of labor tax credits," 2025.

Floden, Martin and Jesper Lindé, "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?," *Review of Economic dynamics*, 2001, *4* (2), 406–437.

Fuchs-Schündeln, Nicola and Matthias Schündeln, "Precautionary Savings and Self-Selection: Evidence from the German Reunification "Experiment"," *The Quarterly Journal of Economics*, 08 2005, *120* (3), 1085–1120.

Gabaix, Xavier, "A behavioral New Keynesian model," *American Economic Review*, 2020, *110* (8), 2271–2327.

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *American economic review*, 2019, *109* (7), 2383–2424.

_ , Fiona Greig, Max Liebeskind, Pascal Noel, Daniel M Sullivan, and Joseph Vavra, "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, (2021-19).

Gormley, Todd A and David A Matsa, "Growing out of trouble? Corporate responses to liability risk," *The Review of Financial Studies*, 2011, *24* (8), 2781–2821.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Guiso, Luigi and Tullio Jappelli, "Anatomy of Consumption Risk," CSEF Working Papers 732, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy July 2024.

_ , _ , and Daniele Terlizzese, "Earnings uncertainty and precautionary saving," *Journal of Monetary Economics*, 1992, *30* (2), 307–337.

Guvenen, Fatih, Alisdair McKay, and Ryan Conor, "A Tractable Income Process for Business Cycle Analysis," *Working Paper*, 2025.

_ , Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621–660.

Harris, Christopher and David Laibson, "Dynamic choices of hyperbolic consumers," *Econometrica*, 2001, *69* (4), 935–957.

Hodrick, Robert J and Edward C Prescott, "Postwar US business cycles: an empirical investigation," *Journal of Money, credit, and Banking*, 1997, pp. 1–16.

Japelli, Tulio and Luigi Pistaferri, *The Economics of Consumption*, Oxford University Press, 2017.

Juelsrud, Ragnar and Ella Wold, "The Importance of Unemployment Risk for Individual Savings," *Management Science*, 2024.

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Kekre, Rohan, "Unemployment insurance in macroeconomic stabilization," Technical Report, National Bureau of Economic Research 2021.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica*, 1990, *58* (1), 53–73.

Kuka, Elira, "Quantifying the benefits of social insurance: unemployment insurance and health," *Review of Economics and Statistics*, 2020, *102* (3), 490–505.

Laibson, David, "Golden Eggs and Hyperbolic Discounting*," *The Quarterly Journal of Economics*, 05 1997, *112* (2), 443–478.

_ , Peter Maxted, and Benjamin Moll, "Present bias amplifies the household balance-sheet channels of macroeconomic policy," Technical Report, National Bureau of Economic Research 2021.

_ , Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices over the Lifecycle," *The Review of Financial Studies*, 07 2024, p. hhae035.

Lusardi, Annamaria, "On the Importance of the Precautionary Saving Motive," *The American Economic Review*, 1998, *88* (2), 449–453.

Maxted, Peter, "Present bias in consumption-saving models: A tractable continuous-time approach," Technical Report, Mimeo 2020.

McKay, Alisdair, "Time-varying idiosyncratic risk and aggregate consumption dynamics," *Journal of Monetary Economics*, 2017, *88*, 1–14.

_ and Ricardo Reis, "The role of automatic stabilizers in the US business cycle," *Econometrica*, 2016, *84* (1), 141–194.

_ and _ , "Optimal automatic stabilizers," *The Review of Economic Studies*, 2021, *88* (5), 2375–2406.

_ , Emi Nakamura, and Jón Steinsson, "The power of forward guidance revisited," *American Economic Review*, 2016, *106* (10), 3133–58.

Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa, "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias," *American Economic Review*, January 2021, *111* (1), 324â63.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, March 1999, *89* (1), 103–124.

Ravn, Morten O and Vincent Sterk, "Macroeconomic fluctuations with HANK & SAM: An analytical approach," *Journal of the European Economic Association*, 2021, *19* (2), 1162–1202.

Schaab, Andreas, "Micro and macro uncertainty," *Available at SSRN 4099000*, 2020.

Shimer, Robert, "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 2012, *15* (2), 127–148.

Taylor, John B, "Discretion versus policy rules in practice," in "Carnegie-Rochester conference series on public policy," Vol. 39 Elsevier 1993, pp. 195–214.

Vimercati, Riccardo Bianchi, Martin S Eichenbaum, and Joao Guerreiro, "Fiscal policy at the zero lower bound without rational expectations," Technical Report, National Bureau of Economic Research 2021.

Werning, Iván, "Incomplete markets and aggregate demand," Technical Report, National Bureau of Economic Research 2015.