

Precautionary saving and aggregate demand

EDOUARD CHALLE

Department of Economics, Ecole Polytechnique and CREST (CNRS)

JULIEN MATHERON

Research Division, Banque de France

XAVIER RAGOT

Department of Economics, Paris School of Economics

JUAN F. RUBIO-RAMIREZ

Department of Economics, Emory University and Federal Reserve Bank of Atlanta

We construct, and then estimate by maximum likelihood, a tractable dynamic stochastic general equilibrium model with incomplete insurance and heterogeneous agents. The key feature of our framework is that cross-sectional heterogeneity remains finite dimensional. The solution to the model thus admits a state-space representation that can be used to recover the distribution of the model's parameters. Household heterogeneity expands the set of observables to cross-sectional moments available at the business-cycle frequency (in addition to the

Edouard Challe: edouard.challe@gmail.com

Julien Matheron: julien.matheron@banque-france.fr

Xavier Ragot: xavier.ragot@gmail.com

Juan F. Rubio-Ramirez: juan.rubio-ramirez@emory.edu

We are grateful to Susanto Basu, Morten Ravn, and Jose-Victor Rios Rull for their feedback on an earlier version of this paper. We also benefited from comments by conference participants at the 2013 (Micro Heterogeneity) and 2014 (Impulse and Propagation) NBER Summer Institutes, the 2013 NBER EFSF Workshop on DGSE Models, the 2014 and 2015 SED Meetings, the 2014 ASSA–AEA Meetings, the 2014 ENSAI Economic Day, the 2014 ES North American Meetings, the INET–Cambridge Conference on Aggregate Demand, the Labor Market and Macroeconomic Policy (September 2014), the 2014 Hydra Workshop on Dynamic Macroeconomics, the 3rd AMSE/BDF Labor Market Conference, the 2015 CREI/IFW/Richmond Fed Conference on New Developments in the Macroeconomics of the Labor Markets, as well by seminar participants at IAE/Barcelona GSE, IMF, EIEF, CEPREMAP, Ecole Polytechnique, Tinbergen Institute, Toulouse School of Economics, San Francisco Fed, USC, UCLA, UC Davis, UC Irvine, University of Warwick, University of Cambridge, and University of Konstanz. Edouard Challe acknowledges financial support from Investissements d'Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) and Chaire FDIR. Xavier Ragot acknowledges financial support from the European Commission FP7-612796 MACFINROBODS. Juan F. Rubio-Ramirez acknowledges financial support from National Science Foundation Grant SES-1227397, Fondation Banque de France pour la Recherche, the Institute for Economic Analysis (IAE), the Programa de Excelencia en Educacion e Investigacion of the Bank of Spain, and the Spanish Ministry of Science and Technology ECO2011-30323-c03-01. The views expressed herein are those of the authors and should not be interpreted as reflecting those of the Federal Reserve Bank of Atlanta, the Federal Reserve System, the Banque de France, the Eurosystem, or OFCE.

usual macro and monetary time series). Incomplete insurance gives rise to a precautionary motive for holding wealth that propagates aggregate shocks via (i) a *stabilizing* aggregate supply effect, working through the supply of capital, and (ii) a *destabilizing* aggregate demand effect coming from the feedback loop between unemployment risk and precautionary saving. Using the estimated model to measure the contribution of precautionary savings to the propagation of recent recessions, we find strong aggregate demand effects during the Great Recession and, to a lesser extent, during the 1990–1991 recession. In contrast, the supply effect at least offsets the demand effect during the 2001 recession.

KEYWORDS. DSGE, incomplete insurance, heterogeneous agents, Bayesian estimation.

JEL CLASSIFICATION. C32, E12, E21, E52.

“A rational expectations equilibrium is a likelihood function.” Thomas J. Sargent (in Evans and Honkapohja (2005)).

A rational expectations equilibrium is a likelihood function: given preferences and technologies, if the aggregate shocks have a distribution, then there is a likelihood function on the structural parameters that obeys the cross-equation restrictions implied by the model. Following this insight, a growing number of researchers have constructed medium-scale New Keynesian models with enough shocks and wedges to fit the data well, which they have estimated using likelihood-based procedures.¹ Likelihood-based procedures have two advantages with respect to traditional calibration approaches. First, they deliver estimates of the paths of the shocks that explain the data and, hence, it is possible to construct counterfactuals. Second, they provide a sense of the parameter uncertainty surrounding the estimates. A common feature of this line of work is the assumption of *perfect insurance* against idiosyncratic income shocks. The reason for this is that imperfect insurance typically generates enormous ex post heterogeneity among agents, which existing solution methods cannot handle without drastically restricting the set of aggregate shocks and endogenous state variables. As a consequence, imperfect-insurance models cannot be estimated by maximum likelihood; rather, they are calibrated by some method of moments using only a subset of the moment conditions.²

In this paper, we formulate a medium-scale, New Keynesian dynamic stochastic general equilibrium model with imperfect insurance, whose solution admits a finite-dimensional state-space representation that can be used to compute the likelihood function. The property that the state-space remains finite-dimensional, despite imperfect insurance, follows from two basic assumptions. The first one concerns the extent of risk sharing; namely, we depart from perfect insurance in a minimal way by assuming

¹Examples are Boivin and Giannoni (2006), Del Negro, Schorfheide, Smets, and Wouters (2007), Ireland (2004), Lubik and Schorfheide (2004), Canova and Sala (2009), DeJong, Ingram, and Whiteman (2000), Schorfheide (2000), Otrok (2001), Smets and Wouters (2007), Justiniano, Primiceri, and Tambalotti (2010), Aruoba and Schorfheide (2011).

²Examples are Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007), Davila, Hong, Krusell, and Ríos Rull (2012).

that households belong to large, representative “families,” within which full risk sharing takes place between *employed* members only, while the unemployed receive unemployment benefits provided by the government. Our second assumption is that the debt limit faced by households is tighter than the “natural” debt limit, that is, the maximum amount that a household can borrow while being able to repay (and always enjoy positive consumption) in the worst possible income history (see, e.g., [Aiyagari \(1994\)](#)). We show that under these two assumptions—partial risk sharing and tight debt limit—the wealth distribution converges to a distribution with a finite number of mass points, which in turn implies that the aggregate state itself remains finite dimensional.³ Incidentally, one attractive feature of our approach is to make it possible to include the time-series dimension of cross-sectional information into the likelihood estimation of the model, in addition to the usual macro and monetary time series. While household-level data have routinely been used to calibrate imperfect-insurance models (starting with [Krusell and Smith \(1998\)](#)), this information has not yet been used as observable variables when estimating medium-scale New Keynesian models. A crucial step in that direction was made in papers like [Iacoviello \(2008\)](#) or [Justiniano, Primiceri, and Tambalotti \(2015\)](#), who calibrate such models to match some key features of the wealth distribution. We push this approach further by using quarterly household-level data on consumption dispersion at the estimation stage, in addition to lower-frequency household-level data at the calibration stage.⁴

Our general framework significantly expands the set of macroeconomic questions that can be investigated via structural, likelihood-based estimation, namely, to any issue where imperfect insurance against idiosyncratic shocks is likely to matter. In the present paper we illustrate our approach by focusing on one such issue: the way households’ “precautionary saving” behavior—their rational savings response in the face of imperfect insurance—propagates aggregate shocks. We focus on precautionary savings against unemployment risk, the main source of time-varying idiosyncratic risk at the business-cycle frequency. While there may be other sources of business-cycle driven change in idiosyncratic risk (e.g., wage risk), those would almost certainly magnify the response of the precautionary motive and hence strengthen the effects that we are after. To capture the main channels by which precautionary saving may affect outcomes, our framework combines three basic frictions: (i) nominal rigidities (in prices and wages), (ii) labor-market frictions, and (iii) imperfect insurance against idiosyncratic unemployment risk. All three frictions are known, even in isolation, to capture some important features of the business cycle. Importantly, their interactions give rise to a feedback loop between precautionary savings and aggregate demand: following aggregate shocks

³[Challe and Ragot \(2014\)](#) construct and calibrate a real business-cycle (RBC) model with imperfect insurance that also features a finite-dimensional wealth distribution, using a period utility function that is linear above a threshold (rather than assuming a form of partial cross-household insurance). The framework that they use ignores aggregate demand effects by construction, and is not suitable for structural estimation.

⁴More specifically, we estimate the model using as an observable a moment of the cross-sectional distribution of consumption (constructed from the U.S. Consumption Expenditure Survey). In a similar spirit, but focusing on firms’ capital structure rather than household heterogeneity, [Ajello \(2014\)](#) uses firm-level data on the capital structure at the estimation stage of a dynamic stochastic general equilibrium (DSGE) model with credit frictions.

that lower demand, job creation is discouraged, unemployment persistently rises, and hence so does idiosyncratic unemployment risk. Imperfectly insured households rationally respond to this rise in idiosyncratic unemployment risk by increasing precautionary wealth, thereby cutting consumption and degrading demand even more. This “aggregate demand” effect of time-varying precautionary savings is, however, usually not the only one at work in economies with imperfect insurance. As is now well understood, time-varying precautionary savings also have an “aggregate supply” effect that tends to *reduce*, not increase, aggregate volatility. Indeed, in a typical recession, as unemployment risk rises, imperfectly insured households save more (for precautionary purposes) than they would if they were perfectly insured. These additional savings lower the equilibrium interest rate relative to the perfect-insurance benchmark, which tends to limit the contraction in investment and the capital stock. Conversely, the reduction in unemployment risk in a typical boom leads to a fall in precautionary savings that raises the equilibrium interest rate and lowers the demand for capital, relative to the perfect-insurance benchmark. The aggregate supply effect of precautionary savings against unemployment risk thus tends to smooth fluctuations in investment, capital, and, ultimately, output (see, e.g., [Krusell and Smith \(1998\)](#)). Hence, in the presence of both the aggregate demand and supply effects of precautionary savings, determining which effect dominates, and hence whether time-varying precautionary savings ultimately makes the economy more or less responsive to aggregate shocks, becomes an empirical question. Our framework allows us to incorporate both effects (see [Krueger, Mitman, and Perri \(2015\)](#)) and to measure their relative strength from the data.

Once the joint posterior distribution of the structural parameters of the model has been recovered, we ask whether the precautionary motive mattered in the propagation of the recent U.S. recessions, including the Great Recession? In these instances, has the aggregate demand effect dominated the aggregate supply effect, making the precautionary motive inherently destabilizing? To answer these questions, we extract the aggregate shocks that affected the U.S. economy during these periods and then feed them into a counterfactual perfect-insurance model; hence, the precautionary motive due to imperfect insurance is not present by construction. For the Great Recession, we find evidence of a powerful feedback loop between idiosyncratic unemployment risk and consumption demand, so that the aggregate demand effect largely dominates the aggregate supply effect (not only does the precautionary motive significantly amplify the fall in aggregate consumption, the latter also feeds back to adverse labor-market conditions). We find qualitatively similar, though quantitatively less important, amplification effects during the 1990–1991 recession. In contrast, we find no evidence of strong aggregate demand effects during the 2001 recession; if anything, the supply effects dominate the aggregate dynamics (that is, there is less aggregate volatility with the precautionary motive than without).

Our analysis relates to several strands of the business-cycle literature. Sticky-price models emphasize the role of aggregate demand as a key driver of the business cycle (see, e.g., [Christiano, Eichenbaum, and Evans \(2005\)](#), [Galí \(2010\)](#), [Smets and Wouters \(2007\)](#), [Woodford \(2003\)](#)). These models have recently been extended to incorporate labor-market frictions; see [Walsh \(2005\)](#), [Trigari \(2009\)](#), [Blanchard and Galí \(2010\)](#),

Gertler, Sala, and Trigari (2008), Heer and Maussner (2010), Leduc and Liu (2014), and Galí (2010) for a survey. We relax the perfect-insurance assumption from this framework.

Krusell, Mukoyama, and Sahin (2010), Nakajima (2012) and, more recently, Kehoe, Midrigan, and Pastorino (2014) analyze imperfect-insurance models with search frictions wherein the idiosyncratic unemployment risk faced by households is endogenized through firms' job creation policy. These models assume flexible prices, implying that only the aggregate supply effect is operative.

Other papers combine nominal frictions with imperfect insurance, but as in Krusell and Smith (1998), treat labor-market flows as exogenous constraints on labor supply. This, by construction, rules out any feedback from aggregate demand to unemployment risk, which is the key amplification mechanism in our model. This class of models includes Guerrieri and Lorenzoni (2011), who study the impact of a tightening of the borrowing constraint, Oh and Reis (2012) and McKay and Reis (2013), who study the impact of fiscal and transfer policies, and McKay, Nakamura, and Steinsson (2015), who examine the effectiveness of "forward guidance" at the zero lower bound.

Two papers consider the same frictions in goods, labor, and asset markets as we do: Gornemann, Kuester, and Nakajima (2012) and Ravn and Sterk (2013). There are important differences between these papers and ours, both in terms of focus and method. Gornemann, Kuester, and Nakajima (2012) are concerned with the redistributive impact of monetary policy shocks. The authors thus construct an imperfect-insurance model with large-dimensional cross-sectional heterogeneity and show that an increase in the policy rate raises income and wealth inequalities, consistent with the empirical findings of Coibion, Gorodnichenko, Kueng, and Silvia (2012). Ravn and Sterk (2013) study how an exogenous shock to the job separation rate can explain the depth and length of the Great Recession. The latter paper illustrates the feedback loop between unemployment risk and aggregate demand, but it has no capital, and hence the aggregate supply effect of precautionary savings is shut down. In contrast to both contributions, we construct the likelihood function and estimate, rather than calibrate, our model.⁵

Our interest in the aggregate demand effect of time-varying precautionary savings is shared by several recent theoretical contributions, most notably Rendhal (2014), Beaudry, Galizia, and Portier (2014), and Heathcote and Perri (2014). Rendhal (2014) shows how a zero lower bound problem coupled with labor-market frictions and rigid nominal wages can cause the economy to fall into a liquidity trap. Beaudry, Galizia, and Portier (2014) show that when the economy has excess capital, precautionary savings against idiosyncratic unemployment risk may cause a demand shortage. The authors' approach is closely related to Heathcote and Perri (2014), who show that the feedback loop between aggregate demand and idiosyncratic unemployment risk may lead to multiple equilibria. Like Beaudry, Galizia, and Portier (2014) and Heathcote and Perri (2014), our paper focuses on the interactions between households' wealth and idiosyncratic unemployment risk, although the specific mechanism by which this occurs in our model is

⁵Den Haan, Rendahl, and Riegler (2015) identify an alternative mechanism generating a feedback loop between unemployment risk and aggregate demand—based on nominal wage rigidities (but flexible prices) and the possibility of hoarding (zero-interest) cash for precautionary purposes—in addition to real assets. Moreover, their model is calibrated and not estimated.

different from theirs and is embedded into the standard sticky-price framework. Moreover, in contrast to all three papers, we design our model to estimate it and extract the strength of the unemployment risk–aggregate demand feedback loop from the data.

On the methodological side, we show how a first-order approximation to an incomplete-insurance model can capture time-varying precautionary savings and can be used for estimation. Our approach differs from the alternative approach of [Reiter \(2009\)](#) in that our framework endogenously generates a finite-dimensional state space, whereas the state space is infinite dimensional and approximated by a finite distribution in [Reiter \(2009\)](#). As a consequence, perturbation methods and likelihood-based estimation are particularly easy to use within our framework.

The rest of the paper is organized as follows. Section 1 describes the model, from agents' behavior to the definition of the recursive equilibrium. Section 2 shows how our assumptions lead to a collapse of the dimension of the state space, while preserving the precautionary motive. Section 3 estimates the model and evaluates its empirical performance. Section 4 discusses our counterfactual experiment and investigates the amplifying role of the precautionary motive during the last U.S. recessions. Section 5 offers some concluding remarks.

1. THE MODEL

1.1 *Model overview*

The model introduces imperfect insurance against time-varying idiosyncratic unemployment risk into a quantitative “New Keynesian” model with labor-market frictions. There are two household types: “workers” and “firm owners.” All households participate in a market for one-period nominal bonds, supply labor when employed, and transit between employment and unemployment. However, only firm owners own the capital stock as well as all firms. Idiosyncratic unemployment risk cannot be perfectly insured by workers, who also face a borrowing constraint (as in, e.g., [Krusell and Smith \(1998\)](#)). Such financial frictions will motivate employed workers' accumulation of assets for precautionary reasons. Hence, to the extent that the idiosyncratic unemployment risk is time-varying, so will be the amount of assets in the economy for precautionary reasons.

The production side has four types of firms, in the spirit of, for example, [Trigari \(2009\)](#) or [Heer and Maussner \(2010\)](#). Labor intermediaries hire labor from the households in a market with matching frictions (modeled as in [Mortensen and Pissarides \(1994\)](#)) and transform it into labor services. Competitive wholesale goods firms buy labor and capital services to produce wholesale goods that are then used as inputs by intermediate goods firms. Every intermediate good firm is the monopolistic supplier of the differentiated good it produces, but faces [Calvo \(1983\)](#)-type nominal frictions when setting nominal prices (as in, e.g., [Christiano, Eichenbaum, and Evans \(2005\)](#), [Smets and Wouters \(2007\)](#)). Finally, intermediate goods firms sell their goods to a competitive final goods sector, which aggregates them into a single final good that is ultimately used for consumption and investment, as well as utilization and vacancy posting costs. Even though intermediate goods firms set nominal prices, we express all prices in real terms

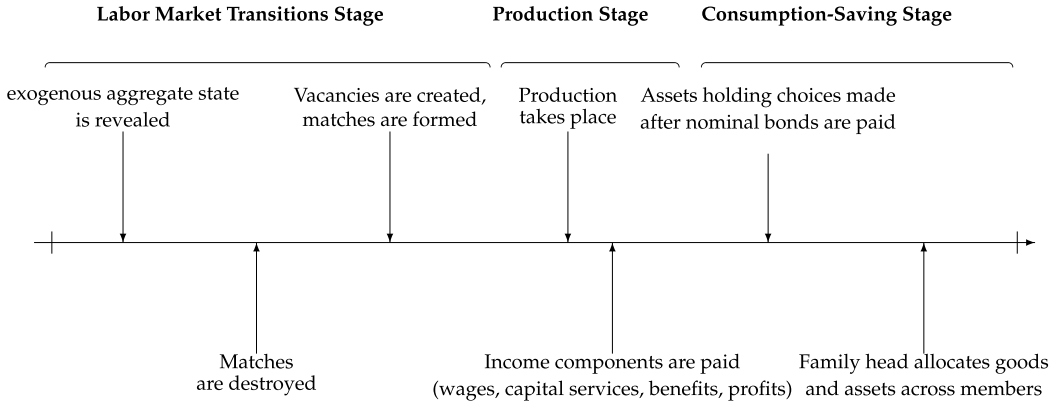


FIGURE 1. Model time line within a period.

using the final good as the numeraire. A central bank determines the nominal interest rate via a Taylor-like rule.

The timing of events within a period is as follows (see Figure 1). A period is divided into three stages: labor-market transitions, production, and consumption–saving. In the first stage, after the innovations to the exogenous aggregate state have been revealed, some existing employment relationships are destroyed, then hiring decisions are made, and new relationships are formed.⁶ In the production stage, production takes place and total income is shared between households in the form of net wages (for employed households, whether workers or firm owners), capital service payments (for firm owners, whether employed or unemployed), unemployment benefits (for unemployed households, whether workers or firm owners), and monopolistic profits (for firm owners, whether employed or unemployed). Finally, households’ assets holdings are determined in the consumption–saving stage, after the nominal bonds issued in the previous period have paid out.

We present the model recursively and use primes to denote the next period’s values. We call the aggregate state X , a vector containing all the relevant aggregate state variables in the model. We assume that all agents know the current value of X as well as its law of motion $X' = \Gamma(X, \epsilon')$, where ϵ' is the innovation to the exogenous aggregate state.⁷ The exogenous aggregate state is Markovian and includes a stochastic productivity trend e^z , where z drifts at rate $\mu_z \geq 0$. For expositional clarity, we summarize the content of X in Section 1.6 below, only after the presentation of the model has been completed. We first present the behavior of the households (Section 1.2), then present that of the firms (Section 1.3), and finally turn to the market-clearing conditions (Section 1.5)

⁶Our timing assumption allows a worker who is separated from the firm in the current period to be re-matched within the same period, in which case the worker does not effectively experience unemployment. This timing is consistent with the fact that labor-market flows occur at a frequency that is higher than the quarterly frequency (see, e.g., Walsh (2005), Galí (2010)).

⁷We will use “exogenous aggregate state” and “aggregate shocks” as interchangeable expressions.

and the definition of the equilibrium (Section 1.6). In the recursive representation of the model, all variables either belong to X or are a function of X . To save on notation when presenting the model we will only make this explicit for the value and policy functions. The rest of the relationships will be clarified when describing the equilibrium.⁸

1.2 Households

There is a unit mass of households, each of which is endowed with one unit of labor that is supplied inelastically during the production stage if the household is employed by the end of the labor-market transitions stage. All households are subject to idiosyncratic changes in their employment status: a share f (resp. s) $\in [0, 1]$ of the households that are unemployed (respectively, employed) before the labor-market transitions stage will be employed (respectively, unemployed) at the end of that stage. We refer to f and s as the job-finding and job-loss rates.

There are two types of households: there is a measure $\Omega \in [0, 1]$ of workers (indexed by W henceforth) and a measure $1 - \Omega$ of firm owners (indexed by F). All households have the same period utility function $u(c - h\mathbf{c}) = \lim_{\tilde{\sigma} \rightarrow \sigma} \frac{(c - h\mathbf{c})^{1 - \tilde{\sigma}} - 1}{1 - \tilde{\sigma}}$, with $\sigma > 0$, where c is consumption, \mathbf{c} is the level of consumption habits, and $h \in [0, 1)$ is a constant habit parameter. Workers and firm owners have subjective discount factors β^W and β^F , respectively, and we assume that

$$0 < \beta^W < \beta^F < e^{(\sigma-1)\mu_z},$$

where the second inequality states that workers are more impatient than firm owners and the third inequality ensures that the intertemporal utility of all households remains bounded.

Habits are external and defined as follows. We let \mathbf{c}^F be the common consumption habit of firm owners in the current period, and it is assumed to be equal to the average consumption of firm owners in the previous period. Regarding workers, we let $\mathbf{c}^W(N)$ denote the habit level of workers in the current period having been continuously unemployed for $N \in \mathbb{Z}_+$ periods. It is assumed to be equal to the average consumption of workers having experienced the same number of consecutive periods of unemployment ($= N$) in the previous period. For example, $\mathbf{c}^W(0)$ is the habit level of currently employed workers, and it is equal to the last period average consumption of employed workers. Similarly, $\mathbf{c}^W(1)$ is the consumption habit of an unemployed worker who was employed in the previous period, and it is equal to the average consumption of those workers who had lost their jobs in the previous period, and so on. This implies that all workers with the same N share the same habit level, while two workers with different N s in general have different habit levels.⁹

⁸In this section, we only describe the households' and firms' problems along with the aggregation and market-clearing conditions. All the optimality conditions are derived in the Technical Appendix, available in a supplementary file on the journal website, http://qeconomics.org/supp/714/code_and_data.zip.

⁹In our model, habits serve the usual purpose of producing an inertial response of aggregate consumption to aggregate shocks, which greatly improves the model's empirical fit; we then assume external rather

1.2.1 *Workers* The only assets that workers can trade are one-period nominal bonds. We let $\tilde{\mu}(a, N)$ denote the cross-sectional distribution of workers over individual assets $a \in \mathbb{R}$ and length of unemployment spell $N \in \mathbb{Z}_+$ at the beginning of the labor-market transitions stage. That is, $\tilde{\mu}(a, N)$ is the share of workers with nominal bond holdings less than or equal to a and having experienced exactly $N \geq 0$ consecutive periods of unemployment at that point in time. This distribution satisfies $\sum_N \int_a d\tilde{\mu}(a, N) = 1$. We let $\mu(a, N)$ characterize the distribution *after* the labor-market transitions stage, that is, at the beginning of the production stage. Note that $\tilde{\mu}(a, N)$ is an element of the aggregate state X , while $\mu(a, N)$ is a function of the aggregate state (since labor-market transition rates f and s are themselves functions of X). Finally, we let $\tilde{\mathbf{n}}^W \equiv \int_a d\tilde{\mu}(a, 0)$ and $\mathbf{n}^W = f(1 - \tilde{\mathbf{n}}^W) + (1 - s)\tilde{\mathbf{n}}^W$ denote the workers' employment rates at the beginning and the end of the labor-market transitions stage, respectively. Obviously we have $\mathbf{n}^W = \tilde{\mathbf{n}}^W$, that is, employment after the labor-market transitions stage is the same as employment before the same stage in the next period.

Employed workers earn the net labor income $(1 - \tau)w$, where w is the real wage received by workers and τ is the social contribution rate. Unemployed workers earn the unemployment benefit $b^u e^z$, where b^u is a constant, while the presence of the stochastic trend e^z will ensure balanced growth. The unemployment insurance (UI) scheme is balanced in every period, that is,

$$\tau w \mathbf{n}^W = b^u e^z (1 - \mathbf{n}^W). \quad (1)$$

We adopt the usual family structure according to which every worker belongs to a “representative family,” with the family head making consumption and saving decisions to maximize the intertemporal welfare of all family members. There is a measure Ω of such families, and each family has measure 1. Crucially, we depart from the standard structure by the amount of insurance that is provided by the family. More specifically, we impose the following restriction on the amount of risk sharing.

ASSUMPTION 1. *The family head can only transfer assets to workers with $N = 0$.*

Assumption 1 restricts the amount of risk sharing that can take place within the family: risk sharing may only take place across employed workers. Note that this implies that there can be no direct transfers either from employed to unemployed workers or across unemployed workers. The unemployed workers are taken charge of by the UI scheme and may also sell some of their assets to provide for current consumption. This risk-sharing arrangement preserves the notion that workers only enjoy imperfect insurance,

than internal habits because the former are much easier to handle than the latter, especially with incomplete markets and heterogenous agents. Note that in the representative agent model, when habits are external the consumption benchmark is just average consumption. By the same token, in two-agent “spender-saver” models, an agent's external consumption benchmark is specific to the group to which the agent belongs. Our specification is a natural generalization to the case with an arbitrary number of groups of agents, where the reference consumption level is (just as in the two previous cases) the average consumption in the previous period of the group to which the household currently belongs.

thereby preserving the precautionary motive for holding assets. As we will see later, Assumption 1 is key in making the model tractable.¹⁰

Imperfect insurance within every family implies that there are family-level, cross-sectional distributions of workers over individual assets and length of unemployment spell. Let $\tilde{\mu}(a, N)$ and $\mu(a, N)$ denote those distributions at the beginning of the labor-market transitions and consumption-saving stages, respectively, and let $\tilde{n}^W = \int_{\mathbb{R}} d\tilde{\mu}(a, 0)$ and $n^W = f(1 - \tilde{n}^W) + (1 - s)\tilde{n}^W$ denote the corresponding employment rates.¹¹ The family head, who cares equally for all the members, solves

$$V^W(\mu, X) = \max_{(a^{W'}(a, N), c^W(a, N))_{(a, N) \in \mathbb{R} \times \mathbb{Z}_+}} \left\{ \sum_N \int_a u(c^W(a, N) - hc^W(N)) d\mu(a, N) + \beta^W \mathbb{E}_{\mu, X} V^W(\mu', X') \right\}$$

subject to the borrowing limit $a^{W'}(a, N) \geq \underline{a}e^z$ and the budget constraints

$$a^{W'}(a, N) + c^W(a, N) = \mathbf{1}_{N=0}(1 - \tau)w + \mathbf{1}_{N \geq 1}b^u e^z + (1 + r)a, \quad \text{for } (a, N) \in \mathbb{R} \times \mathbb{Z}_+,$$

where

$$1 + r = (1 + \mathbf{R}_{-1})e^{\varphi_c} / (1 + \pi) \quad (2)$$

is the real return on nominal bond holdings. In the latter expression, \mathbf{R}_{-1} is the last-period nominal policy rate, π is realized inflation, and $\varphi_c (\in X)$ is a “risk premium” shock that drives a wedge between the gross nominal policy rate and the actual gross nominal return on bonds held by households (see, e.g., Smets and Wouters (2007)). Finally, $\mathbf{1}_{N=0}$ is an indicator function equal to 1 if $N = 0$ and 0 otherwise, and $\mathbf{1}_{N \geq 1} = 1 - \mathbf{1}_{N=0}$.

A few remarks are in order here. First, the family head computes $\mathbb{E}_{X, \mu} V^W(\mu', X')$ with the knowledge of (μ, X) and their laws of motion. In the case of μ , the law of motion involves the transition from μ to $\tilde{\mu}'$ (via asset choices) and then from $\tilde{\mu}'$ to μ' (via labor-market transitions). Second, the presence of a family-level, cross-sectional distribution of workers over individual assets and length of unemployment spell, μ , requires the family head to choose different levels of consumption c^W and assets $a^{W'}$ depending on each member’s individual state (a, N) . Third, according to Assumption 1, perfect insurance does take place between employed workers, and since those workers are symmetric, it is optimal to equalize both assets and consumption equally among them. It follows that $\mu(a, N)$ has a unique mass point in a for $N = 0$. We call this mass point \mathcal{A} ,

¹⁰The described family structure provides some partial insurance to workers, and thus increases their welfare ex ante, relative to a situation where there is no partial risk sharing among employed workers. We have checked this property quantitatively: workers are better off if they join the family (and thus benefit from partial insurance), instead of staying outside the family and self-insuring with their own resources. The methodology is described in the Technical Appendix.

¹¹These are the family-level counterparts of $\tilde{\mathbf{n}}^W$ and \mathbf{n}^W defined above, just as $\tilde{\mu}$ and μ are the family-level counterparts of $\tilde{\mu}$ and μ .

and it is given by

$$A = \frac{(1 - s) \int_{\mathbb{R}} a \, d\tilde{\mu}(a, 0) + f \sum_{N \geq 1} \int_{\mathbb{R}} a \, d\tilde{\mu}(a, N)}{n^W}.$$

To be more specific, at the beginning of the labor-market transition stage, the family-level cross-sectional distribution of workers over individual assets and length of unemployment spell is $\tilde{\mu}$. The numerator is the total amount of assets that is pooled by the workers who are employed at the beginning of the consumption–saving stage. During that stage, a fraction $1 - s$ of the workers remain employed and bring $(1 - s) \int a \, d\tilde{\mu}(a, 0)$ assets. At the same time, a fraction f of the workers who are unemployed before the labor-market transitions stage find a job during that stage. This includes workers of different N , and by the law of large numbers, the total amount of assets that they bring is $f \sum_{N \geq 1} \int_a a \, d\tilde{\mu}(a, N)$.¹² The denominator n^W is the family-level employment rate, that is, the number of workers among whom assets are shared. The optimal policy functions for workers’ assets and consumption have as arguments both their individual state (a, N) and the aggregate state X , that is, $a^{W'} = g_{a^W}(a, N, X)$ and $c^W = g_{c^W}(a, N, X)$.

1.2.2 Firm owners Besides earning either wages or unemployment benefits and participating in the nominal bond market like the workers, firm owners accumulate capital and rent its services out, and they own all the firms (and thus receive all profits). They thus have access to two assets: one-period nominal bonds and the capital stock. In contrast to the workers, we assume that there is perfect insurance among the firm owners. More specifically, we assume that every firm owner belongs to a family whose head freely pools resources across all members and allocates consumption goods and assets so as to maximize the intertemporal utility of all family members. There is a measure $1 - \Omega$ of such families; each family is of measure 1, and full insurance implies that all family members within a family are symmetric. It follows that they consume and save the same regardless of their employment status. The fraction of employed members within every family of firm owners before and after the labor-market transitions stage are denoted by \tilde{n}^F and n^F , respectively. We thus have $n^{F'} = f'(1 - n^F) + (1 - s')n^F$ and $n^F = \tilde{n}^{F'}$. As before, these are family-level variables. The corresponding aggregate variables are denoted $\tilde{\mathbf{n}}^{F'}$ and \mathbf{n}^F .

The intertemporal utility of a family of firm owners at the beginning of the consumption–saving stage is given by

$$V^F(n^F, k, a^F, i, X) = \max_{a^{F'}, c^F, i', v, k'} \{u(c^F - hc^F) + \beta^F \mathbb{E}_{n^F, X}[V^F(n^{F'}, k', a^{F'}, i', X')]\},$$

where k and a^F are the family’s capital stock and one-period nominal bonds holdings, i is its investment in the previous period (which enters $V^F(\cdot)$ due to the presence of investment adjustment costs), c^F is consumption, and $v \in [0, 1]$ is the capital utilization rate. Firm owners face the budget constraint

$$c^F + i' + a^{F'} = w^F n^F + [r_k v - \eta(v)]k + (1 + r)a^F + Y, \tag{3}$$

¹²We assume that the law of large numbers is valid for a continuum. See Miao (2006) for a discussion.

where w^F is the real wage earned by firm owners, r_k is the real rental rate of capital services, and Y is profits from intermediate goods firms and labor intermediaries, both rebated to their owners as dividends (measured in units of the final good).^{13,14} The function $\eta(v)k$ is the real cost that the utilization rate entails in units of the final good, where $\eta(v)$ is such that $\eta(1) = 0$, $\partial\eta(v)/\partial v > 0$ and $\partial^2\eta(v)/\partial^2v > 0$. Finally, the capital stock evolves as

$$k' = (1 - \delta)k + e^{\varphi_i}(1 - S(i'/i))i', \quad (4)$$

where $\delta \in [0, 1]$ is the depreciation rate, φ_i is an investment-specific shock, and $S(\cdot)$ is an investment adjustment cost function satisfying $\partial^2S(\cdot)/\partial(i'/i)^2 > 0$ and $S(g_i) = \partial S(\cdot)/\partial(i'/i)|_{i'/i=g_i} = 0$, where g_i is the steady-state value of i'/i .

The family head maximizes intertemporal utility subject to (3) and (4) and taking as given the laws of motions for X and n^F . We focus on a symmetric equilibrium. Hence, all families of firm owners are identical, that is, the cross-sectional distribution of families over the family-level state vector (n^F, k, a^F, i) is degenerate. We may thus write the optimal policy functions as functions of X only, that is, $a^{F'} = g_{a^F}(X)$, $c^F = g_{c^F}(X)$, $i' = g_i(X)$, $v = g_v(X)$, and $k' = g_k(X)$. Importantly, the homogeneity of firm owners implies that a single pricing kernel serves to price all future profits paid by the firms in the economy. This pricing kernel is given by firm owners' intertemporal marginal rate of substitution (IMRS henceforth), that is,

$$M^{F'} = \beta^F \frac{u_c(c^{F'} - hc^F)}{u_c(c^F - hc^F)}, \quad (5)$$

where we have used the fact that $c^F = g_{c^F}(X) = \mathbf{c}^{F'} \in X'$ (i.e., current consumption determines next period's habit). For example, the bond Euler equation of firm owners is $\mathbb{E}_X[M^{F'}(1 + r')] = 1$.

1.3 Firms

There are four types of firms in the economy.

1.3.1 Final goods firms The final good y is produced by a continuum of identical and competitive firms that combine intermediate differentiated goods according to the production function

$$y = \left(\int_0^1 y_s^{(\theta-1)/\theta} ds \right)^{\theta(\theta-1)}, \quad (6)$$

where y_s is the input of intermediate good s and $\theta > 1$ is the cross-partial elasticity of substitution between any two intermediate goods. Let p_s denote the price of intermediate good s in terms of the final good. This price is taken as given by the final goods firms.

¹³Employed firm owners earn the net labor income $(1 - \tau)w^F$. Unemployed firm owners earn the unemployment benefit $b^{u,F}e^z$, where $b^{u,F}$ is a constant. Since the UI scheme is balanced in every period and we assume perfect insurance within every family of firm owners, we have that $\tau w^F \mathbf{n}^F = b^{u,F}e^z(1 - \mathbf{n}^F)$ and, hence, budget constraint (3).

¹⁴Firm ownership takes the form of a fully diversified portfolio of private (i.e., untraded) equity here. We could equivalently allow the flows of rents to be traded among firm owners.

The program of the representative final goods producer is thus

$$\max_{y, y_s} \left\{ y - \int_0^1 p_s y_s \, ds \right\}, \quad (7)$$

subject to (6). From the optimal choices of final goods firms, one can deduce the demand function faced by an intermediate goods firm s ($\in [0, 1]$):

$$y_s(p_s) = p_s^{-\theta} y. \quad (8)$$

The zero-profit condition for final goods producers implies that

$$\left(\int_0^1 p_s^{1-\theta} \, ds \right)^{1/(1-\theta)} = 1. \quad (9)$$

1.3.2 Intermediate goods firms Intermediate goods firm $s \in [0, 1]$ is the monopolistic supplier of the good it produces, by means of a linear production function with a fixed cost

$$y_s = x_s - \kappa_y e^z, \quad (10)$$

where x_s is the quantity of wholesale goods used in production, and $\kappa_y e^z$ is the fixed cost measured in units of the wholesale good. The term e^z is included to ensure the existence of a balanced growth path, to be defined below. Firm s 's current profit, measured in units of the final good, is given by

$$\Xi = (p_s - p_m) y_s - p_m \kappa_y e^z,$$

where p_m is the price of wholesale goods in terms of the final good (which is taken as given by intermediate goods firms).

Firm s chooses p_s to maximize the present discounted value of future profits, taking as given the demand curve (8). Following Calvo (1983), we assume that in every period every intermediate goods firm can be in one of the following two idiosyncratic states: either the firm can freely reoptimize its price or it cannot and simply rescales the existing price according to the indexation rule

$$\tilde{p}_s = \frac{(1 + \bar{\pi})^{1-\gamma_p} (1 + \pi_{-1})^{\gamma_p}}{1 + \pi} p_{s,-1}, \quad (11)$$

where $\gamma_p \in (0, 1)$ measures the degree of indexation to the most recently available final goods inflation measure, π_{-1} is final goods inflation in the previous period, $\bar{\pi}$ is the steady-state inflation rate, and $p_{s,-1}$ is the last period's relative price of the intermediate good s .¹⁵ The ex ante probability of each firm being able to reoptimize the price in the next period is constant and equal to $1 - \alpha \in [0, 1]$, irrespective of the time elapsed since the period in which the price of the firm was last revised.

¹⁵This indexation rule is the usual nominal-indexing rule $P_s = (1 + \bar{\pi})^{1-\gamma_p} (1 + \pi_{-1})^{\gamma_p} P_{s,-1}$, where the P_s s are nominal prices, but here are formulated in terms of units of the final good.

It follows from this price adjustment mechanism that the behavior of a firm can be described by two Bellman equations, corresponding to the two idiosyncratic states in which the firm can be. The value of a firm that is allowed to reset its price is given by $V^R(X)$ and only depends on the aggregate state. The value of a firm not allowed to reset its selling price and with last period's price $p_{s,-1}$ is denoted as $V^N(p_{s,-1}, X)$. The corresponding Bellman equations are

$$V^R(X) = \max_{p_s} \{ \Xi + \alpha \mathbb{E}_X [M^{F'} V^N(p_s, X')] + (1 - \alpha) \mathbb{E}_X [M^{F'} V^R(X')] \},$$

$$V^N(p_{s,-1}, X) = \Xi + \alpha \mathbb{E}_X [M^{F'} V^N(\tilde{p}_s, X')] + (1 - \alpha) \mathbb{E}_X [M^{F'} V^R(X')],$$

taking as given the demand curve (8), and where \tilde{p}_s is given by (11). This results in the optimal policy function $p^* = g_{p^*}(X)$ for price resetters.

We focus on a symmetric equilibrium wherein the solution to intermediate goods firms' problem is the optimal reset price common to all price resetting firms. After straightforward algebraic manipulations, the first-order conditions associated with the determination of the optimal reset price is $p^* = \frac{K}{F}$, where K and F are defined recursively as

$$K = \mu e^{\varphi_p} p_m y + \alpha \mathbb{E}_X \left[\frac{M^{F'} (1 + \pi')^\theta K'}{(1 + \bar{\pi})^{\theta(1-\gamma_p)} (1 + \pi)^{\theta\gamma_p}} \right],$$

$$F = y + \alpha \mathbb{E}_X \left[\frac{M^{F'} (1 + \pi')^{\theta-1} F'}{(1 + \bar{\pi})^{(\theta-1)(1-\gamma_p)} (1 + \pi)^{(\theta-1)\gamma_p}} \right],$$

where $\mu = \theta/(1 - \theta)$ and where we allow for exogenous variations in the mark-up through the shock $\varphi_p \in X$. This optimal reset price, together with the Calvo price setting mechanism, the zero profit condition (9), and the indexation rule (11), imply the law of motion for inflation

$$\pi = \frac{\alpha((1 + \bar{\pi})^{1-\gamma_p} (1 + \pi_{-1})^{\gamma_p})}{(1 - (1 - \alpha)(p^*)^{1-\theta})^{1/(1-\theta)}} - 1. \quad (12)$$

The price-setting mechanism generates a cross-sectional distribution over prices, since the selling price of a firm not reoptimizing its price depends on the time that has elapsed since the last time the price was reset. However, the price dispersion index $\Lambda \equiv \int_0^1 p_s^{-\theta} ds$ suffices to capture the relevant properties of the distribution, and it evolves according to the law of motion

$$\Lambda = (1 - \alpha)(p^*)^{-\theta} + \alpha \left(\frac{(1 + \bar{\pi})^{1-\gamma_p} (1 + \pi_{-1})^{\gamma_p}}{1 + \pi} \right)^{-\theta} \Lambda_{-1}, \quad (13)$$

where Λ_{-1} is the value of the index in the previous period.

1.3.3 Wholesale goods firms The wholesale good is produced by a continuum of identical and competitive firms. The representative wholesale goods firm produces with the

technology $y_m = \check{k}^\phi (e^z \check{n})^{1-\phi}$, $\phi \in (0, 1)$, where \check{n} and \check{k} denote labor and capital services. It solves

$$\max_{\check{n}, \check{k}} \{ p_m \check{k}^\phi (e^z \check{n})^{1-\phi} - Q\check{n} - r_k \check{k} \}, \quad (14)$$

where Q is the real unit price of labor services. The solution to (14) gives the optimal demands for factor services $\check{n} = g_{\check{n}}(X)$ and $\check{k} = g_{\check{k}}(X)$.

1.3.4 Labor intermediaries and labor-market flows Labor services are sold to wholesale goods firms by labor intermediaries, who hire labor from households in a market with search frictions. More specifically, at the beginning of the labor-market transition stage, a fraction $\rho(\varphi_s)$ of existing employment relationships are destroyed, where φ_s is a job-destruction shock. The workers who loose their jobs on that occasion enter the unemployment pool, where they join the workers who were already unemployed at the end of the previous period. At the same time, labor intermediaries post vacancies, at the unit cost $\kappa_v e^z$ in terms of the final good, where the term e^z is included to ensure the existence of a balanced growth path. One employed worker provides one unit of labor services, but a firm owner provides $\psi > 1$ units of labor services, and we refer to ψ as the *skill premium*.¹⁶ The values to the labor intermediary of a match with a worker and a firm owner are, respectively,

$$J^W = Q - w + \mathbb{E}_X[(1 - \rho')M^{F'}J^{W'}], \quad J^F = \psi Q - w^F + \mathbb{E}_X[(1 - \rho')M^{F'}J^{F'}]. \quad (15)$$

We assume that, when posting a vacancy, labor intermediaries cannot target a particular skill type. Labor intermediaries thus adjust vacancies until the expected payoff on a posted vacancy is equal to its cost, that is,

$$\lambda[\Omega J^W + (1 - \Omega)J^F] = \kappa_v e^z, \quad (16)$$

where λ is the economy-wide vacancy-filling rate.

Let $\tilde{\mathbf{n}} = \Omega \tilde{\mathbf{n}}^W + (1 - \Omega)\tilde{\mathbf{n}}^F$ denote the economy-wide employment rate before the labor-market transitions stage and $\mathbf{n} = \Omega \mathbf{n}^W + (1 - \Omega)\mathbf{n}^F$ is the same rate after the labor-market transitions stage. These two definitions imply that $\tilde{\mathbf{n}}' = \mathbf{n}$. The unemployment pool is made up of workers who are unemployed at the beginning of the labor-market transitions stage (in number $1 - \tilde{\mathbf{n}}$) as well as workers who were employed at the beginning of that stage but lost their job after the job-destruction shock (in number $\rho \tilde{\mathbf{n}}$). The matching technology produces m employment relationships using as inputs the unemployment pool and the aggregate number of vacancies v . This technology has the form

$$m = \bar{m}(1 - (1 - \rho)\tilde{\mathbf{n}})^\chi v^{1-\chi}, \quad (17)$$

where \bar{m} is a scaling parameter and $\chi \in (0, 1)$ the elasticity of m with respect to (w.r.t.) the size of the unemployment pool. Accordingly, the economy-wide job-finding and

¹⁶The extent of consumption dispersion across U.S. households cannot be entirely accounted for by assets dispersion: dispersion in labor income is needed in addition to dispersion in asset income. This is adequately captured by a skill premium (see, e.g., Challe and Ragot (2014)).

vacancy-filling rates are

$$f = \frac{m}{1 - (1 - \rho)\tilde{\mathbf{n}}} \quad \text{and} \quad \lambda = \frac{m}{v}. \quad (18)$$

Since the workers who are separated from the firms can be rematched within the period, the period-to-period job-loss rate s is given by

$$s = \rho(1 - f). \quad (19)$$

As usual, there are two equivalent ways to view labor-market flows. From the point of view of the households, employment dynamics are determined by the flows of job losers and job finders, that is, $\mathbf{n} = f(1 - \tilde{\mathbf{n}}) + (1 - s)\tilde{\mathbf{n}}$. From the point of view of the labor intermediaries, it follows from the natural process of job destruction and the intensity of vacancy postings, that is, $\mathbf{n} = (1 - \rho)\tilde{\mathbf{n}} + \lambda v$.

1.3.5 Wages The presence of labor-market frictions implies that there may exist a full bargaining set over which a labor intermediary and an employee (whether worker or firm owner) find it mutually profitable to be matched. Following Hall (2005), we assume that there are some rigidities in nominal-wage adjustment and we use the implied expression for the base real wage w . More specifically, the latter is given by¹⁷

$$w = \left(\frac{\mathbf{w}_{-1}}{1 + \pi} \right)^{\gamma_w} \left(\bar{w} e^{z + \varphi_w} \left(\frac{\mathbf{n}}{\bar{n}} \right)^{\psi_n} \right)^{1 - \gamma_w}. \quad (20)$$

In equation (20), \mathbf{w}_{-1} denotes last period's real wage rate, \bar{w} denotes a scale factor, γ_w denotes the degree of indexation to past wages, and ψ_n denotes the sensitivity of wages to the business cycle, here measured as the ratio of aggregate employment \mathbf{n} to its steady-state value \bar{n} . The wage equation is also perturbed by a wage shock φ_w , and is appropriately scaled by the technology shock e^z to ensure the existence of a balanced growth path. We assume that the wage premium w^F/w paid to firm owners is equal to the skill premium ψ (as would be the case in a competitive labor market). Finally, we assume that w lies within the appropriate bargaining set (this implies that w^F does too), and will verify that this condition holds over our sample once the model has been estimated.¹⁸

1.4 Central bank

The central bank is assumed to set the nominal interest rate R according to the rule (see, e.g., Christiano, Motto, and Rostagno (2014), Gust, Lopez-Salido, and Smith (2012), Guerrón-Quintana, Fernández-Villaverde, and Rubio-Ramírez (2010))

$$\log\left(\frac{1 + R}{1 + \bar{R}}\right) = \rho_R \log\left(\frac{1 + \mathbf{R}_{-1}}{1 + \bar{R}}\right) + (1 - \rho_R) \left[a_\pi \log\left(\frac{1 + \pi}{1 + \bar{\pi}}\right) + a_y \log\left(\frac{1 + g}{1 + \bar{g}}\right) \right] + \varphi_R, \quad (21)$$

¹⁷The corresponding nominal wage dynamics is recovered by multiplying both sides of (20) by nominal final goods prices and rearranging to eliminate $1 + \pi$.

¹⁸See the Technical Appendix for details. Note that with $w^F = \psi w$, we have $J^F = \psi J^W$, so that the free-entry condition reduces to $\lambda J^W = [\Omega \psi + (1 - \Omega)]^{-1} \kappa_v e^z$.

where \bar{R} is the steady-state nominal interest rate, $\rho_R \in (0, 1)$ is an interest rate smoothing parameter, (a_π, a_y) are the reaction coefficients to inflation and output growth, $g = y/\mathbf{y}_{-1} - 1$ is the growth rate of final output, where \mathbf{y}_{-1} is the last-period final output, and φ_R is a monetary-policy shock.

1.5 Market clearing

1.5.1 Labor services Recall from Section 1.2 that all households face the same labor-market transition rates (f, s) . Hence, in the steady state, the employment rates in every family of workers and firm owners are the same. Assuming that employment is symmetric at the beginning of the date-0 labor-market transition stage, by the law of large numbers they remain symmetric at every point in time, that is,

$$\tilde{n}^W = \tilde{n}^F = \tilde{\mathbf{n}}^W = \tilde{\mathbf{n}}^F \equiv \tilde{\mathbf{n}}, \quad n^W = n^F = \mathbf{n}^W = \mathbf{n}^F \equiv \mathbf{n}. \quad (22)$$

Because a matched firm owner provides ψ times more units of labor services than a worker, the total supply of labor services is $\Omega \mathbf{n}^W + (1 - \Omega)\psi \mathbf{n}^F = (\Omega + (1 - \Omega)\psi)\mathbf{n}$. Denoting by \check{n} firms' demand for labor services, market clearing requires

$$(\Omega + (1 - \Omega)\psi)\mathbf{n} = \check{n}. \quad (23)$$

1.5.2 Assets markets Recall that firm owners are symmetric and in measure $1 - \Omega$. Since each of them supplies νk units of capital services, the total supply of capital services is $(1 - \Omega)\nu k$. Thus, market clearing gives

$$(1 - \Omega)\nu k = \check{k}. \quad (24)$$

All the households may participate in the market for nominal bonds, which are in zero net supply. In symmetric equilibrium, at the end of the consumption–saving stage, all firm owners hold the same amount of assets $a^{F'}$, while workers hold different levels of assets depending on their individual state (a, N) . Clearing of the market for bonds requires

$$(1 - \Omega)a^{F'} + \Omega \sum_N \int_a a^{W'} d\mu = 0. \quad (25)$$

1.5.3 Goods markets The aggregate demand for final goods is made of total investment (by firm owners), the consumption of all households, as well as capital utilization costs (directly paid by firm owners) and vacancy costs (paid by the labor intermediaries). Again, taking into account workers' heterogeneity, we write the market-clearing condition as

$$(1 - \Omega)(c^F + i' + \eta(\nu)k) + \Omega \sum_N \int_a c^W d\mu + \kappa_\nu e^z \nu = y. \quad (26)$$

The intermediate goods sector demands one unit of wholesale goods for any unit of intermediate goods. Hence, the market-clearing condition for the wholesale goods

sector is

$$\int_0^1 x_s \, d\mathfrak{s} = y_m = \check{k}^\phi (e^z \check{n})^{1-\phi}. \quad (27)$$

The total demand for intermediate goods by the final goods sector is $\int_0^1 y_s(X, p_s) \, d\mathfrak{s} = \Lambda y$, where Λ evolves as shown in equation (13). The total supply of wholesale goods is equal to $\int_0^1 x_s \, d\mathfrak{s} - \kappa_y e^z$. Hence, the clearing of the market for intermediate goods requires, using equation (27),

$$\Lambda y = \check{k}^\phi (e^z \check{n})^{1-\phi} - \kappa_y e^z. \quad (28)$$

1.6 Aggregate state and equilibrium

We are now in a position to summarize the content of the aggregate state. Again, we are focusing on a symmetric equilibrium, where family-level variables are identical across families of workers and families of firm owners (e.g., $\mu = \boldsymbol{\mu}$ etc.). The aggregate state is then given by

$$X = \{\tilde{\boldsymbol{\mu}}(\cdot), k, a^F, i, \mathbf{c}^F, \mathbf{c}^W(N)_{N \in \mathbb{Z}_+}, a^e, \mathbf{R}_{-1}, \boldsymbol{\Lambda}_{-1}, \boldsymbol{\pi}_{-1}, \mathbf{y}_{-1}, \mathbf{w}_{-1}, \Phi\}, \quad (29)$$

where $\Phi \equiv \{z, \varphi_i, \varphi_c, \varphi_s, \varphi_R, \varphi_w, \varphi_p\}$ is the exogenous aggregate state.

DEFINITION 1. A *symmetric recursive equilibrium* is a set of value and policy functions, a set of prices, and labor-market flows such that the following statements hold:

(i) *Workers:* Given $r(X), w(X), \tau(X), b^u e^z, \mathbf{c}^W(N)_{N \in \mathbb{N}}, f(X)$, and $s(X)$, the value and policy functions $V^W(\boldsymbol{\mu}, X)$, $g_{a^W}(a, N, X)$, and $g_{c^W}(a, N, X)$ solve the workers' problem.

(ii) *Firm owners:* Given $r(X), r_k(X), w^F(X), \mathbf{c}^F, Y(X), f(X)$, and $s(X)$, the value and policy functions $V^F(n^F, k, a^F, i, X)$, $g_{a^F}(X)$, $g_{c^F}(X)$, $g_i(X)$, $g_v(X)$, and $g_k(X)$ solve the firm owners' problem.

(iii) *Final goods firms:* Given $p_s, s \in [0, 1]$, the demand for intermediate goods $y_s(p_s, X)$ is optimal from the point of view of final goods firms.

(iv) *Intermediate goods firms:* Given $p_m(X), y_s(p_s, X)$, and $M^F(X, X')$, the value functions $V^R(X)$ and $V^N(p_{s-1}, X)$ and the reset price $p^*(X)$ solve the problem of intermediate goods firms.

(v) *Wholesale goods firms:* Given $p_m(X), Q(X)$, and $r_k(X)$, the demand for labor and capital services $\check{n}(X)$ and $\check{k}(X)$ solves the problem of wholesale goods firms.

(vi) *Labor intermediaries:* Given $Q(X), w(X)$, and $M^F(X, X')$, the job values $J^W(X)$ and $J^F(X)$ are given by (15), the free-entry condition (16) determines the vacancy-filling rate $\lambda(X)$, and $m(X), f(X), v(X)$, and $s(X)$ are determined according to (17), (18), and (19).

(vii) *Profits:* The profit function $Y(X)$ results from the optimal decision of the intermediate goods firms and the labor intermediaries.

(viii) *Social contribution rate, real interest rate, stochastic discount factor, wages, and nominal interest rate:* Given $y(X)$, $\pi(X)$, and $b^\mu e^z$, the social contribution rate $\tau(X)$ is such that (1) holds; the real return on nominal bond holdings $r(X)$ follows (2); the stochastic discount factor $M^F(X, X')$ is given by (5); firm owners' wage $w^F(X)$ is equal to $\psi w(X)$, where $w(X)$ is given by (20); the nominal interest rate $R(X)$ is given by (21).

(ix) *Market clearing:* The market-clearing conditions (23)–(28) hold.

(x) *Laws of motion:* Given $p^*(X)$, inflation $\pi(X)$ and price dispersion $\Lambda(X)$ evolve according to (12) and (13), respectively; given $f(X)$, $s(X)$, and $g_{aw}(\cdot)$, the laws of motion from $\tilde{\mu}$ to μ and then from μ to $\tilde{\mu}'$ are given by

$$\begin{aligned} \tilde{\mu} \text{ to } \mu : & \begin{cases} \mu(a, 0) = f(X) \sum_{N \geq 1} \tilde{\mu}(a, N) + (1 - s(X))\tilde{\mu}(a, 0), \\ \mu(a, 1) = s(X)\tilde{\mu}(a, 0), \\ \mu(a, N) = (1 - f(X))\tilde{\mu}(a, N - 1), \quad \text{for } N \geq 2, \end{cases} \\ \mu \text{ to } \tilde{\mu}' : & \tilde{\mu}'(\hat{a}, N) = \int \mathbf{1}_{g_{aw}(a, N, X) \leq \hat{a}} d\mu(a, N), \quad \text{for } N \geq 0. \end{aligned}$$

(xi) *Habits:* Given $g_{c^F}(X)$ and $g_{c^W}(\cdot)$, tomorrow's habit level of a particular household type is equal to the average consumption of this type today, that is,

$$c^{F'} = g_{c^F}(X) \quad \text{and} \quad c^{W'}(N) = \int g_{c^W}(a, N, X) d\mu(a, N).$$

The solution of the model will rely on a linear approximation of its dynamics around a balanced growth path (BGP henceforth). We now provide a formal definition of the BGP and derive its key theoretical properties in the next section.

DEFINITION 2. A BGP is a symmetric recursive equilibrium where the following statements hold:

- (i) Innovations to the exogenous aggregate state (ϵ) are zero at every point in time; therefore, aggregate shocks are absent.
- (ii) The variables $w(X)$, $c^W(N)_{N \in \mathbb{N}}$, $w^F(X)$, c^F , $Q(X)$, $Y(X)$, and $\check{k}(X)$ all grow at rate μ_z .
- (iii) The variables $r(X)$, $r_k(X)$, $f(X)$, $s(X)$, $\lambda(X)$, $m(X)$, $\nu(X)$, $\check{n}(X)$, $R(X)$, $p_m(X)$, and $\pi(X)$ are constant.

2. EQUILIBRIUM DYNAMICS

We now study the equilibrium dynamics of the symmetric recursive equilibrium defined above. We will first show that under some assumptions, the cross-sectional distributions of workers over individual assets and length of unemployment spell in a BGP have finite support. We will then use this result to show that around the BGP the symmetric recursive equilibrium is summarized by a finite number of equilibrium conditions. Finally, we

will isolate the key determinants of workers' precautionary savings and show how they survive when the nonlinear system that characterizes the equilibrium is approximated at the first order.

2.1 *A cross-sectional distribution of workers with finite support*

We now derive our main theoretical results regarding the properties of the cross-sectional distributions of workers (over assets and length of unemployment spell) and workers' consumption–saving choices. We focus on the symmetric recursive equilibrium characterized in Section 1.6, so that the family-level distributions coincide with their aggregate counterparts, that is, $\tilde{\mu} = \tilde{\mu}$ and $\mu = \mu$ (at the beginning of time by assumption, and in every period by implication). We proceed in three steps. First, we show that, under Assumption 1, the cross-sectional distributions $\tilde{\mu}'$ and μ asymptotically tend toward distributions with countable supports. We then show that, as a consequence, workers' consumption–saving choices are summarized by a countable number of Euler conditions, which have exactly the same form as in standard heterogenous-agent models. Finally, we show that, under the additional assumption that the borrowing limit is tighter than the natural limit, the supports of $\tilde{\mu}'$ and μ are not only countable but also finite in and around any BGP. In what follow we state the relevant propositions and leave their proofs for the Appendix at the end of the paper.

PROPOSITION 1. *Under Assumption 1, (a) if the distribution $\tilde{\mu}(a, N)$ has a unique mass point in a for all N , then both $\tilde{\mu}(a, N)$ and $\mu(a, N)$ have a unique mass point in a for all N in the following periods; it follows that all workers with the same $N = 0, 1, \dots$ have the same levels of consumption $c^W(N, X)$ and end-of-period assets $a^{W'}(N, X)$. Additionally, (b) if $\tilde{\mu}(a, N)$ does not have this property, then for any $\check{N} \in \mathbb{Z}_+$, both $\tilde{\mu}(a, N)$ and $\mu(a, N)$ have a unique mass point in a for all $N \leq \check{N}$ after $\check{N} + 1$ periods.*

Proposition 1(a) states that, under Assumption 1, the countability of the support of the cross-sectional distribution of workers over individual assets and length of unemployment spell is preserved over time; that is, the distributions $\tilde{\mu}$ and μ can be characterized by a countable set of pairs (a, N) . By implication, all workers with the same N are indistinguishable: they enter the consumption–saving stage with identical assets $a^W(N)$, consume the same amount $c^W(N)$, and end the consumption–saving stage with the same assets $a^{W'}(N)$.

The intuition for this result is best conveyed in the context of an economy without aggregate shocks, with a zero-debt limit, and where workers are so impatient that in equilibrium they decide to hold at most very little precautionary wealth. The fact that all employed workers pool their assets at the beginning of the consumption–saving stage implies that they all hold the same end-of-period asset wealth; in this sense, all employed workers are alike, despite the fact that they have different individual employment histories. When such an employed worker falls into unemployment, the worker faces a binding debt limit, and hence liquidates his or her entire asset wealth (again, under the assumption that workers are so impatient that they hold at most very little wealth). It follows that the end-of-period wealth of all workers falling into unemployment is zero (i.e.,

the amount afforded by the debt limit). In the next period, such a worker will either find a job or remain unemployed. If the worker finds a job, he or she becomes again identical to all other employed workers (by virtue of the risk-pooling assumption). If the worker remains unemployed, then he or she again faces a binding debt limit and holds zero end-of-period wealth. This precisely makes this unemployed worker, and ultimately all unemployed workers (regardless of how long they have been unemployed), identical to those who are currently falling into unemployment. In this sense, all unemployed workers are alike, just like all employed workers were alike. So in this simple economy we have exactly two types of workers: employed workers, who all have the same positive end-of-period asset wealth, and unemployed workers, who all have the same (zero) end-of-period asset wealth.

The same reasoning can be extended to construct an equilibrium with three, rather than two, wealth states/types of workers. Suppose that workers are slightly less impatient than previously considered, so that those who fall into unemployment choose to keep some wealth (rather than fully liquidating it) as a precautionary buffer against an additional period of unemployment. All workers enter their first unemployment period with the same asset wealth—that inherited from the previous period, when they were employed and hence “identical.” Consequently, they all make the same asset holding choice and remain symmetric at the end of that first unemployment period. The period after, they will either find a job (in which case they will become identical to the existing employed workers) or they will remain unemployed and liquidate what is left of their wealth. In this configuration, there are exactly three worker types: employed workers (with “high” end-of-period wealth), unemployed workers at their first period of unemployment (with “low” end-of-period wealth), and all the other unemployed workers (with zero asset wealth). And again, we can go from three to four wealth states, then from four to five, and so on, by gradually making workers more and more patient.

Proposition 1(b) clarifies the sense in which starting from any $\bar{\mu}_0$ the distribution asymptotically “converges” toward a distribution with countable support (in the sense that the measure of workers whose individual state (a, N) does not belong to a countable subset of $\mathbb{R} \times \mathbb{Z}$ vanishes asymptotically). When all workers with the same N are indistinguishable, we call any worker with length of unemployment spell equal to N periods a *type- N worker*.

PROPOSITION 2. *Assume that $\mu(a, N)$ has a unique mass point in a for all N . Then workers’ intertemporal marginal rates of substitution (IMRS) are given by*

$$\begin{aligned}
 &M^{W'}(0) \\
 &= \beta^W \frac{(1 - s')u_c(c^{W'}(0) - hc^{W'}(0)) + s'u_c(c^{W'}(1) - hc^{W'}(1))}{u_c(c^W(0) - hc^W(0))} \quad \text{for } N = 0, \quad (30)
 \end{aligned}$$

$$\begin{aligned}
 &M^{W'}(N) \\
 &= \beta^W \frac{(1 - f')u_c(c^W(N + 1) - hc^{W'}(N + 1)) + f'u_c(c^W(0) - h^W c^{W'}(0))}{u_c(c^W(N) - hc^W(N))} \\
 &\text{for } N \geq 1, \quad (31)
 \end{aligned}$$

where $\mathbf{c}^W(N)$ satisfies $\mathbf{c}^{W'}(N) = c^W(N)$ (i.e., current consumption determines next period's habits for any N).

As usual, those IMRS consist of ratios of next-period to current marginal utilities. For example, the marginal utility of an employed worker (the denominator of $M^{W'}(0)$) is $u_c(c^W(0) - h\mathbf{c}^W(0))$. Next-period marginal utility (the numerator) must be broken into two idiosyncratic states, because the worker will either stay employed in the next period (with probability $1 - s'$), in which case he will enjoy marginal utility $u_c(c^{W'}(0) - h\mathbf{c}^{W'}(0))$, or fall into unemployment (with probability s') and enjoy marginal utility $u_c(c^{W'}(1) - h\mathbf{c}^{W'}(1))$. The IMRS of currently unemployed workers follows from the fact that a worker with $N \geq 1$ today may stay unemployed in the next period (with probability $1 - f'$), and thus become $N + 1$ with marginal utility $u_c(c^{W'}(N + 1) - h\mathbf{c}^{W'}(N + 1))$, or find a job (with probability f') and obtain marginal utility $u_c(c^{W'}(0) - h\mathbf{c}^{W'}(0))$. By making future marginal utility a weighted average of marginal utilities in each idiosyncratic state, these IMRS tie current and future marginal utilities for both employed and unemployed workers *exactly* as any model with incomplete insurance and encompass the same precautionary motive as they do.

We now make an additional assumption that will further simplify the cross-sectional distribution of workers over individual assets and length of unemployment spell.

ASSUMPTION 2. We have $\underline{a} > \underline{a}^{\text{nat}} \equiv \frac{\beta^F b^u}{\beta^F - e^{(\sigma-1)\mu_z}}$.

Under Assumption 2, the exogenous debt limit $\underline{a}e^z$ is strictly tighter than the “natural” debt limit $\underline{a}^{\text{nat}}e^z$ in any BGP, where the natural limit is defined as the maximum amount that a household can borrow while still being able to repay in the worst possible individual history (Aiyagari (1994)). In our model, this worst possible history corresponds to a history of permanent unemployment, and the present value of the corresponding income stream is

$$\frac{b^u e^{z+\mu_z}}{1+\bar{r}} + \frac{b^u e^{z+2\mu_z}}{(1+\bar{r})^2} + \frac{b^u e^{z+3\mu_z}}{(1+\bar{r})^3} + \dots = \frac{\beta^F b^u}{\beta^F - e^{(\sigma-1)\mu_z}} e^z = -\underline{a}^{\text{nat}} e^z \quad (\geq 0),$$

where we have used the fact that the gross real interest rate is $1 + \bar{r} = e^{\sigma\mu_z} / \beta^F$ in any BGP. We then have the following proposition.

PROPOSITION 3. Under Assumption 2, in a BGP, workers face a binding debt limit after a finite number of unemployment periods. Formally,

$$\exists \hat{N} \in \mathbb{Z}_+, \hat{N} < \infty : \begin{cases} \forall N < \hat{N}, & M^{W'}(N)(1+r') = 1, \\ \forall N \geq \hat{N}, & M^{W'}(N)(1+r') < 1. \end{cases}$$

The intuition for this result is as follows. We know from Proposition 1 that for any $\bar{N} \in \mathbb{Z}_+$, after $\bar{N} + 1$ periods, all unemployed type- N workers, for all $N \leq \bar{N}$, are indistinguishable. Those workers keep on decumulating assets while remaining unemployed, and in so doing they gradually approach the debt limit $\underline{a}e^z$. A worker who would remain

indefinitely unemployed would never actually reach the natural limit because borrowing up to that limit would force zero consumption in every following period, which would be suboptimal. However, whenever the debt limit is tighter than the natural limit (as Assumption 2 states), then that tighter limit can be reached in finite time (and will be, due to worker's impatience) while still allowing strictly positive consumption. This implies the following key corollary.

COROLLARY 1. *In a BGP, under Assumptions 1 and 2, the distribution μ has a finite support.*

This follows directly from the fact that \hat{N} in Proposition 3 is finite, which implies that the support of μ has at most $\hat{N} + 1$ points. Moreover, we can state another corollary.

COROLLARY 2. *Under Assumptions 1 and 2, the number of points in the support of μ is the same in a BGP and in the vicinity of this BGP.*

This corollary comes from the fact that, along the BGP, unemployed workers who face a binding debt limit have their Euler equation holding with strict inequality, and this will also be the case when aggregate shocks are sufficiently small.

2.2 Equilibrium conditions

We now characterize the workers' problem in the symmetric recursive equilibrium in the vicinity of a BGP. From Corollary 2, in the vicinity of a BGP the distribution of workers has finite support, which will allow us to characterize the equilibrium by a finite set of equilibrium conditions.

In every period, workers' end-of-period wealth and consumption are given by the sequences $\{a^{W'}(N), c^W(N)\}_{N \geq 0}$. Since workers face a binding debt limit after \hat{N} consecutive periods of unemployment, we have

$$a^{W'}(N) = \underline{a}e^z \quad \text{for } N \geq \hat{N}. \quad (32)$$

From the workers' budget constraint, we have

$$c^W(N) = (b^u - \underline{a})e^z + (1+r)a(N) \quad \text{for all } N > \hat{N}, \quad (33)$$

while

$$c^W(\hat{N}) = (b^u - \underline{a})e^z + (1+r)a(\hat{N}). \quad (34)$$

The remaining elements of $\{a^{W'}(N), c^W(N)\}_{N=0 \dots \hat{N}-1}$ are determined as follows. First, workers with unemployment spells of length 1 to $\hat{N} - 1$ have the budget constraints

$$a^{W'}(N) + c^W(N) = b^u e^z + (1+r)a(N), \quad N = 1, \dots, \hat{N} - 1, \quad (35)$$

where the wealth of a type- N worker at the beginning of the consumption–saving stage, $a(N)$, is his accumulated assets at the end of the consumption–saving stage of the previous period, when he was an $N - 1$ worker:

$$a(N) = a^{W'}(N - 1). \quad (36)$$

Second, the budget constraint of employed workers is

$$c^W(0) + a^{W'}(0) = (1 - \tau)w + (1 + r)A, \quad (37)$$

where

$$A' = \frac{(1 - s')\mathbf{n}^W a^{W'}(0) + f' \sum_{N=1}^{\hat{N}} \mathbf{n}(N) a^{W'}(N) + f' \left(1 - \mathbf{n}^W - \sum_{N=1}^{\hat{N}} \mathbf{n}(N) \right) \underline{a} e^z}{\mathbf{n}^{W'}}. \quad (38)$$

Third, there are \hat{N} first-order conditions, corresponding to the interior asset holding choices of type-0 to $\hat{N} - 1$ workers:

$$\mathbb{E}_X M^{W'}(N)(1 + r') = 1 \quad \text{for } N = 0, 1, \dots, \hat{N} - 1. \quad (39)$$

Given the processes that workers take as given (see point (i) of Definition 1), equations (35)–(39) form a system of $2\hat{N} + 1$ equations in the $2\hat{N} + 1$ variables $(A', \{a^{W'}(N), c^W(N)\}_{N=0, \dots, \hat{N}-1})$. From the solution of the system, we can find $c^W(\hat{N}) = (b^u - \underline{a})e^z + (1 + r)a(\hat{N})$, since equation (36) implies that $a(\hat{N}) = a^{W'}(\hat{N} - 1)$.

Equations (32)–(39) form the finite set of conditions that characterize the symmetric recursive equilibrium in the vicinity of a BGP. Note that while \hat{N} is constant in and around a BGP, it is an endogenous variable that depends on the underlying parameters of the model. Thus, to construct the equilibrium, we must first conjecture a particular value of \hat{N} , then check that the existence conditions for this equilibrium to exist are verified.

2.2.1 Existence conditions Equations (32)–(39) above determine the dynamics of $a^{W'}(N)$ and $c^W(N)$ as functions of the aggregate state, under the conjecture that the debt limit is binding for all type- N workers such that $N \geq \hat{N}$ and only for them. This requires that equation (32) hold for all workers of type $N < \hat{N}$, while

$$\mathbb{E}_X [M^{W'}(N)(1 + r')] < 1 \quad \text{for } N \geq \hat{N}. \quad (40)$$

Those conditions can be checked empirically for a specific joint distribution of the structural parameters. Given a value of the aggregate shocks, we will provide a posterior probability that the existence conditions (32) and (40) hold.

2.2.2 The case $\hat{N} = 1$ While Proposition 3 establishes the existence of an \hat{N} under Assumptions 1 and 2, it does not give the exact value of \hat{N} in a particular model economy. However, in quantitative applications, the data impose additional discipline on the equilibrium structure. For example, the amount of wealth that workers hold (hence the initial wealth of a worker falling into unemployment) must be consistent with the broad features of the empirical cross-sectional distribution of wealth, and the job transition rates (which are a key determinant of both initial bond holdings and the pace of asset decumulation) must be consistent with their empirical counterparts. We argue in Section 3 below that, given our quantitative exercise and our focus on liquid (rather than

total) wealth and the definition of the period as a quarter, the data favor a specification where workers' wealth is fully liquidated after one period, that is, $\hat{N} = 1$. It follows that there are at every point in time exactly three distinct types of workers: $N = 0$, $N = 1$, and $N \geq 2$, with consumption levels $c^W(0)$, $c^W(1)$, and $c^W(N) = c^W(2)$ (for all $N \geq 2$), respectively. These types are in numbers $\Omega \mathbf{n}^W$, $\Omega s \tilde{\mathbf{n}}^W$, and $\Omega(1 - \mathbf{n}^W - s \tilde{\mathbf{n}}^W)$, respectively. Type $N = 0$ workers save $a^{W'}(0) > \underline{a}e^z$, while types $N = 1$ and $N \geq 2$ all save $\underline{a}e^z \leq 0$. Finally, because there are only three types of workers, there are only three relevant habit levels to keep track of: $\mathbf{c}^W(0)$, $\mathbf{c}^W(1)$, and $\mathbf{c}^W(2)$. Since habit levels are determined by the average consumption of the relevant group in the previous period, we have $\mathbf{c}^{W'}(0) = \mathbf{c}^W(0)$, $\mathbf{c}^{W'}(1) = \mathbf{c}^W(1)$, and $\mathbf{c}^{W'}(2) = \mathbf{c}^W(2)$. The full set of equilibrium conditions and associated existence conditions in the $\hat{N} = 1$ economy are stated in the Technical Appendix.

2.3 Time-varying precautionary savings

We are now in a position to isolate the key determinants of workers' precautionary motive and to show why it survives in a first-order approximation to their optimal consumption–saving choice. We assume that $\hat{N} \geq 1$, so that at least employed workers' choices are interior.¹⁹ An employed worker's incentive to save is summarized by the behavior of its IMRS, as given in equation (30). If we abstract from consumption habits for the clarity of the argument, the IMRS of an employed worker becomes

$$M^{W'}(0) = \beta^W \frac{(1 - s')u_c(c^{W'}(0)) + s'u_c(c^{W'}(1))}{u_c(c^W(0))}. \tag{41}$$

There is uninsured idiosyncratic unemployment risk whenever $s' > 0$ (i.e., the job-loss rate is positive) and $c^{W'}(1) < c^{W'}(0)$ (i.e., falling into unemployment generates a consumption loss). In this case, an increase in the job-loss rate s' raises future marginal utility, that is, it raises the incentive to save. This is the precautionary motive for holding assets. Importantly, the job-loss rate s' affects the precautionary level of assets even if we consider a first-order approximation to the workers' Euler equation. To see this, observe that the log deviation of the IMRS from its steady-state value is given by

$$\hat{M}^{W'}(0) \simeq \underbrace{-\sigma \left(\frac{c^{W'}(0) - c^W(0)}{c^W(0)} \right)}_{\text{perfect-insurance IMRS}} + \underbrace{\sigma \times \Phi \times s'}_{\text{impact of imperfect insurance}}, \tag{42}$$

where Φ is the (constant) mean consumption growth differential between a worker who stays employed from the current to the next period and that of a worker who falls into unemployment in the next period:

$$\Phi \equiv \mathbb{E} \left[\frac{c^{W'}(0) - c^{W'}(1)}{c^W(0)} \right] > 0.$$

¹⁹When $\hat{N} = 0$, all workers, including employed workers, face a binding debt limit, so that there is no precautionary savings. When $\hat{N} \geq 1$, then at least employed workers accumulate precautionary savings in excess of the debt limit and some unemployed workers also do so whenever $\hat{N} \geq 2$.

The first term on the right-hand side of equation (42) corresponds to the usual log-linear IMRS under perfect insurance (i.e., if all workers consumed the same amount $c^{W'}(0)$). The second term on the right-hand side is a correction to the IMRS, coming from the fact that employed workers are imperfectly insured. The impact of the probability of falling into unemployment s' on workers' willingness to save in the current period (as measured by $\hat{M}^{W'}(0)$) is scaled both by workers' risk aversion (σ) and the incidence of unemployment risk on employed workers' consumption growth (Φ). That the precautionary motive to hold assets is preserved in this first-order approximation of $M^{W'}(0)$ implies that it will remain operative when we solve and estimate the linear state-space representation of the full model.

Note that whenever $\hat{N} \geq 2$, not only employed workers but also some of the unemployed workers (those of types $N \in [1, \hat{N} - 1]$) wish to keep a buffer stock of wealth in excess of the debt limit (despite their being impatient). Their willingness to save thus increases with the probability of remaining unemployed, that is, $1 - f'$ (see equation (31)).

3. ESTIMATING THE IMPERFECT-INSURANCE MODEL

In this section, we take our baseline imperfect-insurance model to the data. We first describe the functional forms, the aggregate shocks, and the Bayesian empirical strategy we use. We emphasize our ability to incorporate cross-sectional data in our likelihood-based estimation, next to the traditional aggregate macroeconomic, monetary, and labor-market transitions data. We will split the parameter set into two subsets. The first subset will be chosen to match some unconditional moments related to the steady state of the model. Importantly, this step makes use of both aggregate and cross-sectional data, because we will match some cross-sectional unconditional moments. We call this first subset the calibrated parameters. The second subset will be estimated and called accordingly. For this second set, we describe the prior choice together with the posterior estimates. We then check whether the equilibrium existence conditions hold under the posterior parameter distribution and briefly comment on the model's fit to the data.

3.1 Functional forms and aggregate shocks

Before proceeding, we must specify the functional forms adopted for the utilization cost function, $\eta(\cdot)$, and the investment adjustment cost function, $S(\cdot)$. In particular, we assume

$$\eta(v) = \frac{\bar{r}_k}{\tilde{\nu}_v} [e^{\tilde{\nu}_v(v-1)} - 1] \quad \text{and} \quad S\left(\frac{i'}{i}\right) = \frac{\nu_i}{2} \left(\frac{i'}{i} - \bar{g}_i\right)^2 \quad \text{with } \tilde{\nu}_v, \nu_i > 0.$$

Here \bar{r}_k and \bar{g}_i are the steady-state values of the rental rate of capital services and the growth factor of investment i'/i , while $\tilde{\nu}_v$ and ν_i are the curvatures of the utilization and investment adjustment cost functions. These functional forms ensure that in a steady state both costs vanish. As in Smets and Wouters (2007), we define $\nu_v \equiv \tilde{\nu}_v/(1 + \tilde{\nu}_v)$ and estimate ν_v rather than $\tilde{\nu}_v$. This allows us to eliminate numerical problems at the estimation stage. Moreover, we impose the functional form for the exogenous job-destruction

rate,

$$\rho(\varphi_s) = \frac{1}{1 + e^{\bar{\rho} - \varphi_s}},$$

where $\bar{\rho}$ is a constant that pins down the steady-state value of ρ . With this functional form, we ensure that ρ varies only in the compact set $[0, 1]$.

The considered aggregate shocks are of two classes. First, the technology shock z follows this nonstationary process:

$$z' = \mu_z + z + \varphi'_z.$$

Second, φ_h , for $h \in \{z, i, c, s, R, w, p\}$, are stationary and are assumed to follow AR(1) processes of the form

$$\varphi'_h = \rho_h \varphi_h + \sigma_h \epsilon'_h, \quad \epsilon_h \sim \mathcal{N} \text{ iid}(0, 1).$$

3.2 Empirical strategy and data

Before taking the model to the data, we first need to induce stationarity by normalizing the first-order conditions of the model by e^z . Then we linearize the resulting system in the neighborhood of the normalized steady state.²⁰ Let \hat{X} denote the vector collecting the deviations from the steady state of the normalized state variables, and let ϵ denote the vector collecting the innovations to the aggregate shocks. The law of motion of \hat{X} is of the form

$$\hat{X}' = \mathbf{F}(\boldsymbol{\vartheta})\hat{X} + \mathbf{G}(\boldsymbol{\vartheta})\epsilon', \quad (43)$$

where

$$\boldsymbol{\vartheta} = (\Omega, \sigma, h, \beta^F, \beta^W, \delta, \theta, \phi, \kappa_v, \kappa_y, \bar{\rho}, \chi, \bar{m}, \nu_i, \nu_v, \bar{\eta}, \alpha, \gamma_p, \psi, \bar{w}, \psi_n, \gamma_w, b_u, \underline{a}, \bar{\pi}, \rho, a_\pi, a_y, \mu_z, \rho_x \text{ for } x \in \{c, i, w, s, p, R\}, \sigma_x \text{ for } x \in \{c, i, w, s, p, R, z\})$$

is the vector of all structural parameters of the model. The matrices $\mathbf{F}(\boldsymbol{\vartheta})$ and $\mathbf{G}(\boldsymbol{\vartheta})$ are functions of the model's structural parameters.

The vector of variables used as observable in estimation, O , consists of the growth rates of consumption and total investment, $\Delta \log(c)$ and $\Delta \log(\tilde{i})$, respectively, inflation π , the nominal interest rate R , the nominal wage inflation $\Delta \log(W)$, the job-loss rate s , and the job-finding rate f . Total investment \tilde{i} is defined as the sum of investment i and utilization and vacancy-posting costs, so that $\tilde{i} = i + \eta(v)k + \kappa_v v$. Nominal wage inflation, $\Delta \log(W)$, is defined as real wage inflation plus inflation, that is, $\Delta \log(W) = \Delta \log(w) + \pi$.

The vector O also contains the share of consumption of the 60 percent poorest households in total consumption, that is, c_{60}^*/c^* . Here c_{60}^* is defined as

$$c_{60}^* = \Omega[\mathbf{n}^W c^W(0) + s\tilde{\mathbf{n}}^W c^W(1) + (1 - \mathbf{n}^W - s\tilde{\mathbf{n}}^W)c^W(2)].$$

²⁰See the Technical Appendix for a description of the normalized steady state.

While in our theoretical setup it is the case that c^* and total consumption c coincide exactly, a statistical discrepancy between their data counterpart from the Consumer Expenditure Survey (CEX) and the Bureau of Economic Analysis (BEA) has been documented (see [Heathcote, Perri, and Violante \(2010\)](#)). In an attempt to capture this discrepancy, we append an AR(1) measurement error, u , to the consumption share of the 60 percent poorest households. Formally,

$$O \equiv (\Delta \log(c), \Delta \log(\tilde{i}), \pi, R, \Delta \log(W), s, f, c_{60}^*/c^*)^\top;$$

hence, the measurement equation is

$$O' = \mathbf{M}(\boldsymbol{\vartheta}) + \mathbf{H}(\boldsymbol{\vartheta})\hat{X} + \mathbf{e}_8 u', \quad (44)$$

where \mathbf{e}_8 is the eighth column of the identity matrix of dimension 8, and we use $\mathbf{M}(\boldsymbol{\vartheta})$ to denote the vector of means of observed variables, which are tied to $\boldsymbol{\vartheta}$.

Our sample runs from 1982Q1 to 2007Q4.²¹ We stop our same in 2007Q4 because our likelihood-based estimation approach cannot handle the zero lower bound associated with the Great Recession. We start in 1982Q1 for two reasons. First, the CEX—to be used below—begins in 1980 and, second, we want to avoid the Volcker disinflation period to minimize the effects of parameter uncertainty. Consumption is defined as the sum of personal consumption expenditures on nondurable goods and services as well as government consumption expenditures and gross investment. The resulting series is deflated by the implicit gross domestic product (GDP) deflator. Investment is defined as the sum of gross private domestic investment and personal consumption expenditures on durable goods. The resulting series is also deflated by the implicit GDP deflator. These two series are converted to per capita terms by dividing them by the civilian population, aged 16 and over. Inflation is calculated using the GDP deflator, and the nominal interest rate is defined as the Effective Federal Funds Rate. Finally, we measure nominal wages as the average weekly earnings of production and nonsupervisory employees, from the Current Employment Statistics survey.

For the labor-market transition probabilities f and s , we proceed as follows. First, we compute monthly transition rates using Current Population Survey (CPS) data on unemployment and short-run unemployment, using the approach of [Shimer \(2005, 2012\)](#). Using these series, we construct transition matrices across employment statuses for every month in the sample and then multiply those matrices over the three consecutive months of each quarter to obtain quarterly transition rates.

To construct the consumption share of the 60 percent poorest households, we first aggregate nondurable items in the CEX to compute individual nondurables consumption (using the same categories as [Heathcote, Perri, and Violante \(2010\)](#)), and then sort consumption by income levels to obtain c_{60}^*/c^* .

We follow the Bayesian approach to estimate the model's structural parameters. Based on the state-space representation for the dynamic system represented by (43) and

²¹The data used for estimation come from the BEA, the Federal Reserve Bank of St. Louis' FRED II data base, and the Bureau of Labor Statistics (BLS). See the Technical Appendix for further details on the data.

(44), we (i) evaluate the likelihood of the observed variables at any value of ϑ using the Kalman filter and (ii) form the posterior distribution by combining the likelihood function with a joint density characterizing some prior beliefs.

Given the specification of the model, the posterior distribution cannot be recovered analytically, but we may numerically draw from it using a Monte Carlo Markov chain (MCMC) sampling approach. More specifically, we rely on the Metropolis–Hastings algorithm to obtain a random draw of size 1,000,000 from the posterior distribution of the structural parameters.

3.3 Calibrated parameters

As mentioned, the vector of structural parameters ϑ is split into three subsets ϑ_1 , ϑ_2 , and ϑ_3 . The vector of fixed parameters,

$$\vartheta_1 = (\delta, \theta, \chi, \bar{m}, \Omega, \underline{a}),$$

contains structural parameters that are fixed to the same values standard in the literature. The depreciation rate $\delta = 0.015$ implies a 6 percent annual depreciation of the capital stock. We choose θ such that the markup is 20 percent. The elasticity of the matching function with respect to vacancies is set to 0.5, so that $\chi = 0.5$. The parameter \bar{m} is normalized to 1. Finally, we set the share of workers in the model to $\Omega = 0.6$ and restrict borrowing by setting $\underline{a} = 0$.

The vector of calibrated parameters,

$$\vartheta_2 = (\bar{\pi}, \mu_z, \beta^F, \beta^W, b^u, \phi, \kappa_v, \bar{\rho}, \kappa_y, \psi, \bar{w}),$$

are tied to some restriction implied by the fact that we force the steady state to match some unconditional moments. The steady-state inflation rate $\bar{\pi}$ is set to match the average value of inflation over the sample. The growth rate of technical progress μ_z is set to match the average growth of output over the sample. The subjective discount factor of firm owners $\beta^F = 0.9985$ is set so that, given both μ_z and $\bar{\pi}$, the steady-state nominal real interest rate matches its average empirical counterpart. The subjective discount factor of workers $\beta^W = 0.9835$ is set so that the average nondurables consumption loss when falling into unemployment is 21 percent, as documented by [Chodorow-Reich and Karabarbounis \(2014\)](#). The parameter b^u is set so as to impose a replacement rate of 50 percent. Note that the values of β^W and b^u determine workers' incentive to hold precautionary wealth and hence shape the cross-sectional distribution of wealth. Under our parameters, the workers (i.e., the poorest 60 percent of the households) hold less than 1 percent of total wealth. This extent of wealth inequality is consistent with the data in the sense that the share of nonhome (or "liquid") wealth held by the poorest 60 percent of the U.S. households is typically less than 1 percent in the Survey of Consumer Finances (SCF).²²

²²Liquid wealth is the relevant wealth concept to think about households' ability to smooth nondurables consumption in the face of idiosyncratic income shocks occurring at the business-cycle frequency (see [Challe and Ragot \(2014\)](#), [Kaplan and Violante \(2014\)](#), [Kaplan, Violante, and Weidner \(2014\)](#)). In the 2007 wave of the SCF, the bottom 60 percent of the households in terms of liquid wealth held 0.31 percent of total liquid wealth (or about \$700 on average).

We pin down ϕ so that the labor share in income is 64 percent. We set κ_v so that the share of vacancy costs in output is 1 percent. The parameter $\bar{\rho}$ is pinned down by imposing that the steady-state value of s coincides with its empirical average value. The parameter κ_y is set so that steady-state monopolistic profits are zero. We set the skill premium parameter ψ so as to match the average share of income of the 60 percent poorest households in total income, as backed out from CEX data. This last choice underscores a clear advantage of our setup, in that it allows us to make explicit contact with cross-sectional data at the calibration stage. Given the preceding restrictions, we select \bar{w} to match the average value of f in the data.

Finally, the rest of the parameters contained in \mathfrak{D}_3 will be estimated using Bayesian methods. It is important to note that because most of the mentioned restrictions also involve the structural parameters in \mathfrak{D}_3 , the calibrated parameters \mathfrak{D}_2 will be functions of the posterior draws. More precisely, for each draw of \mathfrak{D}_3 , we readjust \mathfrak{D}_2 to meet all the steady-state restrictions listed above.

3.4 Estimated parameters

The remaining structural parameters contained in \mathfrak{D}_2 are estimated. They are listed in Table 1, together with information on their prior and posterior distributions. In addition to the prior, for each parameter, the table reports the posterior mean and the standard deviation together with the bounds of the 90 percent highest posterior density interval (HPD; labeled “low” and “high”).²³

We will now comment on the estimation results. The posterior means of σ and h are 0.72 and 0.66, respectively. This implies that the coefficient of relative risk aversion along the BGP $-cu_{cc}(c)/u_c(c) \approx (1 + \mu_z)\sigma/(1 + \mu_z - h)$ is equal to 2.1, which is well within the range of available estimates. The posterior mean of transformed curvature of the utilization cost ν_u is 0.58, slightly different from the prior mean. This implies that the actual degree of curvature of the utilization cost function $\tilde{\nu}_u = \nu_u/(1 - \nu_u)$ is about 1.4. The posterior mean for ν_i is 1.9. The degree of price stickiness α has a posterior mean equal to 0.73, implying an average price duration of less than 4 quarters. The posterior means of the degrees of price and wage indexation are equal to 0.34 and 0.82, respectively. The value for γ_p is in the range of previous estimates obtained in the literature. The value for γ_w is quite high, suggesting a relatively strong degree of nominal wage stickiness. In turn, this is partly reinforced by a posterior mean degree of responsiveness of wages to employment of about 1.86, higher than its prior mean. The posterior mean of the degree of interest-rate smoothing is lower than its prior mean, with $\rho = 0.46$. The responsiveness of the nominal interest rate has a posterior mean of 2.01, higher than its prior mean. This is indicative of a strong reaction to inflation on the part of the monetary policy authority, consistent with the historical record over the post-Volcker era. The posterior mean of the responsiveness to economic activity is relatively higher than the prior mean. Finally, the structural parameters controlling the aggregate shocks are within the results obtained in the literature. Here, it is important to bear in mind that our estimate for the

²³Further discussion of our choice of prior distributions can be found in the Technical Appendix.

TABLE 1. Estimation results.

Parameter	Prior Shape	Prior Mean	Prior S.D.	Post. Mean	Post. S.D.	Low	High
σ	Gamma	1.50	0.20	0.72	0.10	0.56	0.89
h	Beta	0.50	0.10	0.66	0.04	0.60	0.72
ν_i	Gamma	2.00	0.20	1.90	0.19	1.59	2.21
ν_u	Beta	0.50	0.10	0.58	0.09	0.44	0.73
α	Beta	0.50	0.10	0.73	0.04	0.66	0.79
γ_p	Beta	0.50	0.10	0.34	0.09	0.18	0.48
γ_w	Beta	0.50	0.10	0.82	0.04	0.76	0.88
ψ_n	Gamma	1.00	0.20	1.86	0.32	1.32	2.38
ρ	Beta	0.75	0.10	0.46	0.06	0.36	0.55
a_π	Gamma	1.50	0.10	2.01	0.10	1.85	2.18
a_y	Gamma	0.13	0.10	0.52	0.16	0.27	0.76
ρ_z	Beta	0.20	0.10	0.42	0.06	0.32	0.52
ρ_c	Beta	0.50	0.10	0.58	0.04	0.52	0.65
ρ_w	Beta	0.50	0.10	0.79	0.06	0.70	0.89
ρ_i	Beta	0.50	0.10	0.86	0.05	0.79	0.93
ρ_p	Beta	0.50	0.10	0.91	0.03	0.86	0.96
ρ_s	Beta	0.50	0.10	0.66	0.05	0.57	0.75
ρ_R	Beta	0.50	0.10	0.49	0.06	0.39	0.59
ρ_u	Beta	0.50	0.10	0.83	0.04	0.77	0.89
σ_c	Inverted gamma	1.00	0.20	0.67	0.09	0.53	0.81
σ_w	Inverted gamma	1.00	0.20	0.52	0.04	0.45	0.58
σ_i	Inverted gamma	1.00	0.20	2.80	0.37	2.20	3.37
σ_p	Inverted gamma	1.00	0.20	1.42	0.23	1.07	1.75
σ_z	Inverted gamma	1.00	0.20	1.17	0.08	1.04	1.29
σ_R	Inverted gamma	1.00	0.20	0.43	0.03	0.37	0.48
σ_s	Inverted gamma	1.00	0.20	6.95	0.47	6.17	7.69
σ_u	Inverted gamma	1.00	0.20	0.80	0.06	0.71	0.89

Note: “Low” and “high” stand for the lower and upper boundaries of the 90 percent HPD interval, respectively.

standard error of the markup shock σ_p is not directly comparable to results discussed in the literature. The reason is that, in general, the markup shock is rescaled by the slope of the New Keynesian Phillips curve, resulting in a relatively low standard error of markup shocks. In the present paper, σ_p is not rescaled. If it were, our estimate would broadly fall in the ballpark of available estimates.

3.5 Verification of the existence conditions

We may now check empirically that the the conditions for the existence of an equilibrium with $\hat{N} = 1$ are satisfied. From (32) and (40), with a zero debt limit, this requires $a^{W'}(N) = 0 \Leftrightarrow \mathbb{E}_X[M^{W'}(1)(1+r')] < 1$ for all $N \geq 1$. However, all workers with $N \geq 2$ are indistinguishable, so it is enough to check the latter inequalities for $N = 1, 2$. The equilibrium also requires positive precautionary savings for employed workers, that is, $a^{W'}(0) > 0$ ($\Leftrightarrow \mathbb{E}_X[M^{W'}(0)(1+r')] = 1$).

From left to right and top to bottom, the panels in Figure 2 report the posterior mean (thick line) of $1 - \mathbb{E}_X[M^{W'}(1)(1+r')]$, $1 - \mathbb{E}_X[M^{W'}(2)(1+r')]$, and $a^{W'}(0)$ (each appropriately normalized), respectively, over the estimation sample, as implied by the smoothed

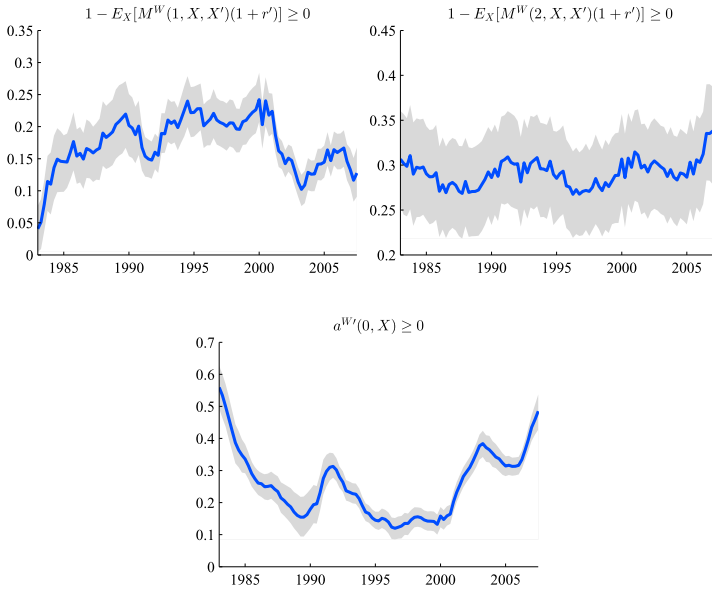


FIGURE 2. Existence conditions.

Note: The thick line is the posterior mean path; the grey area is the 90 percent HPD interval.

values of the state variables. In each panel, we also report the associated 90 percent HPD interval (the grey area delineated by the thin, black dashed lines). Figure 2 makes clear that the posterior probability that the existence conditions are indeed satisfied is very close to 1.

3.6 Empirical performance

In this section, we show that our baseline imperfect-insurance model empirically outperforms the perfect-insurance benchmark, so that taking into account time-varying precautionary savings improves the fit to the data.

The perfect-insurance benchmark is structurally identical to our baseline model, except that all workers enjoy perfect insurance (and not only firm owners). Because workers are impatient relative to firm owners, they then borrow up to the borrowing limit in every period, as in, for example, Kiyotaki and Moore (1997) or Iacoviello (2005). This model does not completely eliminate household heterogeneity (only heterogeneity among workers), so that the cross-sectional distribution of consumption is not degenerated (workers and firm owners consume different amounts), and we can compare the two models using the observable variables described above, that is, including the consumption shares. However, in the perfect-insurance model, workers no longer hold any precautionary wealth in excess of the borrowing limit (by construction), so the aggregate demand and supply effects of time-varying precautionary savings are absent.²⁴

Let \mathcal{M}_{II} and \mathcal{M}_{PI} denote our imperfect-insurance model and its perfect-insurance counterpart, respectively. Both model versions are estimated using the (i) same data,

²⁴See the Technical Appendix for a complete formal description of the perfect-insurance benchmark.

(ii) the same set of calibration restrictions, and (iii) the same prior distributions on the estimated parameters. We compare the fit of the two specifications by comparing their marginal likelihoods. Let $\log(p(O_{1:T}|\mathcal{M}_j))$ denote the log marginal likelihood of model \mathcal{M}_j for $j \in \{\text{II}, \text{PI}\}$, where $O_{1:T}$ denotes the sample observations of the data vector O . We obtain $\log(p(O_{1:T}|\mathcal{M}_{\text{II}})) = -870.5$ and $\log(p(O_{1:T}|\mathcal{M}_{\text{PI}})) = -877.9$.²⁵ Although the evidence is overwhelming in favor of \mathcal{M}_{II} , it is the case that this exercise alone does not tell us where these empirical gains are coming from. In other words, what features of the data is the partial-insurance model fitting better and why? This is a well known drawback of this approach.

It is also the case that these figures show that model \mathcal{M}_{II} is to be preferred to model \mathcal{M}_{PI} , but they do not tell us by how much. To provide an answer to this question, we endow each model with its own prior, $p(\mathcal{M}_j)$, for $j \in \{\text{II}, \text{PI}\}$. Here, we adopt a noninformative choice by setting $p(\mathcal{M}_{\text{II}}) = p(\mathcal{M}_{\text{PI}}) = 0.5$. Armed with these priors, we can then compute the posterior probability of each model specification. Given the above results, we obtain a posterior probability on specification \mathcal{M}_{II} , $p(\mathcal{M}_{\text{II}}|O_{1:T})$, that is very close to 1. Hence, almost all the probability mass is shifted toward the imperfect-insurance model.

Note that the imperfect-insurance model still outperforms the perfect-insurance model when we exclude the consumption share of the 60 percent poorest households from the set of observable variables $\log(p(O_{1:T}^*|\mathcal{M}_{\text{II}})) = -743.6$ and $\log(p(O_{1:T}^*|\mathcal{M}_{\text{PI}})) = -755.3$, where $O_{1:T}^*$ denotes the history of observable variables when the consumption share of the 60 percent poorest households is excluded. Once again, under the non-informative prior on models $p(\mathcal{M}_{\text{II}}) = p(\mathcal{M}_{\text{PI}}) = 0.5$, we obtain a posterior probability $p(\mathcal{M}_{\text{II}}|O_{1:T}^*)$ very close to 1. As an additional conclusion, these results show that considering the unemployment risk is important for the macroeconomy, which justifies the focus of this paper.

4. PRECAUTIONARY SAVINGS DURING POST-VOLCKER RECESSIONS

4.1 *Measuring the contribution of time-varying precautionary savings*

We now use our estimated imperfect-insurance model to measure the contribution of time-varying precautionary savings in the propagation of the recent U.S. recessions (the 1990–1991 recession, the 2001 recession, and the Great Recession). From a theoretical point of view, our model embodies the two aggregate effects of time-varying precautionary savings discussed in the introduction: the aggregate demand effect is operative, because the three basic frictions that we have assumed generate a mutually reinforcing feedback between idiosyncratic unemployment risk and aggregate consumption demand; but the aggregate supply effect is also operative, because our model has capital and, thus, the traditional smoothing effect of imperfect-insurance models (e.g., [Krusell and Smith \(1998\)](#)). In the presence of both effects, the question naturally arises as to which effect dominates, that is, whether time-varying precautionary savings ultimately

²⁵This difference is larger than 7, a bound for DNA testing in forensic science, often accepted by courts of law as evidence beyond reasonable doubt (see [Evetts \(1991\)](#)).

amplify or dampen recessions. Answering this question requires the use of a counterfactual economy; to this purpose we use the perfect-insurance benchmark discussed in the previous section, which by construction has constant (zero) precautionary savings.²⁶

We run our counterfactual experiments as follows. First, using the posterior mean of the distribution of estimated parameters ϑ_2 computed in Section 3, we run the Kalman smoother to extract the sequences of aggregate shocks experienced by the U.S. economy during recession episodes. To that end, we extend the estimation sample to the period covering the Great Recession.²⁷ Second, we feed these shocks into the perfect-insurance counterpart of our model. The perfect-insurance model is the one described in Section 3.6, but with all the structural parameters maintained at the posterior mean of the estimated imperfect-insurance model. We also make sure that the calibrated parameters ϑ_1 have the exact same values as those used in the above smoothing step. Third, we compare the time series generated by the perfect-insurance model (grey dashed line of Figures 3 and 4) with the historical time series (solid lines) for consumption and investment, as well as the job-finding, job-loss, and employment rates (the first two variables are expressed as log deviations from their value at the National Bureau of Economic Research (NBER) peak, while the last four are in level deviations from their value at the peak). Because the precautionary motive is shut down in the perfect-insurance benchmark, this comparison gives us a measurement of the propagating role of time-varying precautionary savings in the last three recessions. Figure 3 analyzes the Great Recession, while Figure 4 considers the 1990 and the 2001 recessions.²⁸

4.2 *The great recession*

Figure 3 compares the actual and counterfactual paths of the abovementioned variables during the Great Recession. The amplification generated by the precautionary motive is striking. The fall in consumption from peak to trough would have been 1.75 times smaller without the precautionary motive than it actually was in the data (it is -3.5% in the data but about -2% in our counterfactual experiment). This fall in consumption

²⁶The perfect-insurance benchmark is a natural counterfactual model because it is identical to the baseline model in every respect except for the assumption of perfect insurance, and differs from it only in that dimension. Despite their proximity, the two models have their own internal structure and parameter distributions, so they transmit structural shocks differently, not only because of the precautionary motive. For example, the workers' marginal propensity to consume is different across the two models.

²⁷At the end of this extended sample, the zero lower bound is binding. We adopt the method advocated by Ireland (2011) to deal with the latter. The procedure is basically as follows. Since we work with a fully linear model, sometimes agents expect the nominal interest rate to become negative because of the Taylor rule. Then, because in the data the interest rate always remains nonnegative, we need to find combinations of shocks that also allow the interest rate to always remain nonnegative in the model. These shocks mostly show up as contractionary monetary policy shocks.

²⁸When running these counterfactual experiments, each model version is initialized with its own historical state variables. Thus the dynamics reported in Figures 3 and 4 reflect that the two model versions (i) have different initial values and (ii) react in different ways to the same shocks. As a robustness exercise, we also considered running the same experiments while forcing each model's state variables to zero at the onset of each of the recessions under study. The results, reported in the Technical Appendix, are quantitatively similar.

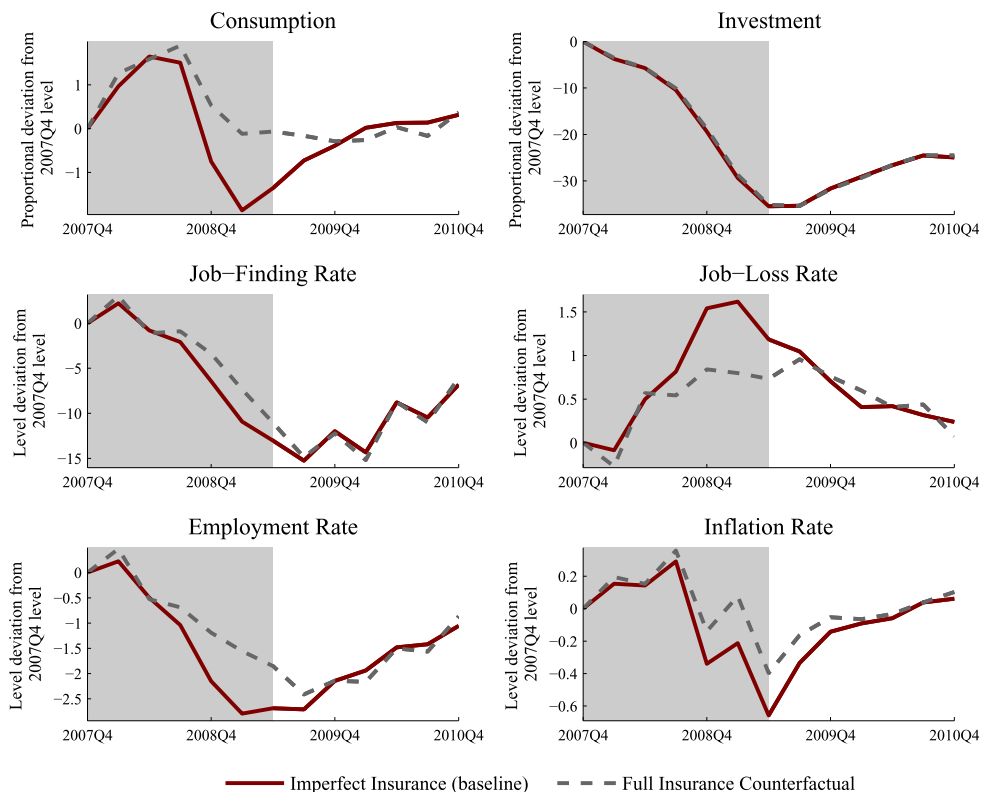


FIGURE 3. The Great Recession.

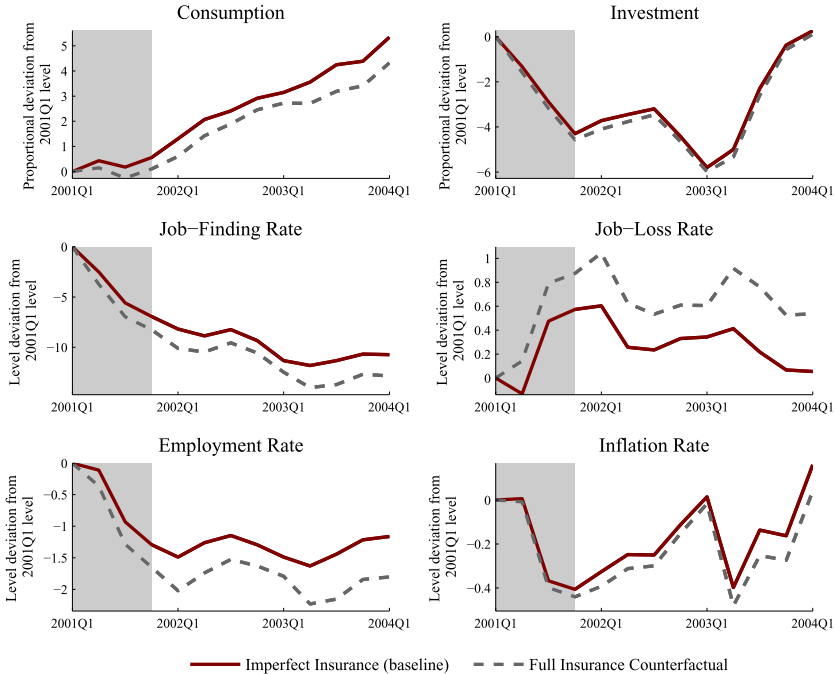
Note: The solid lines correspond to the actual paths of consumption, investment, the job-finding rate, the job-loss rate, and the employment rate. The dashed lines correspond to the counterfactual sample paths. Consumption and investment are reported in proportional deviation from their level at the beginning of the recession. All the other variables are expressed in level deviation from their values at the beginning of the recession. The grey areas indicates the recession dates.

reflects the rational response of imperfectly insured workers to the huge increase in idiosyncratic labor-market risk that they have faced, as is illustrated by the sharp rise in the job-loss rate over the period (about 2.5 standard deviations of the job-loss rate in the pre-recession sample). Moreover, there is evidence of a strong feedback from aggregate demand to idiosyncratic labor-market risk. This can be inferred from the responses of the job-finding and job-loss rates between the data and the perfect-insurance counterfactual. Let us recall that, on the eve of the Great Recession, the quarter-to-quarter job-finding and job-loss rates were 76 percent and 4 percent, respectively, while during the Great Recession the former crashed to 50 percent and the latter went up to 6 percent. Since the model allows for within-period labor-market transitions, the job-loss rate $s = \rho(1 - f)$ combines an exogenous separation component ρ and an endogenous job-finding component f that responds to aggregate demand. During the Great Recession, all of the increase in s is explained by the latter component.²⁹ Our counterfactual analysis indicates that without this feedback, the rise in the job-loss rate would have

²⁹More specifically, having computed quarterly values for f and s on the basis of the monthly rates (Section 3), we then calculated the value taken by the separation shock $\rho = s/(1 - f)$ as a residual. It turns out



(a) 1990Q3 Recession



(b) 2001Q1 Recession

FIGURE 4. The 1990Q3 and 2001Q1 Recessions.

Note: The red lines correspond to the actual paths of consumption, investment, the job-finding rate, the job-loss rate, and the employment rate. The dashed, grey lines correspond to the counterfactual sample paths. Consumption and investment are reported in proportional deviation from their level at the beginning of the recession. All the other variables are expressed in level deviation from their values at the beginning of the recession. The grey area indicates the recession dates.

been half as large as it actually was. But, as explained above, since s in turn drives time-varying precautionary savings and thereby consumption demand, this closes the feedback loop. Combined with the greater fall in the job-finding rate, this manifested itself as a large drop in the employment rate. Hence, as far as consumption and labor-market risk are concerned, the aggregate demand effect of the precautionary motive dominated the supply effect during the Great Recession.

Our counterfactual analysis also indicates that the fall in investment was not significantly affected by the precautionary motive. But this is to be expected, because the precautionary motive has two contradicting effects on investment. On the one hand, larger precautionary wealth in a recession takes down the real interest rate. This drop is transmitted to the market for capital claims by firm owners (who participate in both asset markets) and ultimately stimulates investment. This is precisely the aggregate supply effect of the precautionary motive, which tends to smooth the fall in investment relative to the perfect-insurance case. On the other hand, because of the aggregate demand effect, consumption and output are depressed, which tends to discourage investment. The overall impact of these two forces on investment is a priori ambiguous, and in the present case they roughly offset each other. Similarly, inflation is moderately affected by time-varying precautionary savings. This reflects the fact that the New Keynesian Phillips curve flattens in the post-Volcker period, that is, there are significant nominal price rigidities (see, e.g., [Sbordone \(2007\)](#), [Coibion and Gorodnichenko \(2013\)](#), [Del Negro, Giannoni, and Schorfheide \(2015\)](#)). Hence, large variations in aggregate demand, including those due to the precautionary motive, are associated with limited movements in inflation.

4.3 *The 1990–1991 and 2001 recessions*

We now turn to the other two recessions in our sample. Figure 4 compares the actual and counterfactual paths of consumption, investment, the job-finding rate, the job-loss rate, and the employment rate over the 1990–1991 and 2001 recessions. The impact of the precautionary motive on the propagation of the 1990–1991 recession shares several features of the Great Recession, except, of course, for the size of the effect. More specifically, over the duration of the recession, and also over the three quarters after the recession ended, consumption stagnated. Our counterfactual experiment indicates that it would have kept growing at a moderate pace during this time had the precautionary motive not been active. The feedback from stagnating consumption to depressed labor-market conditions is also apparent from the comparison between the paths of actual labor-market transition rates and those that would have prevailed without the precautionary motive; in particular, the job-loss rate would have roughly stayed at pre-recession level during most of the year 1991. In other words, just as in the Great Recession, the aggregate demand effect of the precautionary motive dominated the supply effect during the 1990–1991 recession.

that ρ has slightly fallen (from 15 percent to 12 percent) during the Great Recession, so the rise in s entirely comes from the fall in f .

The 2001 recession was short and mild, even when compared to the 1990–1991 recession. For that reason, the precautionary motive may have been weak, as is reflected by the small difference between the actual paths and those implied by the counterfactual perfect-insurance model of all the variables of interest. If anything, the perfect-insurance model generates more volatility than the imperfect-insurance model, as is revealed by the dynamics of the labor market. This suggests that the precautionary motive has tended to stabilize the economy, that is, the aggregate supply effect of the precautionary motive dominated the demand effect during this period.

5. CONCLUDING REMARKS

In this paper, we have provided a general framework aimed at incorporating incomplete insurance and heterogeneous agents in an estimatable New Keynesian dynamic stochastic general equilibrium model. Our theoretical framework relies on a minimal set of assumptions about the extent of risk sharing, under which the cross-sectional distribution of workers converges to a distribution with a potentially large, but always finite, support. Incorporating this distribution into the state-space representation of the model makes it possible to compute its likelihood function and thereby to recover the joint distribution of the structural parameters from the data. Our approach also expands the set of observables to times series of cross-sectional moments (e.g., consumption shares by income quantiles), which by construction cannot be used to discipline New Keynesian models with a representative agent.

In the present paper, we have used our estimated baseline model to assess the relative strength of the aggregate demand and supply effects of time-varying precautionary savings. Our analysis shows that the demand effect has largely contributed to the amplification and propagation of the Great Recession, unlike the previous two recessions where this effect was weak and largely offset by the (stabilizing) supply effect of precautionary savings. Obviously, there are many dimensions other than the precautionary motive in which incomplete insurance and household heterogeneity matter, and thus where the approach that we propose can usefully be applied. One area that immediately comes to mind is the impact of transfer-based fiscal policy and its interactions with other macro policies. [Kaplan and Violante \(2014\)](#) and [McKay and Reis \(2013\)](#) have taken the first steps in that direction, but in models that abstract from the unemployment risk/aggregate demand feedback loop that is likely to be activated following fiscal shocks. Note also that it is straightforward to introduce other assets into our framework, including public debt. Cyclical variations in the public debt cannot a priori be neglected when studying fiscal policy in heterogeneous-agent environments, because the very reason why heterogeneity matters—imperfect insurance and borrowing constraints—makes the economy inherently non-Ricardian, causing the induced changes in public debt to have first-order effects on the equilibrium. We leave these themes for future research.

APPENDIX: PROOFS

PROOF OF PROPOSITION 1. We prove (b) by induction, and then show that (a) is a special case of (b). First, recall that the pooling of assets among employed workers within every family implies that, at every point in time, the \mathbf{n}^W employed workers in the economy all hold assets

$$A = \frac{(1 - s) \int_{\mathbb{R}} x \, d\tilde{\boldsymbol{\mu}}(x, 0) + f \sum_{N \geq 1} \int_{\mathbb{R}} x \, d\tilde{\boldsymbol{\mu}}(x, N)}{\mathbf{n}^W}$$

at the beginning of the consumption–saving stage. Formally, $\boldsymbol{\mu}(a, 0) = \mathbf{n}^W \times \mathbf{1}_{a \geq A}$ at all dates.

We now show that if $\boldsymbol{\mu}(a, N)$ has a unique mass point in a for all $N \leq \check{N}$, with $0 \leq \check{N} \leq \infty$, then it is also the case of $\tilde{\boldsymbol{\mu}}'(a, N)$ and, thus, of $\boldsymbol{\mu}'(a, N) \forall N \leq \check{N} + 1$.

Suppose it is the case for $\boldsymbol{\mu}(a, N)$, and denote by $a(N)$ and $\mathbf{n}(N)$ the wealth level of workers with $N \leq \check{N}$ and their numbers, respectively (so that $\boldsymbol{\mu}(a, N) = \mathbf{n}(N) \times \mathbf{1}_{a \geq a(N)}$, while $\mathbf{n}^W = \mathbf{n}(0)$). All these workers have the same individual state vector (a, N) conditional on N for all $N \leq \check{N}$ and, hence, they all save the same amount: $a^{W'}(N, X) = g_{a^W}(a(N), N, X)$. This implies that the end-of-period wealth distribution $\tilde{\boldsymbol{\mu}}'(a, N)$ also has a unique mass point in a up to \check{N} , that is, $\tilde{\boldsymbol{\mu}}'(a, N) = \mathbf{n}(N) \times \mathbf{1}_{a \geq a^{W'}(N, \cdot)} \forall N \leq \check{N}$. The transition from $\tilde{\boldsymbol{\mu}}'$ to $\boldsymbol{\mu}'$ (see Section 1.6) gives $\boldsymbol{\mu}'(a, N + 1) = (1 - f')\tilde{\boldsymbol{\mu}}'(a, N) = (1 - f')\mathbf{n}^W(N) \times \mathbf{1}_{a \geq a^{W'}(N, \cdot)} \forall 1 \leq N \leq \check{N}$ and $\boldsymbol{\mu}'(a, 1) = s'\tilde{\boldsymbol{\mu}}'(a, 0) = s'\mathbf{n}(0) \times \mathbf{1}_{a \geq a^{W'}(0, \cdot)}$, so that $\boldsymbol{\mu}'(a, N)$ has a unique mass point in N for all $N \in [1, \check{N} + 1]$. From this recursion, since $\boldsymbol{\mu}(a, 0)$ necessarily has a unique mass point in a at the beginning of time (whatever $\tilde{\boldsymbol{\mu}}_0$), so do $\boldsymbol{\mu}(a, 0)$ and $\boldsymbol{\mu}(a, 1)$ one period after, $\boldsymbol{\mu}(a, 0)$, $\boldsymbol{\mu}(a, 1)$, and $\boldsymbol{\mu}(a, 2)$ two periods after, and so forth. The statement of part (b) of the proposition directly follows. The proof of (a) is a special case of the proof of (b): if $\tilde{\boldsymbol{\mu}}_0$ has a unique mass point in a for all N up to ∞ , then so does $\boldsymbol{\mu}$ up to $N = \infty$ at the initial date, hence $\boldsymbol{\mu}'$ (by the induction argument above), and the same is true at every period. \square

PROOF OF PROPOSITION 2. Section A.2 of the Technical Appendix provides a detailed derivation of workers’ optimality conditions and how they lead to (30)–(31), so we only provide the main steps of the proof here. If $\boldsymbol{\mu}$ has a unique mass point in $a \forall N \geq 0$, then it is summarized by $(a(N), n(N))_{N \geq 0}$ (i.e., the wealth levels and numbers of family members for each N). We can thus rewrite the workers’ problem as

$$\begin{aligned} & \hat{V}^W((a(N), n(N))_{N \geq 0}, X) \\ &= \max_{(a^{W'}(N), c^W(N))_{N \geq 0}} \sum_N n(N) u(c^W(N) - hc^W(N)) + \beta^W \mathbb{E}_X \hat{V}^W((a'(N), n'(N))_{N \geq 0}, X), \end{aligned}$$

such that $a^{W'}(N) \geq ae^z$ and $a^{W'}(N) + c^W(N) = \mathbf{1}_{N=0}(1 - \tau)w + \mathbf{1}_{N \geq 1}b^u e^z + (1 + r)a(N)$, $N \in \mathbb{Z}_+$. The solution to this problem is characterized by $\#\mathbb{Z}_+$ first-order and envelope

conditions. Combining those generates the $\#Z_+$ Euler conditions

$$\begin{aligned} & u_c(c^W(0) - hc^W(0)) \\ & \geq \beta^W \mathbb{E}_X \{ (1+r')[(1-s')u_c(c^{W'}(0) - hc^{W'}(0)) + s'u_c(c^{W'}(1) - hc^{W'}(1))] \} \quad \text{for } N=0 \end{aligned}$$

and

$$\begin{aligned} & u_c(c^W(N) - hc^W(N)) \\ & \geq \beta^W \mathbb{E}_X \{ (1+r')[f'u_c(c^{W'}(0) - hc^{W'}(0)) + (1-f')u_c(c^{W'}(N+1, X') \\ & \quad - hc^{W'}(N+1))] \} \quad \text{for } N \geq 1, \end{aligned}$$

where the inequalities are strict any time the debt limit is binding. We can rearrange the previous expressions as $\mathbb{E}_X[M^{W'}(N)(1+r')] \leq 1$, $N=0, 1, 2, \dots$, where the $M^{W'}(N)$ s are stated in the proposition. \square

PROOF OF PROPOSITION 3. We provide the main steps of the proof and leave the full proof to the Technical Appendix. The proof is by contradiction. In the economy without aggregate shocks total factor productivity (TFP) grows deterministically at a rate μ_z , the real interest rate is given by $1 + \bar{r} = e^{\sigma\mu_z}/\beta^F$, and we denote by $\hat{c}^W(N) = c^W(N)e^{-z}$ and $\hat{a}^{W'}(N) = a^{W'}(N)e^{-z}$ the detrended consumption and assets of a type- N worker. If unemployed workers never faced a binding debt limit, their Euler equation (as written in the proof of Proposition 2 above) would always hold with equality. Noting that in the absence of aggregate shocks we have $u_c(c^W(N) - hc^W(N)) = (\hat{c}^W(N)(e^{\mu_z} - h))^{-\sigma} e^{-\sigma z + \sigma\mu_z}$, the Euler condition can be written as (after some calculations)

$$\hat{c}^W(N)^{-\sigma} = (\beta^W/\beta^F)(f\hat{c}^W(0)^{-\sigma} + (1-f)\hat{c}^W(N+1)^{-\sigma}) \quad \text{for } N=1, 2, \dots$$

The latter expression defines a recursion on $x(N) \equiv \hat{c}^W(N)^{-\sigma}$, and it can be shown that $x(N) \rightarrow +\infty$ as $N \rightarrow +\infty$, so that $\hat{c}^W(N) \rightarrow 0$ as $N \rightarrow +\infty$. On the other hand, the budget constraint of a type- N worker, expressed in detrended form, is given by

$$\hat{a}^{W'}(N) + \hat{c}^W(N) = b^u + (e^{(\sigma-1)\mu_z}/\beta^F)\hat{a}^W(N) \quad \text{for } N=1, 2, \dots$$

In the limit, with $\hat{c}^W(+\infty) = 0$, it must be that $\hat{a}^{W'}(+\infty) = \underline{a}^{\text{nat}}$, which contradicts Assumption 2 (according to which the debt limit is *strictly* tighter than the natural limit). \square

REFERENCES

- Aiyagari, S. (1994), "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics*, 109, 659–684. [437, 456]
- Ajello, A. (2014), "Financial intermediation, investment dynamics and business cycle fluctuations." Working Papers 2014, Board of Governors of the Federal Reserve System. [437]

- Aruoba, S. B. and F. Schorfheide (2011), “Sticky prices versus monetary frictions: An estimation of policy trade-offs.” *American Economic Journal: Macroeconomics*, 3, 60–90. [436]
- Beaudry, P., D. Galizia, and F. Portier (2014), “Reconciling Hayek’s and Keynes views of recessions.” Working Papers 20101, NBER. [439]
- Blanchard, O. and J. Galí (2010), “Labor markets and monetary policy: A new Keynesian model with unemployment.” *American Economic Journal: Macroeconomics*, 2, 1–30. [438]
- Boivin, J. and M. Giannoni (2006), “DSGE models in a data-rich environment.” Technical working papers, NBER. [436]
- Calvo, G. A. (1983), “Staggered prices in a utility-maximizing framework.” *Journal of Monetary Economics*, 12, 383–398. [440, 447]
- Canova, F. and L. Sala (2009), “Back to square one: Identification issues in DSGE models.” *Journal of Monetary Economics*, 56, 431–449. [436]
- Challe, E. and X. Ragot (2014), “Precautionary saving over the business cycle.” *Economic Journal* (forthcoming). [437, 449, 463]
- Chatterjee, S., D. Corbae, M. Nakajima, and J. V. Ríos-Rull (2007), “A quantitative theory of unsecured consumer credit with risk of default.” *Econometrica*, 75, 1525–1589. [436]
- Chodorow-Reich, G. and L. Karabarbounis (2014), “The cyclicity of the opportunity cost of employment.” Working paper, Harvard University. [463]
- Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005), “Nominal rigidities and the dynamic effects of a shock to monetary policy.” *Journal of Political Economy*, 113, 1–45. [438, 440]
- Christiano, L. J., R. Motto, and M. Rostagno (2014), “Risk shocks.” *American Economic Review*, 104, 27–65. [450]
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2012), “Innocent bystanders? Monetary policy and inequality in the U.S.”, Working Papers 18170, NBER. [439]
- Coibion, O. and Y. Gorodnichenko (2013), “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation.” Working Papers 19598, NBER. [471]
- Davila, J., J. H. Hong, P. Krusell, and J. Ríos Rull (2012), “Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks.” *Econometrica*, 80, 2431–2467. [436]
- DeJong, D. N., B. F. Ingram, and C. H. Whiteman (2000), “A Bayesian approach to dynamic macroeconomics.” *Journal of Econometrics*, 98, 203–223. [436]
- Del Negro, M., M. P. Giannoni, and F. Schorfheide (2015), “Inflation in the great recession and new Keynesian models.” *American Economic Journal: Macroeconomics*, 7, 168–196. [471]

- Del Negro, M., F. Schorfheide, F. Smets, and R. Wouters (2007), "On the fit of new Keynesian models." *Journal of Business & Economic Statistics*, 25, 123–143. [436]
- Den Haan, W., P. Rendahl, and M. Riegler (2015), *Unemployment (Fears) and Deflationary Spirals*. Tech. rep., CEPR. [439]
- Evans, G. W. and S. Honkapohja (2005), "MD interview: An interview with Thomas J. Sargent." *Macroeconomic Dynamics*, 106, 561–583. [436]
- Evvett, I. W. (1991), "Implementing Bayesian methods in forensic science." In *4th Valencia International Meeting on Bayesian Statistics*. [467]
- Galí, J. (2010), "Monetary policy and unemployment." In *Handbook of Monetary Economics*, Vol. 3, Chapter 10 (B. M. Friedman and M. Woodford, eds.), 487–546, Elsevier. [438, 439, 441]
- Gertler, M., L. Sala, and A. Trigari (2008), "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining." *Journal of Money, Credit and Banking*, 40, 1713–1764. [439]
- Gornemann, N., K. Kuester, and M. Nakajima (2012), "Monetary policy with heterogeneous agents." Working Papers 12-21, Federal Reserve Bank of Philadelphia. [439]
- Guerrieri, V. and G. Lorenzoni (2011), "Credit crises, precautionary savings, and the liquidity trap." Working Papers 17583, NBER. [439]
- Guerrón-Quintana, P., J. Fernández-Villaverde, and J. F. Rubio-Ramírez (2010), "The new macroeconometrics: A Bayesian approach." In *The Oxford Handbook of Applied Bayesian Analysis*, Chapter 14 (T. O'Hagan and M. West, eds.), 366–546, Oxford University Press. [450]
- Gust, C., J. D. Lopez-Salido, and M. E. Smith (2012), "The empirical implications of the interest-rate lower bound." Discussion Papers 9214, C.E.P.R. [450]
- Hall, R. E. (2005), "Employment fluctuations with equilibrium wage stickiness." *American Economic Review*, 95, 50–65. [450]
- Heathcote, J., F. Perri, and G. L. Violante (2010), "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics*, 13, 15–51. [462]
- Heathcote, J. and F. Perri (2014), "Wealth and volatility." Report, Federal Reserve Bank of Minneapolis. [439]
- Heer, B. and A. Maussner (2010), "Inflation and output dynamics in a model with labor market search and capital accumulation." *Review of Economic Dynamics*, 13, 654–686. [439, 440]
- Iacoviello, M. (2005), "House prices, borrowing constraints, and monetary policy in the business cycle." *American Economic Review*, 95, 739–764. [466]
- Iacoviello, M. (2008), "Household debt and income inequality, 1963–2003." *Journal of Money Credit and Banking*, 40, 929–965. [437]

- Ireland, P. N. (2004), “A method for taking models to the data.” *Journal of Economic Dynamics and Control*, 28, 1205–1226. [436]
- Ireland, P. N. (2011), “A new Keynesian perspective on the great recession.” *Journal of Money, Credit and Banking*, 43, 31–54. [468]
- Justiniano, A., G. Primiceri, and A. Tambalotti (2015), “Household leveraging and deleveraging.” *Review of Economic Dynamics*, 18, 3–20. [437]
- Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010), “Investment shocks and business cycles.” *Journal of Monetary Economics*, 57, 132–145. [436]
- Kaplan, G. and G. L. Violante (2014), “A model of the consumption response to fiscal stimulus payments.” *Econometrica* (forthcoming). [463, 472]
- Kaplan, G., G. L. Violante, and J. Weidner (2014), “The wealthy hand-to-mouth.” *Brookings Papers on Economic Activity* (forthcoming). [463]
- Kehoe, P., V. Midrigan, and E. Pastorino (2014), “Debt constraints and unemployment.” Report, New York University. [439]
- Kiyotaki, N. and J. Moore (1997), “Credit cycles.” *Journal of Political Economy*, 105, 211–248. [466]
- Krueger, D., K. Mitman, and F. Perri (2015), “Macroeconomics and heterogeneity, including inequality.” Report, University of Pennsylvania. [438]
- Krusell, P., T. Mukoyama, and A. Sahin (2010), “Labour-market matching with precautionary savings and aggregate fluctuations.” *Review of Economic Studies*, 77, 1477–1507. [439]
- Krusell, P. and A. A. Smith (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106, 867–896. [437, 438, 439, 440, 467]
- Leduc, S. and Z. Liu (2014), “Uncertainty shocks are aggregate demand shocks.” Working Papers 2012-10, Federal Reserve Bank of San Francisco. [439]
- Lubik, T. A. and F. Schorfheide (2004), “Testing for indeterminacy: An application to U.S. monetary policy.” *American Economic Review*, 94, 190–217. [436]
- McKay, A., E. Nakamura, and J. Steinsson (2015), “The power of forward guidance revisited.” Report. [439]
- McKay, A. and R. Reis (2013), “The role of automatic stabilizers in the U.S. business cycle.” Discussion Papers 9454, CEPR. [439, 472]
- Miao, J. (2006), “Competitive equilibria of economies with a continuum of consumers and aggregate shocks.” *Journal of Economic Theory*, 128, 274–298. [445]
- Mortensen, D. T. and C. A. Pissarides (1994), “Job creation and job destruction in the theory of unemployment.” *Review of Economic Studies*, 61, 397–415. [440]
- Nakajima, M. (2012), “Business cycles in the equilibrium model of labor market search and self-insurance.” *International Economic Review*, 53, 399–432. [439]

- Oh, H. and R. Reis (2012), “Targeted transfers and the fiscal response to the great recession.” *Journal of Monetary Economics*, 59, S50–S64. [439]
- Otrok, C. (2001), “On measuring the welfare cost of business cycles.” *Journal of Monetary Economics*, 47, 61–92. [436]
- Ravn, M. and V. Sterk (2013), “Job uncertainty and deep recessions.” London. Report, University College. [439]
- Reiter, M. (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control*, 33, 649–665. [440]
- Rendhal, P. (2014), “Fiscal policy in an unemployment crisis.” Report. [439]
- Sbordone, A. M. (2007), “Globalization and inflation dynamics: The impact of increased competition.” In *International Dimensions of Monetary Policy*, 547–579, NBER, NBER Chapters. [471]
- Schorfheide, F. (2000), “Loss function-based evaluation of DSGE models.” *Journal of Applied Econometrics*, 15, 645–670. [436]
- Shimer, R. (2005), “The cyclical behavior of equilibrium unemployment and vacancies.” *American Economic Review*, 95, 25–49. [462]
- Shimer, R. (2012), “Reassessing the ins and outs of unemployment.” *Review of Economic Dynamics*, 15, 127–148. [462]
- Smets, F. and R. Wouters (2007), “Shocks and frictions in US business cycles: A Bayesian DSGE approach.” *American Economic Review*, 97, 586–606. [436, 438, 440, 444, 460]
- Trigari, A. (2009), “Equilibrium unemployment, job flows, and inflation dynamics.” *Journal of Money, Credit and Banking*, 41, 1–33. [438, 440]
- Walsh, C. E. (2005), “Labor market search, sticky prices, and interest rate policies.” *Review of Economic Dynamics*, 8, 829–849. [438, 441]
- Woodford, M. (2003), *Interest and Prices*. Princeton University Press, Princeton, NJ. [438]

Co-editor Rosa L. Matzkin handled this manuscript.

Manuscript received 10 May, 2016; final version accepted 8 September, 2016; available online 28 October, 2016.