Micro and Macro Uncertainty

Andreas Schaab*

November 13, 2020

Job Market Paper

Please find the latest version here.

Abstract

Uncertainty rises sharply during economic downturns at both the micro and macro level. Leveraging a new solution method, I study the interaction between micro and macro uncertainty in a globally solved Heterogeneous Agent New Keynesian (HANK) model with aggregate risk, counter-cyclical unemployment risk, and a zero lower bound (ZLB) constraint on monetary policy. The interaction with micro uncertainty emerges as the dominant transmission channel of macro uncertainty. The overall effect of uncertainty on economic activity is substantially amplified. My model also generates endogenous spikes in uncertainty during bad times as the economy is pushed towards the ZLB. In general equilibrium, a feedback loop emerges that gives rise to an "Uncertainty Multiplier": A contraction in economic activity spurs endogenous uncertainty about the future, which depresses aggregate demand further. The model matches the skewness and kurtosis exhibited by macro uncertainty in the data even in the absence of exogenous second-moment shocks. The interplay between micro and macro uncertainty has ramifications for the nature of zero lower bound spells, the welfare cost of business cycles, and the effectiveness of stabilization policy.

*Harvard University. Email: aschaab@hbs.edu. I am indebted to Emmanuel Farhi, Xavier Gabaix, Matteo Maggiori and Ludwig Straub for their mentorship, guidance, and support throughout my Ph.D. I am also grateful to Isaiah Andrews, Martin Aragoneses, John Campbell, Gabriel Chodorow-Reich, Christopher Clayton, Edward Glaeser, Sam Hanson, Xavier Jaravel, Moritz Lenel, Alisdair McKay, Benjamin Moll, Chris Rycroft, Andrei Shleifer, Adi Sunderam, Jeremy Stein, Paul Tucker, Stephen Terry, Gianluca Violante, Samuel Young, Allen Zhang, as well as seminar participants at Harvard University and HBS for many helpful conversations and comments.

Uncertainty at both the micro and macro level rises sharply during economic downturns. At the micro level, households face elevated unemployment risk during crises (e.g. Storesletten et al. (2004), Shimer (2005) and Guvenen et al. (2014)), and firms experience more volatile sales, cash-flow and productivity growth (e.g. Kehrig (2015) and Bloom et al. (2018)). At the macro level, recessions are associated with higher stock market volatility and increased variability in GDP growth (e.g. Bloom (2014)). Much previous work has studied the implications of uncertainty at either the micro or the macro level. A framework that allows for the joint determination of micro and macro uncertainty has thus far remained elusive, in part because it poses considerable methodological challenges. Modeling uncertainty at the micro level requires cross-sectional heterogeneity, while uncertainty at the macro level presupposes aggregate risk.[1]

In this paper, I make two main contributions. First, I show that accounting for the interaction between micro and macro uncertainty is crucial to characterize the role uncertainty plays in business cycle fluctuations. The dominant transmission channel of macro uncertainty is its interaction with micro uncertainty. Indeed, in general equilibrium (GE) a strong feedback loop emerges between micro and macro uncertainty. As a result, recessions are associated with large endogenous spikes in uncertainty. My second contribution is methodological. I develop a new global solution method for heterogeneous-agent macro models. Leveraging its power, I globally solve a Heterogeneous Agent New Keynesian ("HANK") model with aggregate risk, counter-cyclical unemployment risk, and a zero lower bound (ZLB) constraint on monetary policy. While a rapidly growing body of work has developed first-order perturbation methods to solve heterogeneous-agent macro models with aggregate risk (see Reiter (2009), Winberry (2020), Ahn et al. (2017), Boppart et al. (2018), and Auclert et al. (2019)), progress on global solution methods has been slower.[2]

I start my discussion in Section 1 by studying the transmission mechanism of macro uncertainty in an illustrative two-period model. Households solve a consumption-savings problem, facing both aggregate wage risk and idiosyncratic unemployment risk. At the heart of my analysis is the empirically motivated assumption that a household's employment transition probabilities vary with aggregate economic activity. Indeed, the job separation (finding) rate in the data is strongly counter-cyclical (pro-cyclical): Recessions are times of heightened micro uncertainty over households' job prospects (e.g. Storesletten et al. (2004) and Shimer (2012)). The transmission of macro uncertainty in this setting is governed by a set of direct (partial equilibrium) and indirect (general equilibrium) channels, borrowing the language from Kaplan et al. (2018). The direct effects of uncertainty are those that arise from a change in households' beliefs about future shock realizations holding prices constant. Indirect effects arise in general equilibrium as prices respond to the change in household behavior induced by direct effects.

[1] Bloom et al. (2018) develop a model with cross-sectional firm heterogeneity and study exogenous variation in the micro and macro uncertainty that firms face. My focus, on the other hand, is on the endogenous interaction between uncertainty at the micro and macro level, which is not modeled in their paper.

[2] Variants of the original Krusell and Smith (1998) algorithm largely remain the method of choice. Notable exceptions include Fernández-Villaverde et al. (2019) and Pröhl (2019).

When there is no interaction with micro uncertainty, the only direct effect of macro uncertainty on household behavior to second order is the standard Kimball (1990) precautionary savings motive. A mean-zero spread in next period's aggregate wage rate exposes employed households to symmetric risk in labor income. Since this effect is marginal, its relevance to households is captured by differentiating marginal utility and, in particular, is proportional to $u'''(\cdot)$. Unlike aggregate wage risk, unemployment risk at the micro level takes the form of discrete and asymmetric jumps. Its relevance is not captured by a derivative of marginal utility but, rather, by the jump scaling factor $u'(c^E) - u'(c^U)$, where $u'(c^E)$ and $u'(c^U)$ are marginal utility conditional on employment and unemployment, respectively. It is in this sense that the prospect of job loss is akin to an idiosyncratic rare disaster from the perspective of households.

When we account for its interaction with unemployment risk, macro uncertainty operates through a set of novel channels. Direct interaction effects emerge because a mean-zero spread in the aggregate shock now also implies increased dispersion in employment outcomes. In general equilibrium, an indirect interaction effect emerges as well: The direct effects of micro and macro uncertainty on household behavior lead to an initial fall in consumer spending and aggregate demand. As a result, the demand for labor falls, which consequently raises (lowers) the job separation (finding) rate. For an employed household, an increase in the job separation rate both lowers expected earnings and increases their variance. The former elicits a consumption-smoothing response and the latter a precautionary savings motive. These novel interaction effects tend to increase households' desired savings in response to an increase in macro uncertainty.

I show that these direct and indirect interaction effects are proportional to the jump scaling factor $u'(c^E) - u'(c^U)$ that is characteristic of uncertainty at the micro level. Households in this setting respond to uncertainty at the macro level not so much because it implies a mean-zero spread in aggregate wages but rather because it translates into disaster risk at the micro level. These results suggest qualitatively that accounting for an interaction with micro uncertainty can substantially alter the transmission mechanism of macro uncertainty.

To evaluate these implications quantitatively, I build a business cycle model in which households face uninsurable unemployment risk in the tradition of Huggett (1993), Aiyagari (1994), McKay et al. (2016) and Kaplan et al. (2018). I build on the benchmark HANK model of Kaplan et al. (2018) along three dimensions. First, I model a zero lower bound (ZLB) constraint on monetary policy, thus explicitly introducing a source of non-linearity at the macro level. Second, I allow for aggregate risk in the form of discount rate shocks.[3] Third, and most importantly, I account for the cyclicality in the job finding and separation rates so that the unemployment risk faced by households varies over the business cycle. I estimate the sensitivity of employment transition rates to changes in economic activity in Current

[3]Discount rate shocks are a popular proxy for aggregate demand shocks (e.g. Basu and Bundick (2017) and Auclert et al. (2020)). I have also solved variants of this model with alternative demand and supply (TFP) shocks. In the presence of nominal rigidities, TFP shocks tend to generate counter-cyclical price inflation (e.g. Gali (2015)). Since the source of macro non-linearity in this model is the ZLB constraint on monetary policy, it is important that inflation is pro-cyclical, which is the case in the presence of aggregate discount rate shocks.

Population Survey (CPS) micro data. This approach follows a long tradition of constructing gross flow data for employment transitions from micro data.[4] By using the resulting estimates directly in the model, I take as given the link between employment transitions and economic activity.[5] The sense in which my model generates an endogenous link between micro and macro uncertainty is through economic activity: heightened macro uncertainty depresses aggregate demand which in turn raises unemployment risk. At the micro level, therefore, households face uncertainty over future job prospects. At the macro level, households are exposed to uncertainty over asset prices, interest rates and the aggregate determinants of disposable income, such as the wage rate, transfers, and unemployment insurance (UI) payments. The crucial feature of my model is that uncertainty at the micro level is high precisely at the same time as uncertainty at the macro level.

My first main result is that the transmission mechanism of macro uncertainty in the quantitative model differs starkly from that in a Representative Agent New Keynesian ("RANK") benchmark.[6] In the RANK model, the direct precautionary savings response to heightened aggregate uncertainty contributes the largest share to the overall effect on consumption. Using a conservative relative risk aversion coefficient of $\gamma = 2$ in my calibration, this direct precautionary effect is unsurprisingly small. Therefore, the overall effect of macro uncertainty on economic activity in the RANK baseline is quite modest.[7]

In the quantitative model, this direct effect of macro uncertainty is relatively muted. Instead, the interaction with micro uncertainty emerges as the dominant driver of transmission. The quantitative analysis therefore corroborates the importance of the interaction effects identified in the two-period model. Household behavior is more responsive to uncertainty at the macro level precisely because it translates into unemployment risk. Similarly, and in the spirit of Kaplan et al. (2018), other indirect channels working through portfolio returns and the aggregate determinants of disposable income become important relative to the direct precautionary channel. Therefore, the overall effect of macro uncertainty on economic activity is large in my model precisely because it works through micro uncertainty and other indirect channels. I show that the peak response of output to an increase in macro uncertainty is 5 to 8 times larger than in the associated RANK benchmark. The interaction between micro and

[4]See for example Marston et al. (1976), Abowd and Zellner (1985), Darby et al. (1985), Darby et al. (1986), Poterba and Summers (1986), Blanchard et al. (1990), Shimer (2005), Fujita and Ramey (2009), Elsby et al. (2009), and Shimer (2012).

[5]These empirical estimates based on U.S. micro data are intended as a reduced-form summary of the state dependence in employment transition rates that would result from a search-and-matching micro-foundation. Explicitly implementing a search-and-matching block is beyond the scope of this paper and left for future work.

[6]The focus of my paper is on macroeconomic uncertainty that arises endogenously. However, isolating and decomposing the effects of endogenous uncertainty is challenging in a setting with aggregate risk. To be able to study the transmission mechanism of macro uncertainty, I consider an exogenous shock to fundamental risk (similar to Bloom (2009), Basu and Bundick (2017), Bloom et al. (2018) and Bayer et al. (2019)). The assumption implicit in this strategy is that the transmission mechanism of an exogenous increase in the volatility of discount rate shocks is sufficiently similar to that of an endogenous increase in the volatility of economic activity.

[7]This is consistent with the results in Basu and Bundick (2017). They show that a fundamental risk shock induces co-movement in output, consumption, investment and hours, but they impart households with Epstein-Zin preferences and a relative risk aversion coefficient of 80 to generate quantitatively meaningful responses.

macro uncertainty is the main source of this amplification.

An overarching theme of my analysis is that the behavior of uncertainty changes drastically during economic crises. I show that the peak decline in output in response to a given increase in macro uncertainty is 50% larger when the economy is already at the cusp of the ZLB than during normal times. Close to the ZLB, the relationship between economic activity and aggregate risk exhibits a degree of negative skewness. As a result, a mean-zero spread in aggregate risk leads to a contraction of economic activity in expectation. The implications of such skewness at the macro level are not unlike the skewness households face at the micro level. Indeed, I show that non-linearity at the macro level strongly interacts with non-linearity at the micro level: The relative importance of its interaction with micro uncertainty in the transmission of macro uncertainty further rises during economic crises. Identifying the importance of the ZLB crisis region for the behavior of uncertainty in my model is only possible because I use a global solution method.[8]

My second main result is that the interaction between micro and macro uncertainty has implications not only for the transmission mechanism of macro uncertainty but also for its endogenous responsiveness to changes in economic activity. I show that macro uncertainty in my model rises endogenously during economic downturns. That is, recessions are accompanied by endogenous spikes in uncertainty even in the absence of exogenous second-moment shocks. Crucially, the sensitivity of macro uncertainty to economic activity is dampened substantially when I shut off its interaction with micro uncertainty. Concretely, I show that endogenous macro uncertainty is 4 times more responsive to a negative, first-moment discount rate shock when I account for the interaction with micro uncertainty. Indeed, when I hold unemployment risk constant, macro uncertainty hardly responds at all to discount rate shocks during normal times.

Macroeconomic uncertainty responds endogenously to changes in economic activity through two main channels in my model. The first channel centers around the zero lower bound constraint. When the economy is at the ZLB, monetary policy can no longer accommodate negative demand shocks, whose effects on economic activity are consequently amplified. Even when the nominal interest rate is still positive, expansionary monetary policy moves the economy closer to the ZLB, thus raising the likelihood that policy will be constrained in the future. A given increase in macro uncertainty resulting from proximity to the ZLB has a larger effect on aggregate demand when its transmission works through micro uncertainty, thus pushing the economy even closer to the ZLB. To capture this economic force, it is crucial to employ a global solution method.

[8]Overall, these results highlight that uncertainty can have large effects on consumption even in a setting where household behavior is not perfectly forward-looking. It is well known that the effective planning horizon of households in an incomplete markets setting is shortened in the presence of borrowing constraints (e.g. McKay et al. (2016) and Kaplan et al. (2018)). Households directly at the borrowing constraint consume their disposable income "hand-to-mouth". But even those households that are close to but not at the constraint only try to smooth consumption until they expect to reach the constraint. Unlike the representative household in the RANK benchmark, most households in my model are close to or at their borrowing constraint and have relatively short planning horizons.

The second channel results from the counter-cyclicality in households' average marginal propensity to consume (MPC). As economic activity contracts, households become unemployed and draw down their liquid cash buffers, thus moving closer to their borrowing constraints at the micro level. The prevalence of "hand-to-mouth" behavior grows, which implies an increase in the average household's MPC. Consumer spending, and by implication aggregate activity, thus become more sensitive to further demand shocks, which represents an increase in macro uncertainty. By contrast, the counter-cyclicality of the average MPC is considerably dampened when employment transition rates are held constant over the business cycle. The standard Krusell and Smith (1998) algorithm struggles to account for these shifts in the household distribution. The global solution method developed in this paper therefore plays a key role in my ability to study endogenous uncertainty spikes in this model.

In general equilibrium, therefore, a strong feedback loop emerges between uncertainty and economic activity: When activity contracts, uncertainty about the future rises, which itself depresses aggregate demand further. This feedback loop can be instructively characterized as an "Uncertainty Multiplier", which measures how much endogenous amplification there is in macro uncertainty. I show that this Uncertainty Multiplier is high precisely when we account for the cyclicality of unemployment risk, and when the economy is already in a recession. Indeed, this feedback loop between uncertainty and activity allows my model to match the time series moments of various macro uncertainty proxies in the data without requiring exogenous second-moment shocks: Macro uncertainty in the model is strongly counter-cyclical, highly persistent and exhibits large positive skewness and kurtosis. Overall, uncertainty emerges as both a driver and a byproduct of business cycle fluctuations. Accounting for the interaction between micro and macro uncertainty is thus crucial to understand the broader role that uncertainty plays in macroeconomic fluctuations.

The interaction between micro and macro uncertainty has far-reaching implications for central questions in business cycle analysis, several of which I discuss in Sections 6 and 7, and in the appendix. First, ZLB spells become more frequent and more persistent. I also show that the interaction between the ZLB and a paradox of thrift dynamic, by which households increase savings in anticipation of reaching the ZLB, is amplified. Second, an interplay between uncertainty at the micro and macro level provides a new perspective on the welfare cost of business cycles. I show that the implied share of consumption households are willing to forego to instead "live" in a representative-agent economy is 3.9%, or about two orders of magnitude larger than the original estimates in Lucas (1987) and Lucas (2003). In the appendix, I study a series of policy experiments and show that stabilization policy in my setting operates through a novel set of micro and macro uncertainty channels.

Methodological Contribution. The methodological contribution of my paper is a new global solution method for heterogeneous-agent macro models with aggregate risk. I show that this new method is key to solve the quantitative model with time-varying unemployment risk and occasionally-binding ZLB constraint, and in turn study the interaction between micro and macro uncertainty. The main challenge in numerically solving my model is that the entire cross-sectional distribution of agents, an infinite-dimensional object, becomes part of the aggregate state space. Let $\boldsymbol{x}_t$ denote the vector of idiosyncratic state variables and $g_t(\boldsymbol{x})$ the cross-sectional distribution of agents. In this paper, I work with finite-dimensional distribution approximations of the form

$$F(\alpha_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x}).$$

For illustration, it is easiest to think of F as a set of basis functions that are parameterized by the time-varying $\alpha_t \in \mathbb{R}^N$. While representations of the form $F(\alpha_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x})$ are commonly used in the context of local perturbation methods, they have proven intractable in the context of global methods due to the curse of dimensionality.[9] I make three contributions that help overcome this challenge:[10]

1. Most global methods currently in use, such as the seminal Krusell and Smith (1998) algorithm, work with a finite set of moments to approximate the distribution of agents. The costliest step of these algorithms is finding an internally consistent law of motion for these moments. I show in Section 3.3 that, for a large class of models, the coefficients α_t follow a diffusion process with drift $\mu_{\alpha,t}$ and volatility $\sigma_{\alpha,t}$, and I derive analytical formulas for these objects that can be easily computed. In this sense, and in sharp contrast to the Krusell and Smith (1998) algorithm, finding the consistant law of motion incurs almost no increase in numerical complexity.[11]

2. While $F(\cdot)$ can be chosen from a parametric family, I develop a non-parametric algorithm in Section 3.4 that delivers substantial efficiency gains especially when the idiosyncratic state space of agents is high-dimensional.

3. For most economic applications of interest, accurate approximations of the distribution will require high-dimensional $F(\alpha_t)(\boldsymbol{x})$. Global methods will therefore quickly encounter the curse of dimensionality. In Schaab and Zhang (2020), we develop an adaptive sparse grid library for solving partial differential equations that can overcome the curse of dimensionality in high dimensions. Using this library and leveraging the results developed in this paper, I can solve the benchmark Krusell and Smith (1998) model with a $F(\alpha_t)(\boldsymbol{x})$ representation in over 20 dimensions, that is $\alpha_t \in \mathbb{R}^{20}$.

Literature Review. This paper is most directly related to a long literature studying the role of uncertainty in business cycle fluctuations. One prominent strand of this literature, of

[9]My paper builds on the important contribution of Winberry (2020) who uses a distribution representation of this form in the context of a local perturbation method.

[10]I provide additional details for and discuss the algorithmic and computational aspects of my methodological contribution in a separate Numerical Appendix that can be found here.

[11]I build on the closely related contribution of Ahn et al. (2017) and generalize this state space reduction technique to settings where the cross-sectional distribution of agents itself is stochastic. Asset pricing models with portfolio choice problems are typically of this kind. See Appendix E for an example.

which Bernanke (1983) is an early example, follows the seminal contribution of Bloom (2009) and asks whether uncertainty can drive business cycles. This group of papers considers the implications of an exogenous increase in micro or macro uncertainty, prompting a precautionary savings response among households (Leduc and Liu (2016), Basu and Bundick (2017), Bayer et al. (2019)), a wait-and-see response by firms (Bloom (2009), Bloom et al. (2018)) or a tightening of financial constraints (Gilchrist et al. (2014), Arellano et al. (2019)).[12] An alternative approach to study the effects of uncertainty on economic activity uses vector autoregression (VAR) estimates (see for example Bloom (2009), Ludvigson et al. (2015) or Basu and Bundick (2017)). Relative to this literature, I show that accounting for the interaction between micro and macro uncertainty qualitatively changes the transmission mechanism of an uncertainty shock, substantially amplifying its effect on activity.[13]

A largely distinct strand of literature argues that uncertainty arises endogenously as a byproduct of economic crises.[14] Several channels have been proposed through which a contraction in economic activity may spur uncertainty: When economic activity falls, economic agents interact less frequently, stifling the spread of information (Van Nieuwerburgh and Veldkamp (2006), Fajgelbaum et al. (2017), Straub and Ulbricht (2017)); policymakers may resort to adopting untested policies, raising uncertainty (Pástor and Veronesi (2013)); firms may take riskier and more experimental actions during bad times (Bachmann et al. (2011)). The endogenous responsiveness of uncertainty to economic activity in my model works largely through two channels: the ZLB, and counter-cyclical MPCs. Plante et al. (2018) similarly make the argument that, during bad times as the economy approaches the ZLB, policy becomes further incapacitated, which raises uncertainty.[15]

My focus on time-varying micro uncertainty in the form of counter-cyclical unemployment risk is shared by a large body of work that documents the cyclicality of earnings risk and employment transitions in the data (e.g. Storesletten et al. (2004), Shimer (2005) and Guvenen et al. (2014)), and studies its implications analytically and quantitatively (e.g. Ravn and Sterk (2016), Schmidt (2016), McKay (2017), Acharya and Dogra (2020)). In this paper, I study the interaction between unemployment risk and macro uncertainty: Households respond to an increase in macro uncertainty in large part because it translates into micro uncertainty. Patterson (2019) documents systematic heterogeneity in household exposure to cyclical earnings risk. Taking into account heterogeneous incidence is left for future work.

My paper also adds to the burgeoning heterogeneous-agent New Keynesian (HANK) literature.[16] This is the first paper of this class that studies macroeconomic uncertainty and

[12] See Fernández-Villaverde et al. (2015) for a quantitative analysis of a policy uncertainty shock.

[13] There is also a large literature that tries to measure uncertainty in the data. See Bloom (2014) for an overview.

[14] A smaller body of work seeks to determine whether the counter-cyclicality of uncertainty is a result of exogenous shocks or rather an endogenous response. See for example Ludvigson et al. (2015) and Berger and Vavra (2019).

[15] While they make this point in a representative-agent New Keynesian model, I show that accounting for cross-sectional heterogeneity and, in particular, the interaction between micro and macro uncertainty is crucial.

[16] See Oh and Reis (2012), Guerrieri and Lorenzoni (2017), McKay and Reis (2016), McKay et al. (2016), Werning (2015), Challe et al. (2017), Kaplan et al. (2018), Auclert et al. (2019), Auclert et al. (2018), Auclert et al.

features an occasionally binding macro constraint; both of these features require a higher-order or indeed global solution method.[17] My main contribution to this literature is to show that the transmission mechanism of uncertainty changes qualitatively when we account for cross-sectional household heterogeneity: the effect through micro uncertainty becomes the new dominant transmission channel of a macro uncertainty shock.

Finally, I build on an extensive body of work that has developed global solution methods for heterogeneous-agent macro models.[18] My paper also builds on the important contributions of Winberry (2020) and Ahn et al. (2017) who propose a similar finite-dimensional distribution representation as I do in the context of local perturbation methods.

1 Illustrative Two-Period Example

The goal of this section is to develop intuition for the economic mechanism driving my results: the interaction between micro and macro uncertainty and its implications for household behavior. I characterize the effect of uncertainty on household consumption in a stylized two-period model of consumption and savings decisions. Households in this setting face unemployment risk at the micro level and uncertainty over wage growth at the macro level. The critical assumption I make is that households' employment transition probabilities are a function of aggregate economic activity.

Setting. There are two periods, $t = 0, 1$. Aggregate risk is represented by the normal random variable $\sigma\epsilon \sim \sigma\mathcal{N}(0, 1)$, which is realized at the beginning of period 1. While I leave the macro block of the model largely unspecified, I assume and work with a notion of aggregate economic activity which I denote by Y_t. The only structure I have to impose is that economic activity in period 1 responds to the realization of aggregate risk, that is $Y_1 = Y_1(\sigma\epsilon)$.[19]

Households. A continuum of households i make consumption and savings decisions, facing

(2020), Bayer et al. (2019), Ottonello and Winberry (2017), Acharya and Dogra (2020), and Bilbiie (2020). This list is non-exhaustive.

[17] There is a large literature that studies the zero lower bound constraint quantitatively in representative-agent settings. See for example Christiano et al. (2011), Fernández-Villaverde et al. (2015), Nakata (2017) and Plante et al. (2018).

[18] For example, see Den Haan (1996), Den Haan and Others (1997), Krusell and Smith (1998), Reiter (2010), Algan et al. (2008), Algan et al. (2014), Brunnermeier and Sannikov (2014), Brumm and Scheidegger (2017), Duarte (2018), Pröhl (2019), and Fernández-Villaverde et al. (2019).

[19] More formally, the discussion in this section is valid as long as the general equilibrium block satisfies the following restriction. Let X_t denote the vector of all macroeconomic aggregates in period t. Then there are sets of equations, which I dub the macro block, $\mathcal{H}_0(X_0, \mathbb{E}_0(X_1)) = 0$ and $\mathcal{H}_1(X_0, X_1, \sigma\epsilon) = 0$. In particular, the assumption that only the first moment of future aggregate states of the economy affects the allocation in period 0 guarantees that macro uncertainty has no effect on household behavior to first order. In this setting, we have $Y_1 = Y_1(X_0, \sigma\epsilon)$. In response to an increase in σ, there is a direct (partial equilibrium) effect on Y_1 through its second argument and an indirect (general equilibrium) effect through its first argument.

uncertainty at both the micro and macro level. Household i's budget constraints are given by

$$c_{i,0} + a_{i,1} = a_{i,0} + y_{i,0} z_{i,0}$$
$$c_{i,1} = R a_{i,1} + y_{i,1} z_{i,1}.$$

In period 0, households consume, $c_{i,0}$, and save, $a_{i,1}$, out of initial wealth, $a_{i,0}$, and total labor income given by $y_{i,0} z_{i,0}$. In period 1, household consumption is equal to the gross return on savings and labor income. Labor supply is inelastic.

The key object in this stylized setting is household labor income, which comprises a component that loads on the aggregate state, $y_{i,t}$, and a purely idiosyncratic term, $z_{i,t}$. The idiosyncratic component corresponds to the household's employment status, with $z_{i,t} \in \{0, 1\}$. Conditional on employment, $z_{i,t} = 1$, $y_{i,t}$ can be thought of as a wage. In particular, I assume that a household's wage is directly proportional to economic activity, $y_{i,t} = \gamma_i Y_t$, as in Werning (2015).

The main assumption I make, which introduces a meaningful interaction between micro and macro uncertainty, is that the unemployment risk faced by households is counter-cyclical. In particular, I assume that the probability that household i becomes or remains employed in period 1 directly depends on economic activity, with

$$\mathbb{P}(z_{i,1} = 1 \mid z_{i,0}, \epsilon) = p_i(Y_1),$$

where p_i corresponds to the job finding rate when $z_{i,0} = 0$ and one minus the job seperation rate when $z_{i,0} = 1$.

Household preferences are defined over consumption, given by

$$\mathbb{E}_0 \sum_{t=0}^{1} \beta^t U(c_{i,t}).$$

The household's consumption and savings decision is then characterized by a standard Euler equation. Namely,

$$U'(c_{i,0}) = \beta R \mathbb{E}_0 \Big[U'(c^u_{i,1}) \Big(1 - p_i(Y_1)\Big) \Big] + \beta R \mathbb{E}_0 \Big[U'(c^e_{i,1}) p_i(Y_1) \Big]$$

where I have used the law of total probability to split expected marginal utility in period 1 into an unemployment and an employment term. $U'(c^u_{i,1})$ is the marginal utility of consumption conditional on unemployment, and it is multiplied by the probability of unemployment given economic activity Y_1. Similarly, $U'(c^e_{i,1})$ is the marginal utility of consumption conditional on employment, and it is multiplied by the probability of employment.

To think about the effects of uncertainty on household behavior, I consider a compariative static in σ, which is analogous to a macro uncertainty shock in this illustrative model. The following result characterizes how household consumption depends on σ locally around $\sigma \approx 0$. That is, I perform a Taylor expansion around an economy that features micro uncertainty in the form of employment transitions but no aggregate risk.

Proposition 1. *To second order,*

$$c_{i,0}(\sigma) \approx c_{i,0}(0) + \frac{1}{2}\frac{d^2 c_{i,0}}{d\sigma^2}(0)\,\sigma^2,$$

where

$$\frac{d^2 c_{i,0}}{d\sigma^2} = \Bigg\{ \overbrace{\frac{U'''(c^e_{i,1})\gamma_i^2}{U''(c_{i,0})} p_i\,\mathbb{E}_0\left[\left(\frac{\partial Y_1}{\partial \sigma}\right)^2\right]}^{\text{① Pure macro (wage) uncertainty}} + \overbrace{\frac{U'(c^e_{i,1}) - U'(c^u_{i,1})}{U''(c_{i,0})}}^{\text{Micro uncertainty}} \Bigg(\overbrace{p_i''\,\mathbb{E}_0\left[\left(\frac{\partial Y_1}{\partial \sigma}\right)^2\right]}^{\text{② Direct (PE) effect}} + \overbrace{p_i'\,\mathbb{E}_0\left[\left(\frac{d^2 Y_1}{d\sigma^2}\right)\right]}^{\text{③ Indirect (GE) effect}} \Bigg)$$

$$+ \underbrace{2\frac{U''(c^e_{i,1})\gamma_i}{U''(c_{i,0})} p_i'\,\mathbb{E}_0\left[\left(\frac{\partial Y_1}{\partial \sigma}\right)^2\right]}_{\text{④ Cov micro} \times \text{macro risk}} \Bigg\} \times \beta R\, MPS_{i,0} \quad + \underbrace{\Delta_i}_{\substack{\text{"Standard" effects} \\ \text{(no interaction with uncertainty)}}}$$

All objects are evaluated in the limit as $\sigma \to 0$, and $MPS_{i,0}$ is household i's marginal propensity to save in period 0.

Macro uncertainty affects household consumption in this model through four channels:

1. As in any representative-agent model, aggregate risk elicits the standard precautionary savings motive that is proportional to $U''' > 0$ (Kimball (1990)), channel ①. Only employed households are exposed to wage risk, so this term is scaled by the probability of employment, p_i. Their exposure is proportional to the sensitivity of wages to the aggregate shock ϵ, which, to second order, is given by $(\gamma_i Y_1')^2$.

The remaining three channels characterize the effect of macro uncertainty on household consumption *through* its interaction with micro uncertainty.

Aggregate wage risk exposes households to marginal changes in consumption, whose relevance is captured by differentiating marginal utility. Micro uncertainty in the form of employment transitions, on the other hand, represents an *idiosyncratic rare disaster* for households: it leads to discrete jumps in consumption and, therefore, marginal utility. Its relevance is not captured by a derivative of marginal utility but, rather, by a *jump* scaling factor. The effect on marginal utility of a transition from unemployment to employment is thus given by $U'(c^e_{i,1}) - U'(c^u_{i,1})$. It is in precisely this sense that the prospect of job loss is akin to a rare disaster, and consequently much more prominent, from the perspective of households.

2. Both channels ② and ③ represent the effect of macro uncertainty on micro uncertainty, and are thus scaled by $U'(c^e_{i,1}) - U'(c^u_{i,1}) < 0$. Channel ② corresponds to the *direct* or partial equilibrium effect.[20] This channel is operative only when $p_i'' \neq 0$, so that

[20] I use "direct" and "indirect" in the sense of Kaplan et al. (2018).

there is a non-linearity at the micro level in the response of employment transition rates to economic activity. I present some empirical evidence for such non-linearities in Section 4.2. Consider the case of an employed household. Intuitively, if the job *separation* rate is strictly convex as a function of economic activity, $p_i'' < 0$, then a mean-zero spread in aggregate risk leads to an increase in the household's expected probability of unemployment. Through this interaction with unemployment risk, macro uncertainty has a first-moment effect on expected earnings and induces desired savings.

3. Macro uncertainty has an effect on micro uncertainty to second order, even if the job finding and separation rates are only linear, $p_i' \neq 0$. Channel ③ represents this *indirect* or general equilibrium effect. An increase in uncertainty will, to second order, elicit precautionary savings (e.g. channel ①) and thus lead to a contraction in aggregate economic activity. If this contraction in activity is persistent and propagates into period 1, then households expect an associated GE effect on employment transition rates.[21] Simply put, if a macro uncertainty shock leads to a persistent recession, then households expect elevated unemployment risk going forward.[22] The indirect GE effect of channel ③ is the dominant transmission channel of macro uncertainty in the quantitative model: households respond to an increase in macro uncertainty precisely because it translates into micro uncertainty.

4. Finally, there is a covariance between risks at the micro and macro level. Formally, channel ④ represents the effect of σ, to second order, on the covariance between marginal utility and employment transition rates. When the probability of transitioning into employment state j is high for precisely those realizations of ϵ that also imply a high marginal utility of employment state j, then this channel represents another source of risk for households and induces precautionary savings. If the covariance structure is reversed, then this represents a hedging term.

The behavioral response in consumption is also subject to other effects, specifically a set of

[21] In the continuous-time quantitative model of Section 2, this GE effect will be *contemporaneous*. Therefore, an indepth discussion of the sources of persistence and propagation is unnecessary at this point.

[22] When the economy's general equilibrium block takes the form discussed in the previous footnote, then this GE effect can formally be further decomposed into two components. When $Y_1 = Y_1(X_0, \sigma\epsilon)$, we have

$$\mathbb{E}_0\left[\frac{d^2 Y_1}{d\sigma^2}\right] = \underbrace{\mathbb{E}_0\left[\frac{\partial Y_1}{\partial X_0}\right]\frac{d^2 X_0}{d\sigma^2}}_{\text{Endogenous GE effect}} + \underbrace{\mathbb{E}_0\left[\frac{\partial^2 Y_1}{\partial \sigma^2}\right]}_{\text{Macro non-linearity}} .$$

The first term highlights that, to second order, a macro uncertainty shock adversely affects aggregates in period 0, $d^2 X_0 / d\sigma^2$. This can lead to a persistent recession in period 1, $\mathbb{E}_0[\partial Y_1 / \partial X_0]$, and raise the expected probability of job loss. The second term emphasizes that, even holding the other macro aggregates X_0 fixed, there will be a similar effect if there is skewness at the aggregate level. When the economy is close to or in a crisis region, then a mean-zero spread in aggregate risk will have a first-moment impact on activity because, due to negative skewness, a negative shock is amplified more than a similarly sized positive shock. The large negative skewness exhibited by U.S. business cycles speaks to the relevance of this economic force.

interest rate and earnings effects, that I suppress here. These effects are "standard" in the sense that they would also emerge in a representative-agent setting and do not interact with uncertainty in meaningful ways.

2 A HANK Model with Micro and Macro Uncertainty

The quantitative model takes as its starting point the heterogeneous-household model proposed by Kaplan et al. (2018). Households face uninsurable earnings risk but can trade liquid and illiquid assets under incomplete markets. I depart from this benchmark in three important ways. First, I model aggregate uncertainty in the form of discount rate shocks. Second as in the illustrative model in Section 1, I assume that unemployment risk varies over the business cycle. Third, I introduce a zero lower bound (ZLB) constraint on monetary policy.

2.1 Households

The economy is populated by a continuum of households. Facing both idiosyncratic unemployment risk and aggregate uncertainty, households make consumption, savings and portfolio allocation decisions across time. The idiosyncratic state of a household consists of its portfolio position, made up of liquid assets a_t and illiquid assets k_t, and its employment status z_t.

Household preferences are defined over consumption and labor, given by

$$\max \mathbb{E}_0 \int_0^\infty e^{-\int_0^t (\rho_s+\zeta)\, ds} u(c_t, h_t) dt, \tag{1}$$

where c_t is the rate of consumption and h_t denotes the rate at which work hours are supplied. As in the canonical New Keynesian model, c_t is a consumption basket which is itself comprised of intermediate goods. This final good basket is priced at the consumer price index (CPI) P_t. Households die at rate ζ. At the same time, an equal mass of new households is formed with zero liquid and illiquid assets.

Discount rate shocks. The household's effective discount rate is given by $\rho_t + \zeta$. All households share a time-varying discount rate, ρ_t, which is the source of aggregate demand shocks in my model.[23] It follows a continuous-time AR(1) process, given by

$$d\rho_t = \theta_\rho(\bar{\rho} - \rho_t)dt + \sigma_\rho dB_t,$$

where B_t is a standard Brownian motion.[24] Discount rate shocks are a popular proxy for aggregate demand shocks.

[23] Basu and Bundick (2017), an important reference point for my quantitative results, also build a model with aggregate discount rate shocks. Using the same underlying shock process makes a direct comparison of results easier. Auclert et al. (2020) also use a discount rate shock as one of their demand shocks.

[24] I have solved versions of the model with alternative demand and supply shocks. In Appendix F, I present a variant of the model with supply (TFP) shocks and financial constraints.

Budget constraint. Households trade in two asset markets, one for bonds (liquid assets) and one for capital (illiquid assets).I denote the household's liquid asset position by a_t and its illiquid asset position by k_t. Households are endowed with an investment technology that transforms $Q_t \iota_t + P_t \psi(\iota_t, k_t)$ units of the numeraire into ι_t units of capital, where Q_t is the price of capital investment and P_t is the CPI. I denote the real price of capital by $q_t = Q_t / P_t$. ψ_t represents an investment adjustment cost and is the source of capital's illiquidity. Households are thus the direct owners of capital in this model, which they rent to firms in an economy-wide, competitive rental market.

Following Blanchard (1985), I introduce perfect annuity markets, in which households can trade claims on their remaining wealth at time of death. They pledge this wealth to a risk-neutral insurance company that, in turn, compensates households with a flow annuity payment at a rate ζ times their current asset positions. This is exactly the payment rate that makes the insurance company break even in expectation. See Appendix A.6 for details. Introducing household death rates is a commonly used technique to ensure stationarity in the wealth distribution.

A household's liquid asset position evolves according to

$$\dot{a}_t = (r_t + \zeta) a_t + k_t \frac{dR_t}{dt} + e_t - q_t \iota_t - \psi(\iota_t, k_t) - c_t. \tag{2}$$

This budget constraint is derived from its nominal analog (see Appendix A.1.1 for details). The real rate of return on the liquid asset consists of the real riskfree rate r_t and the rate of annuity payments ζ. Capital earns a real rate of return dR_t, which I will further specify in Section 2.6 after presenting the rest of the model. Wage payments and rebates are collected in the earnings variable e_t, given by

$$e_t = (1 - \tau^{\text{lab}}) z_t w_t h_t + \tau_t^{\text{lump}} + \tau^{\text{UI}}(z_t),$$

where $w_t = W_t / P_t$ is the real wage, τ_t^{lump} is a set of lump-sum rebates and $\tau^{\text{UI}}(z_t)$ denotes unemployment insurance payments that explicitly depend on the household's employment status. Finally, households consume and buy illiquid assets at relative price q_t, subject to the adjustment cost ψ.

The household's illiquid asset position evolves as

$$\dot{k}_t = (\zeta - \delta) k_t + \iota_t, \tag{3}$$

where δ denotes capital depreciation and ζ the rate of annuity payments. Since households are the direct owners of capital, they also incur depreciation. For convenience, I will later use the shorthand notation s_t and m_t to refer to the drift in the household's liquid and illiquid asset positions, respectively.

Finally, households' portfolio allocation is subject to a borrowing constraint on liquid assets, $a_t \geq \underline{a}$ with $0 > \underline{a}$, and a short-sale constraint on capital, $k_t \geq 0$.

The household's dynamic problem is therefore to maximize (1) subject to budget constraints (2) and (3), as well as the borrowing and short-sale constraints. Households take as given the laws of motion for their employment status and macroeconomic aggregates. I denote the resulting policy functions in terms of the household's idiosyncratic state variables as $c_t(a,k,z)$, $h_t(a,k,z)$ and $\iota_t(a,k,z)$.

Appendix A provides additional details on the household problem. I derive a recursive representation and the associated consumption Euler equation in Appendix A.1 and A.2. In Appendix A.4, I discuss the implications of inflation risk for the household's portfolio equations. In Appendix A.5, I discuss alternative assumptions for transfers and rebates.

Investment adjustment cost. I adopt the functional form for the household's investment adjustment cost that is used by Kaplan et al. (2018). In particular,

$$\psi(\iota_t, k_t) = \psi_0|\iota_t| + \psi_1 \left(\frac{|\iota_t|}{k_t} \right)^2 k_t.$$

Unemployment risk. Households face uninsurable earnings risk that is encoded in the state variable z_t, which follows a two-state Markov process. These two states are given by $z_t \in \{z^E, z^U\}$ and are thus meant to represent employment and unemployment. In practice, I set $z^E = 1$ and $z^U = 0$ as in the two-period model.

Formally, let N_t be a standard Poisson process and let $j \in \{E, U\}$ index the household's current earnings state. Then the evolution of the earnings state follows

$$dz_t^j = (z^{-j} - z^j)\, dN_t\left(\lambda_t^j\right),$$

where λ_t^j is the Poisson arrival rate. These transition rates are the continuous-time analog of the transition probabilities p_i in the two-period model. For now, I specify λ^E and λ^U as arbitrary functions on the aggregate state space, which the household takes as given. In Section 4.2, I estimate the sensitivity of employment transition rates to changes in economic activity in Current Population Survey (CPS) micro data. By using the resulting estimates directly in the model, I take as given the link between employment transitions and economic activity.[25] The sense in which my model generates an endogenous link between micro and macro uncertainty is through changes in economic activity: heightened macro uncertainty depresses aggregate demand which in turn raises unemployment risk.

2.2 Firms

The structure of the goods producing sector is as in the standard New Keynesian model. Monopolistic intermediate producers, which I will simply refer to as firms, sell differentiated varieties to a retailer that aggregates these into a composite final consumption good.

[25]These empirical estimates are intended as a reduced-form proxy for the state dependence in employment transition rates that would result, for example, from a search-and-matching micro-foundation. Explicitly implementing a search-and-matching block is beyond the scope of this paper and left for future work.

Retailer. The aggregation technology of the retailer is given by

$$Y_t = \left(\int_0^1 Y_t(j)^{\frac{\epsilon^f - 1}{\epsilon^f}} dj \right)^{\frac{\epsilon^f}{\epsilon^f - 1}},$$

where $Y_t(j)$ is the output produced by intermediate firm j and ϵ^f denotes the elasticity of substitution across intermediate varieties. The retailer's cost minimization problem gives rise to the standard demand function for intermediate inputs, $Y_t(j) = (P_t(j)/P_t)^{-\epsilon^f} Y_t$, where $P_t(j)$ is the price of firm j.

Firms. Intermediate goods producers operate a technology that combines capital and labor. Firm j's production function is given by

$$Y_t(j) = K_t(j)^{1-\beta} L_t(j)^{\beta}, \tag{4}$$

where β denotes the labor share.

Firm j demands work hours at rate $L_t(j)$, for which it pays the nominal wage rate W_t. There is an economy-wide rental market for capital in which firms rent capital from households at nominal rental rate i_t^k. All firms face the same prices and take them as given. Firm j's nominal profits are thus given by $\Pi_t(j) = P_t(j)Y_t(j) - W_t L_t(j) - i_t^k K_t(j)$. Standard cost minimization implies a composite nominal marginal cost of

$$MC_t = \frac{1}{\beta^\beta (1-\beta)^{1-\beta}} \left(i_t^k \right)^{1-\beta} \left(W_t \right)^{\beta},$$

and I define the real marginal cost as $mc_t = MC_t / P_t$, so that the real rental rate of capital is analagously given by $r_t^k = i_t^k / P_t$.

Dynamic price setting. While the choice of factor inputs is a static one, firms are monopolistically competitive and face a dynamic pricing problem subject to price adjustment costs. I adopt the quadratic specification of Rotemberg (1982).

Define $\pi_t(j)$ to be the instantaneous rate of inflation chosen by firm j. This rate is chosen to maximize an appropriately discounted sum of all future expected profits subject to an adjustment cost which I specify in terms of firm utility. This dynamic problem is given by

$$\max_{\{\pi_t(j)\}} \mathbb{E}_0 \int_0^\infty e^{-\int_0^t i_s^k ds} \Big[(1 - mc_t) P_t(j) Y_t(j) - \Lambda(\pi_t(j)) \Big] dt, \tag{5}$$

where the effective discount rate is in terms of the cost of capital i_t^k. See Appendix A for additional details on the firm pricing problem.

2.3 Capital producer

For expositional clarity, I explicitly specify a capital producing sector.[26] The capital producer is owned by households. It is endowed with a technology that transforms the final consumption

[26] Capital production could be subsumed in the household problem, as in Brunnermeier and Sannikov (2014) for example, as long as one is careful to distinguish between the idiosyncratic and aggregate adjustment costs.

good into capital. Concretely, it requires $I_t + \Phi(I_t/K_t)K_t$ units of the final good to produce I_t units of capital. The capital producer sells all new capital to households at the nominal price $Q_t = P_t q_t$.

The real rate of profit earned by the capital producer, and paid out to households, is thus given by

$$\Pi_t^Q = q_t I_t - I_t - \Phi\left(\frac{I_t}{K_t}\right)K_t, \tag{6}$$

where I_t should be interpreted as the rate at which new capital, i.e. the latest vintage, is generated. The associated optimality condition for static profit maximization is then given by

$$q_t = 1 + \Phi'\left(\frac{I_t}{K_t}\right). \tag{7}$$

Finally, accounting for the depreciation of capital incurred by households, the evolution of the economy's aggregate capital stock is given by

$$\dot{K}_t = I_t - \delta K_t. \tag{8}$$

2.4 Nominal wage stickiness

I follow a long tradition in the wage rigidity literature and model household labor supply h_t as determined by union labor demand (see Erceg et al. (2000), Schmitt-Grohé and Uribe (2005), and Auclert et al. (2020)). There is a continuum of labor unions indexed by $k \in [0, 1]$. Households supply labor to each of these unions, which in turn bundle labor and pass on a differentiated labor variety to an aggregate labor packer. This structure is meant to mimic the analogous structure in the standard price stickiness setup where intermediate producers sell differentiated varieties to a final retailer. The final labor packer aggregates all labor in the economy into a composite labor factor, which is then used by firms in the production of intermediate goods.

Since each household supplies labor to each of the k unions, we have $h_t = \int h_{k,t} dk$. Union k's aggregation technology combines all households' effective work hours, $z_t h_{k,t}$, into a union-specific labor variety $L_{k,t}$. An aggregate, competitive labor packer combines these intermediate inputs into an aggregate labor supply basket according to

$$L_t = \left(\int L_{k,t}^{\frac{\epsilon^w - 1}{\epsilon^w}} dk\right)^{\frac{\epsilon^w}{\epsilon^w - 1}}.$$

It sells this composite good at the nominal wage rate W_t to intermediate goods producers. The demand function of the labor packer for union k's input is given by $L_{k,t} = (W_{k,t}/W_t)^{-\epsilon^w} L_t$.

I assume that a union faces a quadratic utility cost when adjusting its nominal wage $W_{k,t}$. This cost is given by $\frac{\chi^w}{2}(\pi_{k,t}^w)^2 L_t$, where $\pi_{k,t}^w = \dot{W}_{k,t}/W_{k,t}$ is the rate of nominal wage k inflation. Labor unions act according to an equal-weighted sum of household preferences.

Unions are small and take as given households' policy functions. Their dynamic wage setting problem is thus given by

$$\max_{\pi^w_{k,t}} \mathbb{E}_0 \int_0^\infty e^{-\int_0^t (\rho_s + \zeta)\, ds} \left[\int u(c_t, h_t) g_t d(a,k,z) - \frac{\chi^w}{2} (\pi^w_{k,t})^2 L_t \right] dt. \tag{9}$$

Importantly, union k directly controls the evolution of its wage path, so that there are no stochastic innovations and nominal wages, $W_{k,t}$, are locally deterministic. Furthermore, unions are small and take as given the consumption policy function of the household, the demand function of the labor packer, as well as the distribution of households and its evolution. In the zero-inflation steady state, the real wage and marginal cost will be pinned down by $(1-\tau^{\text{lab}}) w \int u'(c) g = \frac{\epsilon^w}{\epsilon^w - 1} v'(H)$. See Appendix A for additional details on the union problem.

2.5 Government policy and ZLB

In the baseline model, the scope of fiscal policy is limited. The government budget constraint is simply given by

$$r_t B^G + \tau_t^{\text{lump}} + \int \tau^{\text{UI}}(z) g_t(a,k,z) d(a,k,z) + G_t = \Pi_t + \tau^{\text{lab}} w_t L_t. \tag{10}$$

Fiscal expenditures include lump-sum rebates and unemployment insurance payments to households. The government is assumed to be a net debtor, with a constant level of riskfree debt outstanding given by B^G, on which it pays the real interest rate r_t. Finally, the government purchases the final consumption good at rate G_t. The government finances these expenditures with a labor income tax and with the corporate profits collected from the goods producing sector, Π_t.

Monetary policy in this model follows a simple Taylor rule and is subject to the zero lower bound (ZLB) constraint on nominal interest rates, so that

$$i_t = \max\left\{ r^* + \bar{\pi} + \lambda_\pi \pi_t + \lambda_Y y_t,\ 0 \right\}, \tag{11}$$

where r^* is the steady state riskfree rate, and $y_t = \log(Y_t / Y^*)$ denotes the output gap.

2.6 Aggregation and market clearing

I denote by $g_t(a,k,z)$ the cross-sectional household distribution at time t over liquid assets, a, illiquid assets, k, and employment status, z. Aggregation in this economy, for example of consumer expenditures, then takes the form

$$C_t = \int c_t(a,k,z) g_t(a,k,z) d(a,k,z).$$

The aggregate stocks of liquid and illiquid assets held by households are defined, respectively, as $A_t = \int a g_t(a,k,z) d(a,k,z)$ and $K_t = \int k g_t(a,k,z) d(a,k,z)$.

Market clearing. Formally, there are five markets in this economy that must clear at all times. The clearance of two of these markets, however, the rental market for capital and the labor market, is already implicit in the notation I have adopted: The aggregate stock of capital used in production must equal the aggregate stock of the illiquid asset held by households, and aggregate hours must equal the aggregate labor supply bundled by unions and the labor packer.

This leaves the bond market, which clears when the stock of liquid assets held by households equals the total riskfree debt issued by the government,

$$A_t = B^G.$$

In particular, as long as government debt is constant the stock of liquid assets is also constant. Second, the market for the final consumption good clears when

$$Y_t = C_t + I_t + \Phi_t + \Psi_t + G_t,$$

where C_t is aggregate consumption, $I_t + \Phi_t$ denotes gross investment expenditures including adjustment costs, and $\Psi_t = \int \psi(\iota_t, k) g_t(a,k,z) d(a,k,z)$ is the aggregate adjustment cost paid by households to trade in the illiquid asset.

This leaves, finally, the market for new capital, where households trade with the capital producing firm. All new capital produced must be purchased by households, and the price of capital Q_t adjusts to clear this market,

$$I_t = \int \iota_t(a,k,z) g_t(a,k,z) d(a,k,z).$$

In Appendix A.7, I provide an illustrative derivation of Walras' law.

Profits and the return on capital. What remains to be specified is the dissemination of profits and the composition of the rate of return on capital, dR_t. In the baseline model, I assume that pure profits earned by the goods producing sector are paid out uniformly to all households, via the government's lump-sum rebate τ_t^{lump}. Pure profits earned by the capital producing sector, on the other hand, are disseminated in proportion to a household's illiquid asset holdings. That is, they constitute part of the return on capital.

Therefore, the real rate of return on capital is given by

$$dR_t = \left(r_t^k + \frac{\Pi_t^Q}{K_t} \right) dt \tag{12}$$

where r_t^k is the real rental rate paid by firms and Π_t^Q / K_t is the profit from capital production per unit of aggregate capital. The rate of return on the *stock* of capital features no capital gains term. If the household's portfolio equations are rewritten in terms of *net worth*, as is more common in the context of asset pricing, a capital gains term emerges that loads on the aggregate risk factor (see Appendix A.1.2 for details).

3 A Global Solution Method

Why are models like the one presented in Section 2 difficult to solve numerically? In any rational expectations equilibrium, agents must base their actions today on forecasts of future prices that are consistent with the economy's true law of motion. And in any heterogeneous-agent model, future prices will, through market clearing, depend on the future cross-sectional distribution of agents. Therefore, agents must consistently forecast the evolution of this distribution, which becomes part of the aggregate state space. The state space of the economy presented in Section 2 is therefore infinite-dimensional. Before moving on to my solution method, I will develop this argument formally and present the definition of a recursive equilibrium for my economy. The discussion in Section 3 is focused exclusively on the paper's methodological contribution. Readers who wish to proceed directly to the economic results may skip to Section 4 without loss.

The aggregate state space of my model is given by $\Gamma_t = (\rho_t, g_t)$, where ρ_t is the exogenous discount rate shock and g_t is the cross-sectional household distribution. The stationary value function of a household can be written as $V_t(a,k,z) = V(a,k,z,\Gamma_t)$, with slight abuse of notation. The household value function thus takes an infinite-dimensional, or measure-valued, input. While a, k and Γ are continuous state variables, z is discrete. I will therefore denote by $V^j(a,k,\Gamma) = V(a,k,z_j,\Gamma)$ the value function of a household with employment status j and by $g^j(a,k,\rho) = g(a,k,z_j,\rho)$ the mass of households of employment type j with portfolio (a,k) in aggregate state Γ. Since the household distribution g is an argument of the value function, I will first characterize its evolution over time. For convenience, I introduce the shorthand notation $s^j(a,k,\Gamma)$ and $m^j(a,k,\Gamma)$ for the drift in liquid and illiquid portfolio positions, respectively.

The household distribution $g_t(a,k,z)$ evolves through time according to a Kolmogorov forward equation given by

$$\begin{aligned}\frac{dg^j(a,k,\rho)}{dt} =& -\partial_a\Big[s^j(a,k,\Gamma)g^j(a,k,\rho)\Big] - \partial_k\Big[m^j(a,k,\Gamma)g^j(a,k,\rho)\Big] \\ &- \lambda^j(\Gamma)g^j(a,k,\rho) + \lambda^{-j}(\Gamma)g^{-j}(a,k,\rho),\end{aligned} \tag{13}$$

where ∂_x denotes the partial derivative operator with respect to state variable x, and $-j$ denotes the employment type that *is not* j. See Appendices B.1 and B.2 for details. For compacter notation, equation (13) can also be written $(\mathcal{A}^* g^j)(a,k,\rho) \equiv \frac{d}{dt} g^j(a,k,\rho)$.[27]

[27] $\mathcal{A}^*$ is a functional operator that I will discuss further below.

Formally, V^j solves a system of Hamilton-Jacobi-Bellman (HJB) equations given by

$$
\begin{aligned}
(\rho+\zeta)V^j(a,k,\Gamma) = \max_{c^j,h^j,\iota^j} & \left\{ u(c^j,h^j) + s^j\partial_a V^j(a,k,\Gamma) + m^j \partial_k V^j(a,k,\Gamma) \right\} \\
& + \lambda^j(\Gamma)\left[V^{-j}(a,k,\Gamma) - V^j(a,k,\Gamma)\right] + \theta_\rho(\bar{\rho}-\rho)\partial_\rho V^j(a,k,\Gamma) \\
& + \frac{\sigma_\rho^2}{2}\partial_{\rho\rho}V^j(a,k,\Gamma) \; + \; \underbrace{\sum_l \int \frac{\delta V^l}{\delta g}(\mathcal{A}^* g^l)(a,k,\rho) d(a,k)}_{\substack{\text{Effect of changes in cross-sectional distribution} \\ \text{on household value function}}}
\end{aligned}
\tag{14}
$$

subject to the household budget, borrowing and short-sale constraints. Recall that s^j and m^j are only used as shorthand notation for the drift in the household's liquid and illiquid asset positions, respectively. See Appendix A.2 for a formal derivation. Equation (14) is a characterization of the *stationary* household value function and therefore internalizes the effects of changes in the aggregate state of the economy. The last term in the third row captures the effect changes in the household cross-sectional distribution have on this value function. In particular, since g is infinite-dimensional, a functional Gateaux derivative δ is used.[28] This is precisely the term that cannot directly be implemented in any numerical scheme to solve this value function and must therefore be approximated.[29]

We are now ready to define a recursive equilibrium of this economy. The definition of recursive equilibrium formally expresses the intuitive account with which I started this section: Since households must consistently forecast future aggregates, the infinite-dimensional

[28] Formally, the value function $V(a,k,z,\rho,g)$ is defined on the product space comprised by the finite-dimensional state space for (a,k,z,ρ) and the space of measures for g. The partial differential equation (14) that characterizes V on this product space is also known as the "master equation" (see for example Cardaliaguet et al. (2015) and, originally, Lasry and Lions (2007) and Lions (2011)). The solution method I propose in this paper can therefore be viewed as a numerical implementation of the master equation. To my knowledge, this paper presents one of the first global numerical algorithms to solve the master equation of a mean field game with common noise.

[29] To build intuition, imagine that the cross-sectional distribution was a vector $g = (g_1, \ldots, g_N) \in \mathbb{R}^N$ rather than an infinite-dimensional function. This illustration is similar to one presented in Ahn et al. (2017). The HJB equation then becomes

$$
\begin{aligned}
(\rho+\zeta)V^j(a,k,g,\rho) = \max_{c^j,h^j,\iota^j} & \left\{ u(c^j,h^j) + s^j\partial_a V^j(a,k,g,\rho) + m^j \partial_k V^j(a,k,g,\rho) \right\} \\
& + \lambda^j(g,\rho)\left[V^{-j}(a,k,g,\rho) - V^j(a,k,g,\rho)\right] + \theta_\rho(\bar{\rho}-\rho)\partial_\rho V^j(a,k,g,\rho) \\
& + \frac{\sigma_\rho^2}{2}\partial_{\rho\rho}V^j(a,k,g,\rho) + \sum_{i=1}^N \sum_l \frac{\partial V^l(a,k,g,\rho)}{\partial g_i^l} \dot{g}_i^l.
\end{aligned}
$$

The problematic term becomes a sum of partial derivatives, which can in principle be implemented numerically since everything is now finite-dimensional. One contribution of my solution method is to generalize this approach to any arbitrary finite-dimensional approximation of g.

cross-sectional distribution becomes part of their state space.[30]

Definition. (*Recursive Equilibrium*) A recursive competitive equilibrium of this economy is defined as the sets of functions $\{V^j, g^j, c^j, h^j, \iota^j\}(a, k, \Gamma)$ for $j \in \{E, U\}$ and $\{r, r^k, q, \tau^{\text{lump}}, H, w, Y, K, C, I, L, \pi, \pi^w\}(\Gamma)$ such that:

(i) (Household optimization) $\{V^j, c^j, h^j, \iota^j\}$ solve the HJB (14) given all aggregates.

(ii) (Firm and union optimization) Given aggregates, $\pi(\Gamma)$ solves the firm problem (5) and $\pi^w(\Gamma)$ solves the union problem (9) in each aggregate state Γ.

(iii) (Aggregation) For each aggregate state Γ, the description of the aggregate household sector, i.e. $C(\Gamma)$, $I(\Gamma)$ and $H(\Gamma)$, is consistent with aggregation from the micro level, using the cross-sectional distribution $g^j(a, k, \Gamma)$.

(iv) (General equilibrium) All markets clear, and the remaining macroeconomic aggregates solve the model's general equilibrium conditions.

(v) (Distribution and rational expectations) The cross-sectional household distribution evolves according to (13) and household behavior is based on forecasts that are consistent with the true law of motion of the aggregate state of the economy.

3.1 Finite-dimensional distribution representations

Any numerical solution of a heterogeneous-agent model with aggregate risk must approximate the aggregate state space in a finite-dimensional subspace. In other words, any numerical solution method will implement an approximate value function that takes as an input a finite-dimensional approximation of the distribution g_t. In this paper, I focus on the class of finite-dimensional representations given by

$$F(\alpha_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x}), \tag{15}$$

where $\hat{g}_t(\boldsymbol{x}) = F(\alpha_t)(\boldsymbol{x})$ is then interpreted as the approximate cross-sectional distribution parameterized by α_t. For compacter notation, I use $\boldsymbol{x}_t = (a_t, k_t, z_t)$ to denote the vector of idiosyncratic state variables. While the representation (15) is quite general, it is most illustrative to think of F as a set of basis functions defined over $\boldsymbol{x}$ that are parameterized by α_t. Before proceeding, I present an example for illustration.[31]

[30] One could also define a sequence equilibrium for this economy. For example, Brunnermeier and Sannikov (2014) define a sequence equilibrium for an economy that is structurally similar to a variant of mine in which the distribution is a degenerate two-point delta function.

[31] In the context of a local solution, Winberry (2020) uses an approximate distribution representation that also takes the form (15). In this sense, my paper can be viewed as taking an approach similar to Winberry (2020) but in the context of a global solution method. I present several analytical and numerical tools below that help overcome the curse of dimensionality that oftentimes becomes a hurdle for global solution methods.

Example 2. Denote by $T^n(\boldsymbol{x})$ the nth Chebyshev polynomial over $\boldsymbol{x}$. We can stack these basis functions in a row-vector denoted $T(\boldsymbol{x})$. Letting $\alpha_t \in \mathbb{R}^N$ denote the column-vector of coefficients, equation (15) becomes

$$F(\alpha_t)(\boldsymbol{x}) = T(\boldsymbol{x})\alpha_t = \sum_{n=1}^{N} \alpha_t^n T^n(\boldsymbol{x}) \approx g_t(\boldsymbol{x}).$$

While the set of selected basis functions F remains constant over time, time variation in the parameters α_t allows to capture approximately the evolution of the true cross-sectional distribution g_t over time. In the context of the above example: The researcher picks the N Chebyshev polynomials $T^n(\boldsymbol{x})$ ex ante and they don't change over the course of a simulation. The basis function coefficients α_t^n, on the other hand, are time-varying. In particular, we will want to specify the time variation in α_t^n so as to match the time variation in $g_t(\boldsymbol{x})$ as closely as possible.

Approximate economy. How does the approximation (15) make the model of Section 2 tractable? I will refer to the model solved under the finite-dimensional representation (15) as an *approximate economy*. That is, an approximate economy uses $\hat{g}_t(\boldsymbol{x}) = F(\alpha_t)(\boldsymbol{x})$ as its cross-sectional distribution.

Recall the source of intractability in the household's value function equation (14): As part of the aggregate state space, the infinite-dimensional distribution g_t enters V as an argument. The aggregate state space of the approximate economy no longer features the true cross-sectional distribution but rather its finite-dimensional approximation, $\hat{g}_t(\boldsymbol{x})$. That is, the *approximate aggregate state space* becomes $\hat{\Gamma}_t = (\boldsymbol{X}_t, \hat{g}_t)$. Under the representation $\hat{g}_t(\boldsymbol{x}) = F(\alpha_t)(\boldsymbol{x})$, we can further simplify the aggregate state space: All information we require to track the evolution of $\hat{g}_t(\boldsymbol{x})$ over time is equivalently encoded in α_t. Therefore, the approximate aggregate state space becomes $\hat{\Gamma}_t = (\boldsymbol{X}_t, \alpha_t)$.[32]

Reducing the aggregate state space from Γ_t to $\hat{\Gamma}_t$ makes the household problem (14) tractable. The value function that characterizes household behavior in the approximate economy is now given by

$$\begin{aligned} V(\boldsymbol{x}, \rho, g) &\approx V(\boldsymbol{x}, \rho, \hat{g}) \\ &= \hat{V}(\boldsymbol{x}, \rho, \alpha). \end{aligned}$$

[32] My method also builds on previous work that makes use of *distribution selection functions* (e.g. Algan et al. (2008) and Reiter (2010)). Like the original Krusell-Smith algorithm, these papers use moments to characterize the distribution. In the spirit of my approach, however, they postulate a mapping like (15) that is consistent with their moments. This distribution selection function associates a cross-sectional distribution with each possible realization of moments. These algorithms share several important features with the method I propose here, and I show in the Numerical Appendix that several of my results below can also be useful for those algorithms. For example, I show how to modify the Algan et al. (2008) algorithm so that it can directly make use of Proposition 3.

In particular, the household value function now only takes finite-dimensional arguments. In the approximate economy households only have to forecast the approximate distribution $\hat{g}_t(\boldsymbol{x})$ or, equivalently, the new state variables α_t. I show below that, under quite general conditions, equilibrium in the approximate economy is characterized by the law of motion

$$d\alpha = \mu_\alpha(\hat{\Gamma})dt + \sigma_\alpha(\hat{\Gamma})dB.$$

That is, α_t follows an Ito diffusion process, and μ_α and σ_α are equilibrium objects that we must solve for. In an approximate economy based on (15), therefore, household behavior is now characterized by

$$\begin{aligned}(\rho+\zeta)\hat{V}^j(a,k,\hat{\Gamma}) = \max_{\hat{c}^j,\hat{h}^j,\hat{\iota}^j} &\left\{ u(\hat{c}^j,\hat{h}^j) + \hat{s}^j \partial_a \hat{V}^j(a,k,\hat{\Gamma}) + \hat{m}^j \partial_k \hat{V}^j(a,k,\hat{\Gamma}) \right\} \\ &+ \hat{\lambda}^j(\hat{\Gamma})\Big[\hat{V}^{-j}(a,k,\hat{\Gamma}) - \hat{V}^j(a,k,\hat{\Gamma})\Big] + \theta_\rho(\bar{\rho}-\rho)\partial_\rho \hat{V}^j(a,k,\hat{\Gamma}) \\ &+ \frac{\sigma_\rho^2}{2}\partial_{\rho\rho}\hat{V}^j(a,k,\hat{\Gamma}) + \underbrace{\mu_\alpha(\hat{\Gamma})\partial_\alpha \hat{V}^j(a,k,\hat{\Gamma}) + \frac{1}{2}\sigma_\alpha(\hat{\Gamma})^T \partial_{\alpha\alpha}\hat{V}^j(a,k,\hat{\Gamma})\sigma_\alpha(\hat{\Gamma})}_{\substack{\text{Effect of cross-sectional distribution on household} \\ \text{value function in approximate economy}}}.\end{aligned} \tag{16}$$

The new terms in the third row are now readily computable given the functions $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$.

Working with a finite-dimensional distribution representation of the form (15) offers numerous advantages.[33] While this approach is commonly used in the context of local perturbation methods, it has proven intractable in the context of global methods due to the curse of dimensionality. That is, accurate approximations $\hat{g}_t(\boldsymbol{x}) = F(\alpha_t)(\boldsymbol{x})$ oftentimes require high-dimensional α_t. In the remainder of this section, I discuss several analytical and numerical tools that help overcome this challenge. In the interest of brevity and accessibility, I only sketch technical arguments in the main text. Details are provided in Appendix B and a separate Numerical Appendix.[34]

[33] A key advantage of being able to work with an approximate distribution object $\hat{g}_t(\boldsymbol{x})$ on the grid is that it allows us to evaluate market clearing conditions. If we instead approximates the cross-sectional distribution with moments, this is generally not possible. For example, to evaluate the goods market clearing condition in Section 2 requires first evaluating aggregate household consumption, $\hat{C}(\Gamma) = \int \hat{c}(\boldsymbol{x},\Gamma)\hat{g}(\boldsymbol{x},\Gamma)d(\boldsymbol{x})$, which is only possible given $\hat{g}$. With a distribution object "on the grid", we are able to solve for prices and other general equilibrium objects directly on the grid. By contrast, algorithms like Krusell and Smith (1998) must resort to a costly "simulate-estimate" approach, which I describe below, to solve for these general equilibrium objects.

[34] An important question is under what conditions the scheme in (15) is a *consistent* approximation to the system of coupled PDEs, given by the HJB and KF equations, that characterizes the true model. In other words, under what conditions does a sequence of approximating economies characterized by $\{F^n(\alpha_t^n)\}_n$ converge to the true economy as $n \to \infty$. I am currently working on proving such a convergence result for particular classes of basis functions F.

3.2 How to overcome the curse of dimensionality?

For most economic applications of interest, accurate approximations of the distribution will require high-dimensional $F(\alpha_t)(x)$. Global methods will therefore quickly encounter the curse of dimensionality. In Schaab and Zhang (2020), we have developed an adaptive sparse grid library for solving partial differential equations in continuous time. Sparse grid methods aim to combat the curse of dimensionality when representing functions on high-dimensional grids. Regular sparse grids reduce the complexity of a grid in d dimensions from $\mathcal{O}(n^d)$, where n denotes the number of grid points per dimension, to $\mathcal{O}(n \log(n)^{d-1})$ (see for example Bungartz and Griebel (2004)). They have found occasional use in economics (see for example Krueger and Kubler (2004) and Judd et al. (2014)). Truly adaptive sparse grids were first introduced to economics by Brumm and Scheidegger (2017) but had previously enjoyed more popularity in physics and applied math.

Grid adaptation leverages the insight that not all grid points are "equally valuable" when representing a function on a grid. A rule of thumb is that an accurate function representation requires more grid points in areas where the function is particularly concave. A linear function in one dimension, for example, can be represented perfectly using only two grid points and linear interpolation elsewhere. Adaptive sparse grid algorithms automatically add and drop grid points to maximize the efficiency of grid point placement for a given application.

The main contribution of Schaab and Zhang (2020) relative to Brumm and Scheidegger (2017) is to develop a robust adaptive sparse grid infrastructure for solving (partial) differential equations in continuous time. While a self-contained description of our library is beyond the scope of this paper, I will illustrate the power of adaptive sparse grids in the context of the model of Section 2. For simplicity, I focus on the model's deterministic steady state, formally defined as $V^{j,0}(\cdot) = \lim_{\sigma_\rho \to 0} V^j(\,\cdot\,; \sigma_\rho)$. Figure 1 displays fully converged solutions of $V^{U,j}(a,k)$, an unemployed household's value function in the deterministic steady state, across four increasingly adapted grids.[35] A dense grid consists of a full set of equidistantly spaced grid points. The key insight of sparse grid methods is that not all of these grid points are equally valuable in the representation of V. In Panel (b), the value function is solved on a so-called regular sparse grid, which starts from a dense grid and removes grid points according to a prespecified pattern. Panels (c) and (d), finally, adapt the grid and place grid points efficiently to conform to the concavity in V. Panel (d) illustrates clearly that grid points are placed in the region where a and k are low and household behavior at the micro level exhibits non-linearities. While Figure 1 demonstrates that sparse grid methods are already useful for two-dimensional grids, they become especially powerful in higher dimensions.

[35]For illustration, Figure 1 uses more grid points in the idiosyncratic household dimensions than I use to solve the quantitative model in higher dimensions when $\sigma_\rho > 0$.

Figure 1: Household value function on increasingly adapted sparse grids

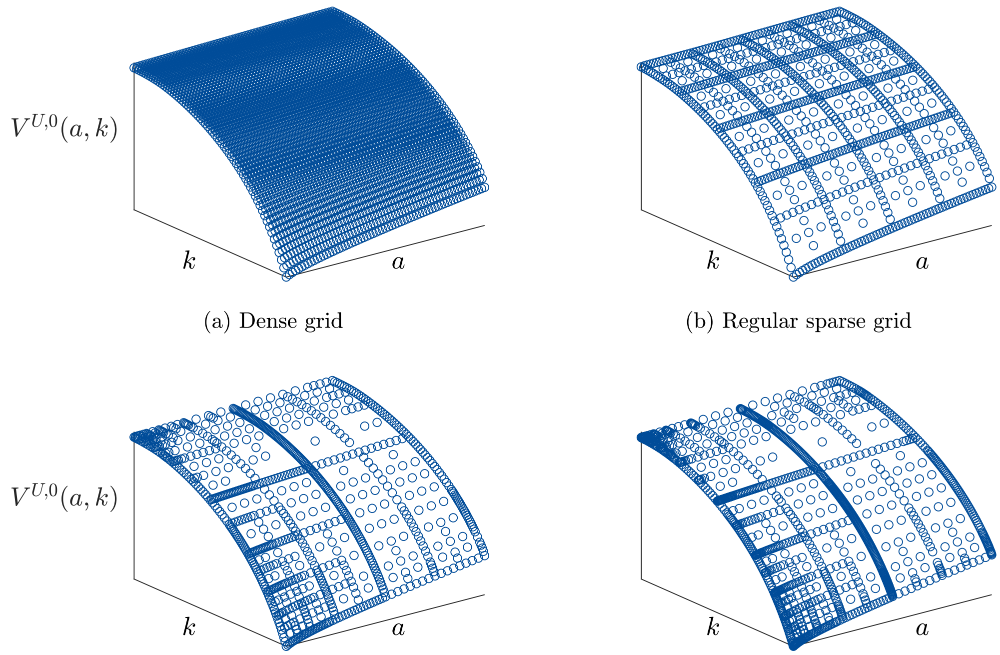

(a) Dense grid

(b) Regular sparse grid

(c) Rounds of adaptation: 3

(d) Rounds of adaptation: 5

Notes. Each panel displays the value function of an unemployed household in the deterministic steady state, i.e. $V^{U,0}(a,k) = \lim_{\sigma_\rho \to 0} V^U(a,k\,;\sigma_\rho)$. Panel (a) solves V on a dense grid consisting of a full set of equidistantly spaced grid points. Panel (b) solves V on a regular sparse grid which removes grid points from the dense grid according to a prespecified pattern. Panels (c) and (d) display the converged value function on increasingly adapted sparse grids.

3.3 How to find a consistent law of motion for α_t?

To solve the household problem in the approximate economy, we must solve for the equilibrium law of motion of α_t. This is highlighted by equation (16) which depends directly on μ_α and σ_α.

The traditional approach to solving for the law of motion of α_t is as follows: (i) Start with a guess for $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$. (ii) Given this guess, solve equation (16) and the rest of the model on the aggregate state space. (iii) Simulate the model. (iv) Update the original guess for $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$ using estimates from the simulated data, and start again at step (i). These steps are repeated until convergence. This strategy was popularized by Krusell and Smith (1998) and remains an integral part of most global solution methods currently in use.

While this "simulate-estimate" approach is, of course, valid and oftentimes useful in algorithms based on (15), it is also very costly.[36] In practice, it typically requires hundreds of

[36] In practice, I use a combination of the "simulate-estimate" approach and Proposition 3 to solve the quantitative model of Section 2. In the Numerical Appendix, I show that exclusively using Proposition 3 can lead to significant performance gains in simpler models such as the Krusell and Smith (1998) model or a

(outer) iterations in the fixed point algorithm sketched above to consistently estimate the law of motion of α_t. In each step, the entire model must be solved, including the high-dimensional household value function, and simulated.

An advantage of an algorithm based on (15) is that it offers a more efficient alternative to find this law of motion. Indeed, Proposition 3 below provides analytical formulas for $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$ that can be easily computed. In this sense, and in sharp contrast to the Krusell and Smith (1998) algorithm, finding the consistant law of motion incurs almost no increase in numerical complexity.

To develop intuition for this result, I present a heuristic but hopefully illustrative derivation for the special case where $\sigma_\alpha(\hat{\Gamma}) = 0$. This turns out to be true for the model of Section 2. Proposition 3 below provides more general formulas for $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$, which are developed formally in Appendices B.1 through B.4.

According to equation (13), the true law of motion of the distribution can be written as $dg_t(\boldsymbol{x}) = -\partial_a(sg) - \partial_k(mg) - \partial_z(\mu_z g) \equiv (\mathcal{A}^* g_t)(\boldsymbol{x})dt$. Under the approximation $\hat{g}_t(\boldsymbol{x}) \approx g_t(\boldsymbol{x})$, we have $d\hat{g}_t \approx dg_t$. The law of motion of the approximate cross-sectional distribution can therefore also be written as

$$d\hat{g}_t(\boldsymbol{x}) = (\mathcal{A}^* \hat{g}_t)(\boldsymbol{x})dt \tag{17}$$

where, abusing notation slightly, $\mathcal{A}^*$ is now evaluated in the approximate economy. Conjecturing that $d\alpha_t = \mu_\alpha(\hat{\Gamma})dt$ and differentiating $\hat{g}_t(\boldsymbol{x}) = F(\alpha_t)(\boldsymbol{x})$ with respect to time, we have

$$d\hat{g}_t(\boldsymbol{x}) = F_\alpha(\alpha_t)(\boldsymbol{x})\mu_\alpha(\hat{\Gamma})dt$$

where F_α is the gradient of F with respect to α. Using equation (17) and matching coefficients, we arrive at a functional equation that characterizes $\mu_\alpha(\hat{\Gamma})$ via

$$F_\alpha(\alpha_t)(\boldsymbol{x})\mu_\alpha(\hat{\Gamma}) = (\mathcal{A}^* F(\alpha_t))(\boldsymbol{x})$$

To "invert" this equation and solve for $\mu_\alpha(\hat{\Gamma})$ we must assume a specific (estimation) norm that we want to minimize. Under the $\mathbb{L}^2(\boldsymbol{x})$ norm, we arrive at $\mu_\alpha = (F_\alpha^T F_\alpha)^{-1} F_\alpha^T (\mathcal{A}^* F)$.

Proposition 3 is presented for the choice of norm $\mathbb{L}^2(\boldsymbol{x})$, and under the additional assumption that $\boldsymbol{x}_t$ follows a diffusion process.[37]

Proposition 3. *Let $\hat{\Gamma}_t = \{\rho_t, \alpha_t\}$ denote the aggregate state of an approximate economy under a distribution representation (15). Then:*

1. *The law of motion of α_t is given by*

$$d\alpha = \mu_\alpha(\hat{\Gamma})dt + \sigma_\alpha(\hat{\Gamma})dB \tag{18}$$

one-asset HANK model.

[37] A more general version of Proposition 3 where, for example, $\boldsymbol{x}_t$ is allowed to follow a jump-diffusion process, is derived in Appendices B.3 and B.4.

2. *The choice of μ_α and σ_α that minimizes forecast errors in the $\mathbb{L}^2$-norm is given by*

$$\mu_\alpha = (F_\alpha^T F_\alpha)^{-1} F_\alpha^T \left[\mathcal{A}^* F - \frac{1}{2} \sigma_\alpha^T F_{\alpha\alpha} \sigma_\alpha \right] \tag{19}$$

and

$$\sigma_\alpha = (F_\alpha^T F_\alpha)^{-1} F_\alpha^T \mathcal{B}^* F, \tag{20}$$

where F_α and $F_{\alpha\alpha}$ are the Jacobian and Hessian of F, respectively, $\mathcal{A}^*$ is the adjoint of the infinitesimal operator defined by the HJB equation, and $\mathcal{B}^*$ is the diffusion coefficient in the stochastic Kolmogorov forward equation for g_t.

Proposition 3 promises a shortcut in the design of algorithms to solve heterogeneous-agent models with aggregate risk: Equations (19) and (20) provide formulas to compute the internally consistent law of motion of α_t directly.[38] The most surprising and useful aspect of Proposition 3 is that each term that features in equations (19) and (20) is readily computable inside the household's value function iteration step. For illustration, consider applying Proposition 3 to the model of Section 2, which yields $\sigma_\alpha(\hat{\Gamma}) = 0$ (as conjectured above) and

$$\mu_\alpha = -(F_\alpha^T F_\alpha)^{-1} F_\alpha^T \left[\partial_a \left(\hat{s} F \right) + \partial_k \left(\hat{m} F \right) + \partial_z \left(\mu_z F \right) \right]. \tag{21}$$

For a given choice of $F(\cdot)$ and with α on the grid as an aggregate state variable, the terms $F(\alpha)$, F_α and $F_{\alpha\alpha}$ can be computed using standard numerical derivatives even before the value function iteration step. In many relevant cases, F_α and $F_{\alpha\alpha}$ will, in fact, be available as closed form analytical expressions. All remaining terms in equation (21) are comprised either of household policy functions, $\hat{s}$ and $\hat{m}$, or the employment transition rates, μ_z or equivalently λ^j.[39]

Algorithm structure. In light of these observations, I propose to compute $\mu_\alpha(\hat{\Gamma})$ and $\sigma_\alpha(\hat{\Gamma})$ directly as part of the household value function iteration (VFI) step. My algorithm therefore adds an additional step to an otherwise standard VFI procedure. The kth iteration follows these four steps:

1. Start with this iteration's guess $\hat{V}^k$.
2. Compute policy functions $\hat{s}^k$ and $\hat{m}^k$ using $\hat{V}^k$ and the household's FOCs.
3. Use $\hat{s}^k$ and $\hat{m}^k$ to compute $\mu_\alpha^k(\hat{\Gamma})$ and $\sigma_\alpha^k(\hat{\Gamma})$ according to equations (19) and (20).

[38] Ahn et al. (2017) make a similar argument. Relative to their paper, Proposition 3 applies to settings where the cross-sectional distribution of agents itself is stochastic, in which case $\mathcal{B}^* \neq 0$ and $\sigma_\alpha(\hat{\Gamma}) \neq 0$. Formally, I nest models in which the distribution's law of motion is given by a stochastic rather than an ordinary partial differential equation. Asset pricing models with portfolio choice problems are typically of this kind. See Appendix E for an example.

[39] In Appendices B.3 and B.4, I show that $\mathcal{A}^*$ and $\mathcal{B}^*$ are, more generally, comprised of household policy functions and transition rates.

4. Use (16) to find $\hat{V}^{k+1}$, and repeat from step (1.) until convergence.

Notice that, in iteration k of this modified VFI algorithm, $\mu_{\alpha}^{k}(\hat{\Gamma})$ and $\sigma_{\alpha}^{k}(\hat{\Gamma})$ are only required to evaluate (16) in step (4.) and update the value function $\hat{V}^{k+1}$. In other words, equations (19) and (20) are computed *after* the household policy functions and can therefore be readily evaluated. Taking general equilibrium prices as given, this modified VFI algorithm solves for both $\hat{V}$ and an internally consistent law of motion for α_t at the same time.[40] By contrast, the traditional approach that follows Krusell and Smith (1998) constructs an outer fixed point to solve for μ_{α} and σ_{α}.

3.4 How to Choose Efficient Representations $F(\cdot)$?

The next natural question is how F should be chosen in practice. One straightforward approach is to choose F directly from a parametric family.[41] When the functional form $F(\alpha)$ is chosen ex ante in this way, no simulation step is required at any point in the algorithm, which is one of its most compelling features.[42]

The approximation mapping (15) can also be implemented non-parametrically. The non-parametric approach can deliver large performance gains, especially when $\boldsymbol{x}_t$ is high-dimensional, but comes at the cost of an outer simulation step. For illustration, I will discuss the special case where $F(\cdot)$ is affine in $\alpha_t \in \mathbb{R}^N$, that is $F(\alpha_t)(\boldsymbol{x}) = C(\boldsymbol{x}) + T(\boldsymbol{x})\alpha_t$ for some functions C and T. In practice, I simply set $C(\boldsymbol{x})$ equal to the distribution associated with the deterministic steady state, that is $C(\boldsymbol{x}) = g^0(\boldsymbol{x}) \equiv \lim_{\sigma_\rho \to 0} g(\boldsymbol{x}; \sigma_\rho)$. As in Example 2, $T(\cdot)$ should be interpreted as a row-vector of basis functions and α_t as the column-vector of coefficients.

Algorithm structure. The non-parametric approach constructs an outer fixed point whose nth iteration follows three steps:

1. Given this iteration's distribution representation,

$$F^n(\alpha_t)(\boldsymbol{x}) = g^0(\boldsymbol{x}) + \sum_{k=1}^{n} T^k(\boldsymbol{x})\alpha_t^k, \qquad \alpha_t \in \mathbb{R}^n$$

[40] What is particularly nice about applying Proposition 3 in practice is that it prescribes a method for computing $\mu_{\alpha}(\hat{\Gamma})$ and $\sigma_{\alpha}(\hat{\Gamma})$ that is independent from the researcher's choice of F. In practice, therefore, the researcher can switch out different candidates for F with ease to determine which representation works best. In other currently popular global solution methods, the algorithm step to find the law of motion of α_t (or other moments) is oftentimes highly dependent on the particulars of the distribution approximation. It has therefore been difficult in practice to switch between and compare different solution methods (e.g. candidates for F) seamlessly.

[41] See Appendix B.5 for details. Parametric families that are commonly used include hat functions (nodal basis), Chebyshev polynomials, radial basis functions, generalized beta density functions and splines.

[42] The complexity of solving a model like the seminal Krusell and Smith (1998) benchmark is therefore reduced to a single value function iteration step that jointly computes $\hat{V}$ and the law of motion of α_t (see previous subsection). When the model requires an additional general equilibrium fixed point to solve for prices, the solution complexity is reduced to a single iteration of this fixed point.

solve the nth *approximate economy* that uses $\hat{g}_t^n(\boldsymbol{x}) = F^n(\alpha_t)(\boldsymbol{x})$ as distribution.[43]

2. Given a solution of approximate economy n, simulate the model to obtain simulated data for the distribution, $g_t^{n,\text{sim}}$. Compute the forecast error $F^n(\alpha_t)(\boldsymbol{x}) - g_t^{n,\text{sim}}(\boldsymbol{x})$.
3. Find the basis function $T^{n+1}(\boldsymbol{x})$ that minimizes forecast errors under a prespecified norm, subject to equation (18) for α_t. Update F^{n+1} and iterate until a desired accuracy is reached.

Intuitively, step (3.) looks for a kind of functional principal component with explanatory power for the forecast residual. At the cost of having to simulate the model, this non-parametric approach delivers substantial efficiency gains especially when the idiosyncratic state space of agents is high-dimensional.[44] To achieve a given level of accuracy, the number of basis functions n required in this non-parametric approach is often far smaller than in the parametric approach. In Appendix B.6, I provide additional details, including on the choice of forecast error norms and the resulting estimation optimality conditions.

4 Taking the Model to the Data

In this section, I calibrate the model and confront its main predictions with data. The model is calibrated on a quarterly frequency, based on U.S. data.

4.1 Calibration strategy

Table 1 provides a summary of all model parameters except for the employment transition rates that are discussed in Section 4.2. Households have CRRA preferences in consumption and leisure, with a relative risk aversion coefficient of $\gamma = 2$, and an inverse Frisch labor supply elasticity of $\eta = 2$. I calibrate $\bar{\rho}$, households' quarterly discount rate in the risky steady state, to match a natural rate of interest of roughly 2% (2.14% in the model) in the risky steady state. The deathrate ζ is calibrated to imply an average life span of 45 years, corresponding roughly to the average working life.

Households' portfolio choice is primarily affected by the adjustment cost on illiquid investments and the borrowing constraint on liquid assets. I set the latter to $\underline{a} = -1$, implying that households have access to roughly one quarter of average income in unsecured borrowing as in Kaplan et al. (2018). I calibrate the two parameters of the adjustment cost function to match the top 10% wealth share in liquid wealth (86% in both data and model), as well as the share of households with non-positive liquid net worth but a positive illiquid asset position

[43]This step requires solving the household problem (16), the union problem and the firm problem, as well as finding the general equilibrium prices that clear all markets. As discussed in the previous subsection, finding an internally consistent law of motion for α_t is subsumed in the modified VFI step.

[44]In the context of algorithms that work with a set of moments to approximate the distribution $g_t(\boldsymbol{x})$, the fixed point I describe is akin to the idea of successively adding moments to improve forecast accuracy. In my setting, this approach is complicated by the flexibility of being able to choose any $F(\cdot)$.

Table 1: List of Calibrated Parameters

	Parameters	Value	Target / Source
	Preferences		
$\bar{\rho}$	Discount rate (p.q.)	2.2 %	Riskfree rate
γ	Relative risk aversion	2	Standard
η	Inverse Frisch elasticity	2	Standard
ζ	Deathrate	0.556 %	Average life span 45 years
	Household portfolio choice		
$\underline{a}$	Borrowing constraint	-1	1.244 q average income
δ	Capital depreciation (p.q.)	1.25 %	
ψ_0	Linear adjustment cost	0.044	Top 10% (liquid) wealth share
ψ_1	Convex adjustment cost	0.956	% households with $a \leq 0$ and $k > 0$
	Firms		
α	Capital share	0.38	
κ	Aggregate capital adjustment cost	40	$\Delta I/\Delta Y$ after $\Delta\sigma_\rho$
	Nominal rigidities		
$\frac{\epsilon^f}{\epsilon^f-1}$	Elasticity of substitution	1.10	CEE (2005)
χ^f	Price adjustment cost	350	ACEL (2011)
$\frac{\epsilon^w}{\epsilon^w-1}$	Elasticity of substitution	1.05	CEE (2005)
χ^w	Avg. duration of wage contracts	0	Flexible-wage limit
	Government		
B^G	Government debt outstanding	0.06	Average aggregate MPC
G	Government spending	0	
λ_π	Taylor rule weight on inflation	1.5	Standard
λ_Y	Taylor rule weight on output	0.5	Standard
τ^{lab}	Income tax rate	0.2	
τ^{UI}	Unemployment insurance	0.20	Chodorow-Reich and Karabarbounis (2016)
	Aggregate uncertainty		
σ_ρ	Volatility: discount rate shock	0.003	Volatility of GDP growth
θ_ρ	Persistence: discount rate shock	0.22	Autocorrelation of GDP growth

(30% in the data, 27% in the model's risky steady state). In calibrating these parameters, I stay as close as possible to Kaplan et al. (2018) since the wealth distribution is not the primary focus of my contribution.

I set depreciation to 1.25% per quarter and the capital income share α to 0.38. I calibrate the aggregate capital adjustment cost parameter, κ, to match the business cycle sensitivity of investment, relative to output. In particular, the peak decline of aggregate investment following a fundamental risk shock is currently three times as large as that of output while aggregate consumption is less responseive.

The parameters governing nominal price rigidity are set equal to literature benchmarks.

In particular, I set the elasticity of substitution so that $\frac{\epsilon^f}{\epsilon^f - 1} = 1.1$ as in Christiano et al. (2005), and I set χ^f to imply an average price duration of roughly 7 quarters. Altig et al. (2011) estimate an average price duration of 9.36 quarters. I solve the baseline model in the limit as wages are fully flexible, $\chi^w \to 0$.[45]

The baseline model does not feature government spending, and the stock of government debt outstanding is constant. One weakness of the model with only two earnings types is that it is difficult to jointly match high MPCs and a high stock of liquid assets. In particular, matching an aggregate household MPC of roughly 15% requires setting the stock of outstanding government debt close to zero ($B^G = 0.06$). I set the income tax rate to 20%, and the effective unemployment insurance replacement rate to 20%. This is consistent with recent estimates accounting for take-up in Chodorow-Reich and Karabarbounis (2016). The central bank follows a Taylor rule with the standard literature weights on inflation and output of, respectively, 1.5 and 0.5.

Finally, I calibrate the volatility and persistence of the discount rate shock process to match the business cycle moments of GDP growth, with a quarterly auto-correlation of 0.85 and a standard deviation of 0.017.

4.2 The ins and outs of unemployment

The employment transition rates faced by households play a central role in the quantitative results of this paper. I use microdata on employment transitions from the Current Population Survey (CPS) to estimate the sensitivity of the job finding and separation rates over the business cycle.[46] I use the resulting reduced-form estimates for $\lambda^j(\cdot)$ directly in the model. For additional details, see Appendix C.1.

Exploiting the rotating panel nature of the CPS, I match household observations across subsequent months to construct monthly gross flows into and out of employment in a sample from 1996 to 2019. I employ a broad notion of non-employment including unemployed and marginally attached households, as well as those in involuntary part-time employment for economic reasons. This definition corresponds to the U6 unemployment rate. I show that these groups also exhibit sizable transition rates into and out of employment, which, following the logic of my model, motivates their inclusion.

Having constructed matched gross flow data at a monthly frequency, I use a continuous-time conversion formula to compute exact data counterparts to the Poisson arrival rates λ_t^E and λ_t^U in the model. In the spirit of Shimer (2012), this conversion also addresses a time aggregation bias in discrete-time flow data. Panels (c) and (d) of Figure 2 display the data time series for these job finding and separation rates.

[45] The cost of assuming flexible wages is, of course, the standard problem that the profits of goods producing firms are counter-cyclical. In the next iteration of the paper, I plan to allow for both price and wage rigidity.

[46] There is a long tradition of constructing gross employment flows from matched CPS micro data. See for example Marston et al. (1976), Abowd and Zellner (1985), Darby et al. (1985), Darby et al. (1986), Poterba and Summers (1986), Blanchard et al. (1990), Shimer (2005), Fujita and Ramey (2009), Elsby et al. (2009), and Shimer (2012). In particular, I closely follow Shimer (2012) in the construction of monthly gross flow data.

Figure 2: Employment Transition Rates over the Business Cycle

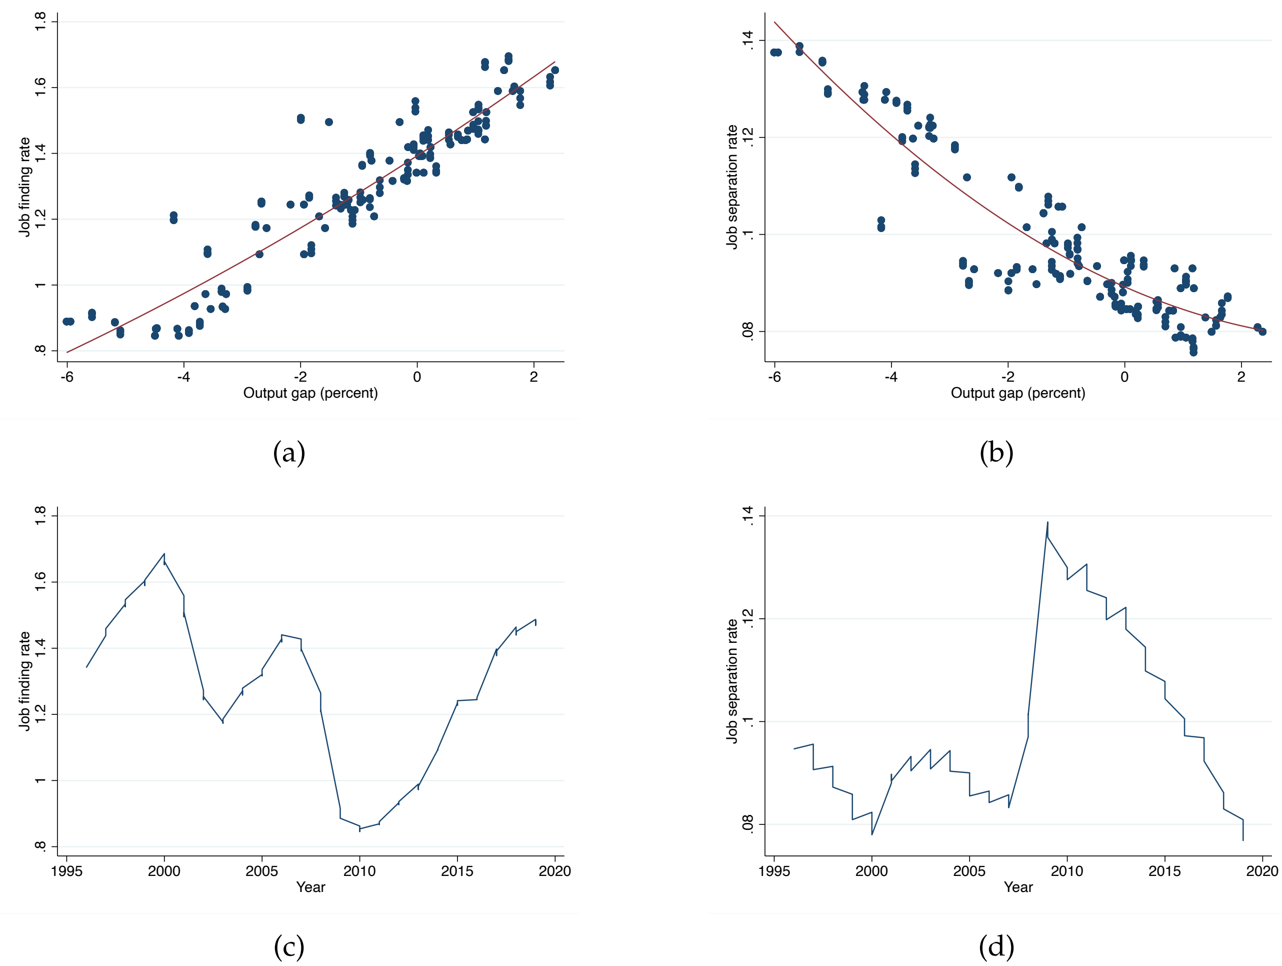

(a) (b) (c) (d)

Notes. The raw CPS micro data for job finding and separation rates from 1996 to 2019 are displayed in Panels (c) and (d). The data are expressed as instantaneous, quarterly transition rates after a conversion from monthly gross flow data to continuous time. The underlying definition of unemployment used is broader than U3 headline unemployment and corresponds closely to U6. Panels (a) and (b) scatter the raw data against output gap estimates. The solid red lines correspond to a quadratic fit.

Since these data objects correspond exactly to the instantaneous transition rates (at quarterly frequency) faced by households in the model, I use an external calibration approach to specify how λ_t^E and λ_t^U move over the business cycle. In the baseline calibration, I use the shortcut specification that the job finding and separation rates are only functions of the output gap, so that $\lambda_t^j = \lambda^j(y_t)$. I implement this specification in the data using a Taylor expansion

$$\text{Job finding rate :} \quad \lambda_t^U = \underset{(114.96)}{1.39} + \underset{(13.66)}{0.115}\, y_t + \underset{(1.08)}{0.0026}\, y_t^2 + \ldots \tag{22}$$

$$\text{Job separation rate :} \quad \lambda_t^E = \underset{(88.67)}{0.89} - \underset{(-6.85)}{0.0053}\, y_t + \underset{(3.62)}{0.0006}\, y_t^2 + \ldots \tag{23}$$

and report the associated T-statistics after a Newey-West correction. Panels (a) and (b) of Figure 2 display the associated scatter plots. There is strong evidence of systematic time variation in employment transition rates over the business cycle. The job finding (separation) rate falls (increases) in bad times. Similarly, equation (23) shows that at least the job separation rate exhibits a non-linearity that is statistically significant.

In my baseline calibration, I use equations (22) and (23) directly to specify λ^j as a function of the economy's aggregate state.[47] In Appendix D, I confirm that the main quantitative results of Section 5 are robust to alternative functional form assumptions for λ_t^E and λ_t^U.

4.3 Business cycle fluctuations

The model matches the asymmetry of business cycles in U.S. postwar history: recessions have been deeper and more pronounced in the data than economic booms.

My calibration only targets the volatility of GDP growth (0.017 in the data and 0.014 in the model) and the auto-correlation of GDP growth (0.85 in the data and 0.88 in the model). The model also matches other untargeted business cycle moments at least qualitatively. In Section 5.1, I show that demand shocks generate co-movement in output, consumption, investment and hours in the model. In particular, aggregate investment is more volatility and responsive than output, while aggregate consumption is less responsive.

Qualitatively, the model also matches higher moments of key business cycle aggregates. Figure 3 compares the distribution of GDP growth realizations in model simulations (yellow bars) against U.S. postwar data (blue bars). The distribution of GDP growth in U.S. data since 1953 has fatter tails than in my model simulations. Nonetheless, output, consumption, investment and hours in model simulations exhibit significant negative skewness and positive kurtosis. While the skewness and kurtosis generated by the model are larger than their respective data counterparts, this is not unexpected in a meaningful sense. While the ZLB has only become a binding constraint for U.S. monetary policy quite recently, most recessions in my model simulations drive the economy to the ZLB. The model is therefore intended to represent this more recent episode in U.S. economic history, characterized by a structurally much lower natural rate of interest. Indeed, if I recompute the U.S. business cycle moments

[47] The empirical transition rate series I have constructed do not exactly aggregate back up to the aggregate unemployment rate (U6) observed in the data. To ensure that the model matches the average unemployment rate, I adjust the constant term in equation (22). That is, my model matches the average job separation rate and the average U6 unemployment rate, as well as the cyclicality exhibited by the job finding and separation rates in the sample from 1996 to 2019.

Figure 3: Distribution of GDP growth in U.S. postwar data and model simulations

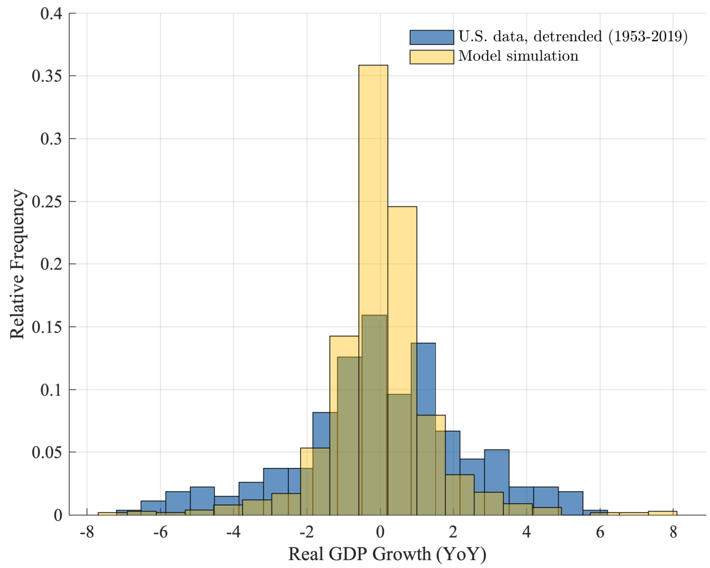

Notes. Relative frequency histogram of distribution of year-over-year GDP growth in U.S. postwar data since 1953 (blue bars) and model simulations (yellow bars). Data is detrended using a Kalman filter.

over a sample of the Great Moderation, starting in 1980, the implied skewness of output is almost as high as that exhibited by my model. Of course, the high negative skewness in that sample is driven almost entirely by the Great Recession.

Simulations of the model imply a zero lower bound (ZLB) incidence between 20 and 30%. Between 1981 and 2018, the U.S. spent 7 out of 48 years, or 14.5% of the time, at the ZLB. ZLB incidence is likely dependent on the natural rate of interest, since a lower equilibrium rate implies a closer proximity to the constraint. Laubach and Williams (2003) and Laubach and Williams (2016) have estimated the natural real rate to have declined from between 3 and 4% in the 1980s to below 1% since the Great Recession. If this decline proves persistent, the incidence of the ZLB will likely be much higher in the future.

5 Micro and Macro Uncertainty

This section develops the paper's main results on micro and macro uncertainty. Macroeconomic uncertainty is an endogenous equilibrium object in this model. Let $\boldsymbol{Z}_t$ denote the vector of all aggregate variables, such as the wage rate, the capital price, inflation etc., and $Z_t \in \boldsymbol{Z}_t$. The notion of uncertainty I adopt in this paper is one of conditional forecast errors, as is typical in most of the uncertainty literature. Specifically, I define

$$\mathcal{U}_t(Z_{t+s}) = \sqrt{\frac{1}{s}\mathbb{E}_t\left[\Big(Z_{t+s} - \mathbb{E}_t(Z_{t+s})\Big)^2\right]} \tag{24}$$

as uncertainty over future realizations in the aggregate process $\{Z_t\}$. While the macroeconomic uncertainty faced by households extends to the broader set of all components of the stochastic process $\{\boldsymbol{Z}_t\}$, I will oftentimes use uncertainty over economic activity, in short GDP $\{Y_t\}$, as a proxy for macroeconomic uncertainty. And since output follows the process $dY_t = \mu_Y(\Gamma_t)dt + \sigma_Y(\Gamma_t)dB_t$ in equilibrium, I concretely refer to

$$\mathcal{U}_t(Y_{t+dt}) = \underbrace{\sigma_Y(\Gamma_t)}_{\text{"Macroeconomic uncertainty"}}$$

as macro uncertainty at time t.

Isolating the effects of macro uncertainty. The first goal of this section is to quantitatively assess the novel transmission channels identified in Section 1. This requires a procedure to isolate and decompose the effects of variation in macro uncertainty. While this is possible for equilibrium objects that explicitly appear in the equilibrium conditions of the model, I am not aware of such a procedure for endogenous uncertainty, which is instead a property of the model's probability distribution.[48]

To study the transmission mechanism of macro uncertainty in the quantitative model, I instead propose to consider an exogenous fundamental risk shock, $\Delta\sigma_\rho$. This requires adding σ_ρ explicitly to the model's aggregate state space (see Appendix D for details). The assumption implicit in this strategy is that the transmission mechanism of an exogenous increase in the volatility of discount rate shocks, σ_ρ, is sufficiently similar to that of an endogenous increase in the volatility of economic activity, σ_Y. Indeed, I show formally in Section 5.3 that endogenous macroeconomic uncertainty is proportional to exogenous fundamental risk, which justifies this approach. Overall, the following analysis of an exogenous fundamental risk shock should be understood primarily as a means to better understand the transmission mechanism of endogenous variation in macroeconomic uncertainty.

Comparison benchmarks. To illustrate my main results, I compare three distinct model benchmarks. The first, which I refer to as the "HANK (full)" benchmark, is the full quantitative model presented in Section 2. The second benchmark, which I refer to as "HANK (constant λ^j)", holds employment transition rates constant and therefore shuts off the interaction between micro and macro uncertainty. Importantly, these two model benchmarks are identical in all other regards. Therefore, a comparison between these two models allows us to approximately isolate the implications of the interaction between micro and macro uncertainty. Finally, I also relate my results to a "RANK" baseline, which is the representative-household New Keynesian model most closely associated with the full HANK model of Section 2.

[48] Kaplan et al. (2018) propose a procedure to decompose the partial effects of equlibrium objects that explicitly appear in the household problem, such as the interest rate r_t. Their method applies to the sequence representation of a model without aggregate risk. I generalize their approach to state space representations of models with aggregate risk. However, neither approach allows for a decomposition of the partial effect of endogenous uncertainty. This is because, unlike the aggregate price r_t, uncertainty over its future realizations, σ_r, does not explicitly appear in the household problem and can therefore not "be held fixed".

5.1 Transmission mechanism of macro uncertainty

As discussed above, I study the transmission mechanism of an exogenous perturbation in fundamental risk, $\Delta\sigma_\rho$, in order to proxy for and better understand the effects of macroeconomic uncertainty in this model.[49] Higher uncertainty elicits an increase in households' desired precautionary savings. This direct effect of uncertainty, borrowing the language of Kaplan et al. (2018), leads to an initial contraction in aggregate demand. In the two-period model of Section 1, this effect comprises both the direct macro uncertainty channel ① and the direct interaction channels ② and ④. As household consumption falls, firms' demand for labor decreases in general equilibrium. Consequently, the job finding rate, λ_t^U, falls while the separation rate, λ_t^E, rises. Households are now confronted with an endogenous increase in micro uncertainty: The employed are more likely to be laid off and the unemployed are less likely to find a new job. This GE effect that follows the contraction in aggregate demand corresponds to channel ③ in the two-period model.

Figure 4 plots both the first- and second-moment response of employment transition rates after an increase in fundamental risk. In expectation, the fundamental risk shock makes it more likely for households to be laid off ($\lambda^E \uparrow$) and less likely to find jobs ($\lambda^U \downarrow$). At the same time, the conditional variance over future realizations of the job finding and separation rates increases substantially. In short, a fundamental risk shock at the macro level translates into uncertainty over job prospects at the micro level.

From the household's perspective, the prospect of job loss is akin to an idiosyncratic rare disaster as I discussed in Section 1: It looms large when compared to relatively small and transitory changes in the aggregate wage rate, for example. Endogenous variation in the transition rates into and out of employment has both a first- and second-moment implication for households' future earnings. Consider an employed household. An increase in the job separation rate lowers the expected value of future income because the probability of unemployment rises. At the same time, the variance of earnings increases as well.

The overall effect of a fundamental risk shock is summarized in Figure 5. It compares the impulse responses of output, consumption, investment and hours across the three model benchmarks. As in Basu and Bundick (2017), a fundamental risk shock generates co-movement across the desired aggregates. The main result of this exercise is that the effect of macro uncertainty on aggregate demand is substantially amplified and indeed driven by its interaction with micro uncertainty. The on-impact decline in output in the full model is up to 8 times larger than in the comparison benchmark that holds employment transition

[49] While my interest in studying the exogenous perturbation $\Delta\sigma_\rho$ is as a proxy for variation in endogenous macro uncertainty, this approach of course builds on a long tradition of studying exogenous second-moment shocks. Following the Great Recession, much work has studied uncertainty as a potential driver of business cycle fluctuations (e.g. Bloom (2009) and Bloom et al. (2018) in the context of firms, Bayer et al. (2019) focuses on micro uncertainty shocks). Basu and Bundick (2017) is a natural reference point for my paper. They show that, in a representative-agent New Keynesian (RANK) model, a fundamental risk shock generates co-movement in output, consumption, investment and hours worked. In a setting where households have Epstein-Zin preferences with a relative risk aversion coefficient of 80, they find that a 90% increase in the volatility of discount rate shocks generates a 0.2% fall in output due to an increase in desired precautionary savings.

Figure 4: Impulse responses of micro uncertainty to fundamental risk shock, $\Delta\sigma_\rho$

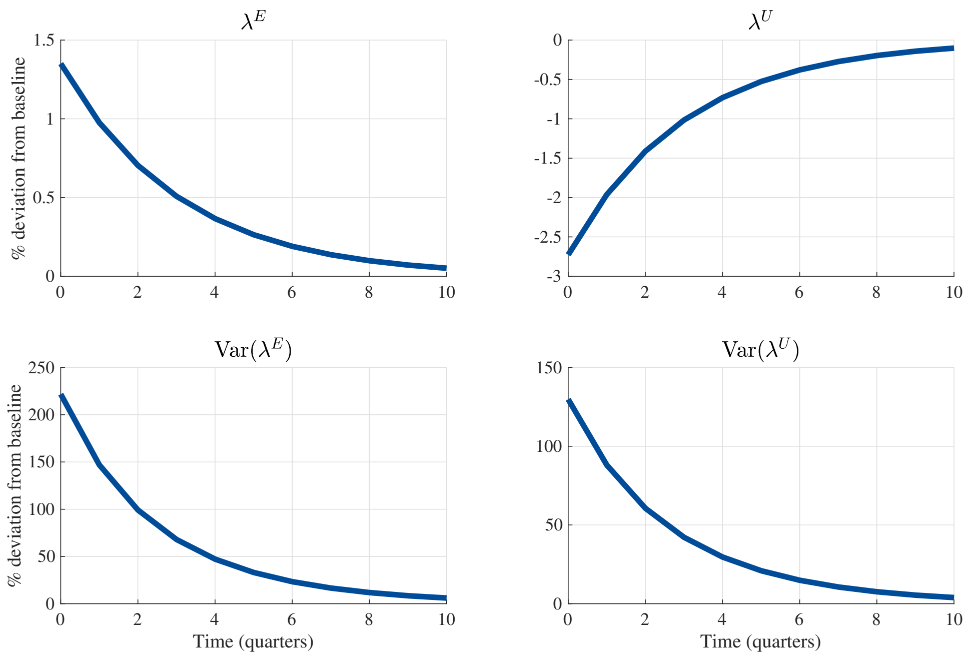

Notes. Comparison of impulse responses of employment transition rates to fundamental risk shock. The shock has a half-life of 1 quarter and is initialized at half the size of the Basu and Bundick (2017) shock to facilitate comparison. λ^E and λ^U correspond to the job separation and finding rates, respectively. The bottom row panels plot the response of conditional forecast errors, $\mathrm{Var}_t(\lambda^j_{t+dt})$.

rates constant and thereby shuts off any interaction between micro and macro uncertainty. Similarly, the on-impact decline in output in the full model is rougly 5 times larger than in the associated RANK benchmark.[50]

Result 1. *The effects of an exogenous increase in fundamental risk, $\Delta\sigma_\rho$, on economic activity are substantially amplified – by a factor of up to 8 – when accounting for the interaction between micro and macro uncertainty.*

Table 2 provides a decomposition of the transmission mechanism of a fundamental risk shock. I decompose the shock's effect on aggregate consumption into a direct uncertainty channel

[50] Basu and Bundick (2017) is an instructive reference point for these quantitative results. They study a fundamental risk shock that is about twice as large as the one I consider here and show that, when households have Epstein-Zin preferences with a relative risk aversion coefficient of 80, the on-impact decline in output is about 0.2%. Their model is, of course, different from and in many aspects richer than the RANK baseline I consider here. Nonetheless, it is noteworthy that I reproduce quantitatively similar magnitudes in a setting where households have CRRA preferences with a relative risk aversion coefficient of 2: For a given level of risk aversion, households in my setting are more responsive to a change in macro uncertainty precisely because it translates into micro uncertainty.

Figure 5: Impulse responses of business cycle aggregates to fundamental risk shock, $\Delta\sigma_\rho$

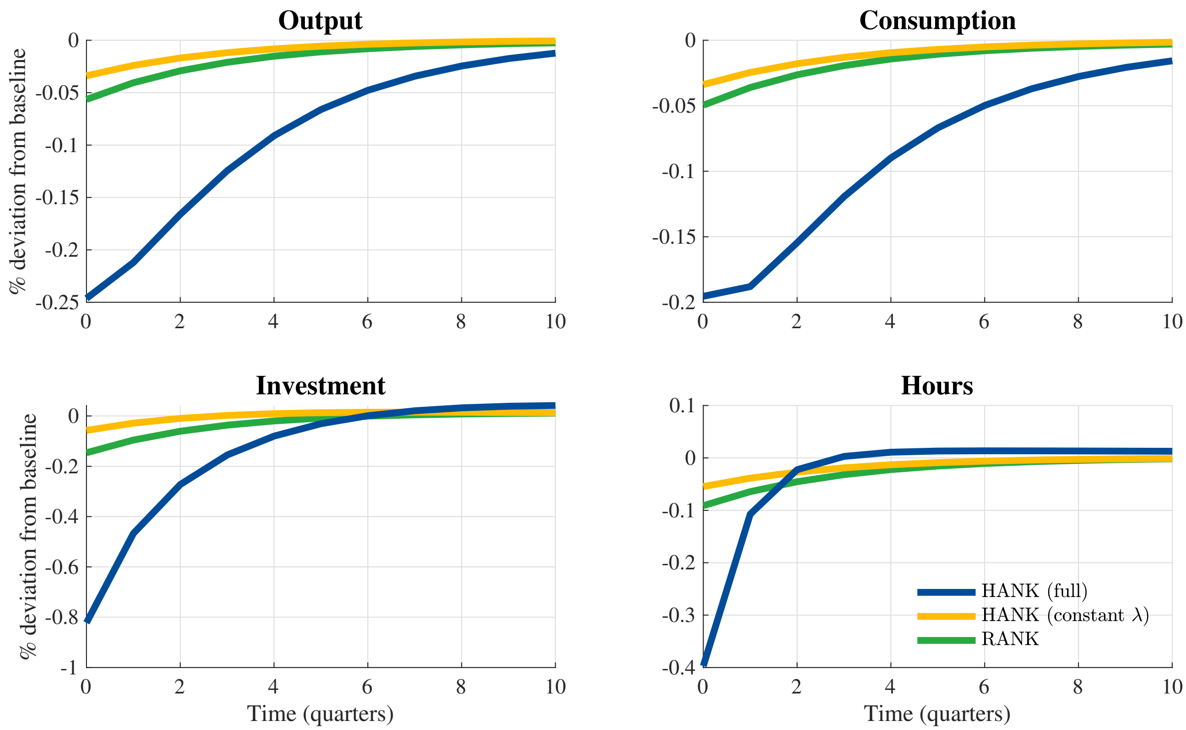

Notes. Comparison of impulse responses of business cycle aggregates to fundamental risk shock across three models. The shock has a half-life of 1 quarter and is initialized at half the size of the Basu and Bundick (2017) shock to facilitate comparison. The blue, yellow and green lines correspond, respectively, to the baseline model, the model that shuts off the interaction between micro and macro uncertainty, and the associated RANK model.

and a set of indirect (GE) channels. This classification extends Kaplan et al. (2018) to a setting with aggregate uncertainty. When considering the direct effect of uncertainty, I hold fixed all aggregate variables that would respond in general equilibrium. Households still realize, however, that volatility going forward is heightened. When computing the contribution of indirect channels, I turn back on, one by one, the GE response of the macroeconomic aggregates households care about.

The two rows of Table 2 that are most relevant to this discussion are the first (the direct uncertainty effect) and the second (the indirect effect through micro uncertainty). These conceptually correspond, respectively, to the sum of channels ①, ② and ④ (direct) and channel ③ (indirect) in the two-period model of Section 1. Table 2 demonstrates that the transmission mechanism of macro uncertainty changes completely relative to a representative-agent benchmark. In the RANK model, the direct macro uncertainty effect (channel ①) is the dominant transmission channel.

The transmission mechanism of macro uncertainty in the full HANK model differs in three key regards. First, the direct uncertainty effect on household behavior (the first row of Table 2) is larger in absolute terms but muted relative to other transmission channels. This is

Table 2: Decomposition of the effect of fundamental risk shock $\Delta\sigma_\rho$ on aggregate consumption

	HANK		RANK
Contribution to % change in C_0	*Normal times*	*Crisis region*	*Normal times*
Direct effect: uncertainty (micro and macro)	−0.19	−0.22	−0.05
Indirect effect: micro uncertainty	−0.47	−0.51	0.00
Indirect effect: disposable income	0.58	0.66	0.02
Indirect effect: portfolio returns	0.04	−0.12	−0.03
Other effects	−0.16	−0.12	0.01
Total effect (% change in C_0)	−0.2	−0.31	−0.05

Notes. Numbers correspond to on-impact responses, and partial effects add up to the total effect in the last row. The first two columns correspond to the full model; the last column corresponds to the associated RANK model. Columns 1 and 2 compare the partial and total effects when the economy is initialized in the risky steady state (normal times) and on the cusp of the ZLB (crisis region), respectively. The disposable income channel comprises the effects through w_t, τ_t and H_t, and the portfolio returns channel those through r_t, r_t^k and Q_t.

consistent with the observation from the two-period model that, while ① is the only direct channel operative in the RANK model, the direct interaction effects ② and ④ emerge in the HANK model. Second, the indirect interaction effect through micro uncertainty (channel ③ in the two-period model, and the second row of Table 2) emerges as a dominant transmission channel, especially during times of crisis. Finally, and in the spirit of Kaplan et al. (2018), other indirect channels become relatively more important than in the RANK baseline.[51]

Result 2. *Accounting for the interaction between uncertainty at the micro and macro level substantially changes the transmission mechanism of the fundamental risk shock $\Delta\sigma_\rho$. Its indirect effect through micro uncertainty becomes a dominant channel in the transmission of macro uncertainty to aggregate consumption spending.*

Studying a fundamental risk shock offers a first glimpse at another major theme of this paper. The behavior of uncertainty changes substantially during times of crisis. The effects of macro uncertainty on economic activity are significantly larger when the economy is close to the ZLB. The dashed lines in Figure 5 trace the impulse responses to the same fundamental risk shock as above in an economy that is initialized on the cusp of the ZLB. The shock's on-impact effect on output is 50% larger than during normal times. Similarly, Table 2 highlights that the transmission mechanism of macro uncertainty works even more strongly through micro

[51] The strong positive response through disposable income is largely driven by profit rebates. In the next iteration of this paper, I plan to address this well-known issue by solving the model with both price and wage stickiness.

uncertainty when the economy is in a crisis. This finding speaks to substantive interactions between non-linearities at the micro and macro level.[52]

When the economy is in a crisis state, the relationship between macroeconomic aggregates exhibits a degree of skewness.[53] In a crisis, a given negative demand shock realization generates a larger contraction in output than the expansion that would follow a similarly sized positive shock realization. In other words, a zero-mean spread in the distribution of future discount rate shocks, which is how households perceive the fundamental risk shock, no longer simply leads to a second-moment but also to a first-moment direct effect. In particular, due to negative skewness, this mean-zero spread leads to a fall in economic activity in expectation. The effect of macroeconomic uncertainty on aggregate demand is therefore much larger when the economy is in a crisis state, as showcased in Figure 5.

The implications of such skewness at the macro level are not unlike the implications of skewness at the micro level, which I discussed earlier. Indeed, when the economy is in a crisis, this non-linearity at the macro level starts to interact with non-linearity at the micro level. Uncertainty over job prospects at the micro level increases more strongly when the initial fundamental risk shock has a direct first-moment effect at the aggregate level. As a result, the relative importance of the interaction between micro and macro uncertainty becomes even larger in the transmission of fundamental risk shocks.

Result 3. *The behavior of uncertainty changes drastically during economic crises. The effect of a fundamental risk shock on economic activity is 50% larger when the economy is already on the cusp of the ZLB. The interaction with micro uncertainty becomes a relatively even more important channel in the transmission of macro uncertainty.*

The model exhibits a non-linearity at the aggregate level during times of crisis that is nearly impossible to capture in a HANK model without aggregate risk or one that is solved using (lower-order) perturbation methods. Indeed, even higher-order perturbation methods may fail because they still take the economy's normal state as the point of approximation. The macro non-linearity that gives rise to my third result does not manifest while the economy is far away from the crisis region, thus requiring a global solution method.

Of the main results in this paper the ones presented in this subsection are most closely framed with reference to previous work on uncertainty shocks. I will thus conclude by briefly situating my contribution. In the household context, this paper is the first to jointly consider variation in micro and macro uncertainty. Relative to Basu and Bundick (2017), for example, I show that its interaction with unemployment risk emerges as the dominant transmission

[52] Much previous work has studied the implications of the ZLB constraint on monetary policy in the context of representative-agent New Keynesian (RANK) models (e.g. Christiano et al. (2011), Fernández-Villaverde et al. (2015) or Plante et al. (2018)). My focus here is primarily on the interaction between non-linearities at the micro and macro level. Indeed, I show first in Section 5.2 and later in Section 6 that the amplification that derives from the ZLB constraint is itself more pronounced when we account for the interaction between micro and macro uncertainty.

[53] As is true in U.S. business cycle data, my model generates significant negative skewness and positive kurtosis in output, investment, consumption and hours. See Section 4.3.

Figure 6: Impulse responses of business cycle aggregates to $\Delta\sigma_\rho$ shock during normal times and crises

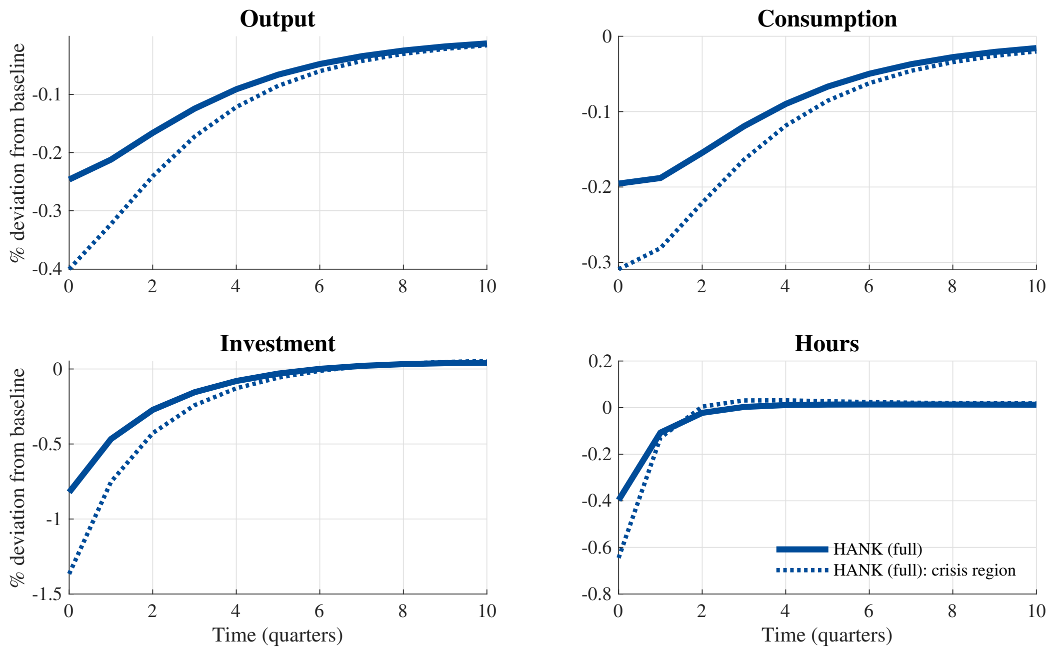

Notes. Comparison of impulse responses of business cycle aggregates to fundamental risk shock across regions of the state space. The shock has a half-life of 1 quarter and is initialized at half the size of the Basu and Bundick (2017) shock to facilitate comparison. The solid lines plot IRFs when the economy is initialized in the risky steady state; the dashed lines plot IRFs when the economy is initialized on the cusp of the ZLB.

channel of a macro uncertainty shock. And since micro uncertainty is much more prominent from the household's perspective, my paper matches the magnitude of their quantitative results with a much lower relative risk aversion coefficient. Bayer et al. (2019) study micro uncertainty shocks but abstract from time variation in fundamental risk at the aggregate level. They also linearize the macro block of their model and therefore do not capture endogenous variation in macro uncertainty. In the firm context, Bloom et al. (2018) model exogenous variation in both the micro and macro uncertainty that firms face. My main insight is that the interaction with *endogenous* micro risk becomes an important transmission channel of macro uncertainty, an interaction that is not modeled in their paper. Lastly, this subsection has demonstrated that the behavior of uncertainty changes in the crisis region of the model, which motivates my emphasis on a global solution method.

5.2 Endogenous uncertainty spikes during crises

In this subsection, I show that endogenous variation in uncertainty becomes substantially more pronounced when we account for the interaction between micro and macro uncertainty. In particular, the model generates large endogenous uncertainty spikes in the crisis region. Therefore, a two-way feedback loop emerges between macro uncertainty on the one hand,

and micro uncertainty and aggregate demand on the other: Recessions are times when a contraction in economic activity spurs endogenous spikes in uncertainty about the future, which in turn depresses aggregate demand further.

The model features at least two channels through which macroeconomic uncertainty becomes endogenous to economic activity. The first channel centers around the zero lower bound constraint. When the economy is at the ZLB, monetary policy can no longer accommodate negative demand shocks, whose effects on economic activity are consequently amplified. Even when the nominal interest rate is still positive, expansionary monetary policy moves the economy closer to the ZLB, thus raising the likelihood that policy will be constrained in the future. This channel is particularly strong in my baseline model because I do not explicitly take into account the endogeneity of fiscal policy stepping in as a substitute for monetary policy during deep recessions.[54]

The second channel results from the counter-cyclicality in households' average marginal propensity to consume (MPC). As economic activity contracts, households become unemployed and draw down their liquid cash buffers, thus moving closer to their borrowing constraints at the micro level. The prevalence of "hand-to-mouth" behavior grows, which implies an increase in the average household's MPC. Consumer spending, and by implication aggregate activity, thus become more sensitive to further demand shocks. The economy's sensitivity to fundamental risk is precisely the definition of macro uncertainty in my framework, which therefore rises.

This second channel interacts strongly with the fact that the length and severity of a recession are uncertain themselves. Therefore, unlike in settings with no aggregate risk, it becomes difficult for households to properly time the exhaustion of their liquid insurance funds. The closer households get to running out of their cash buffers, the more risk averse they become in the face of a potentially prolonged crisis.

What factors determine how sensitive macroeconomic uncertainty is to a given contraction in economic activity? The main novel insight my paper establishes in this context is that first-moment demand shocks have a much larger effect on endogenous macro uncertainty when we account for its interaction with micro uncertainty. Figure 7 shows the impulse response of macro uncertainty, $\sigma_Y(\Gamma_t)$, to a negative discount rate shock $\Delta\sigma_\rho$. The left panel shows that the response of endogenous macro uncertainty to a negative demand shock is substantially dampened when we shut off the interaction between micro and macro uncertainty. Indeed in the comparison baseline that holds employment transition rates constant, uncertainty hardly responds at all to the demand shock during normal times. This stark difference further highlights why accounting for the interaction between uncertainty at the micro and macro level is crucial in order to understand its role in business cycle fluctuations.

Proximity to the crisis region emerges as another key determinant of the sensitivity of macro uncertainty to economic activity. The right panel traces out the impulse response of macro uncertainty to the same underlying demand shock when the economy starts on the

[54] Plante et al. (2018) show that the ZLB becomes a source of endogenous uncertainty in the context of a representative-agent New Keynesian model.

Figure 7: Impulse response of endogenous macro uncertainty to first-moment discount rate shock

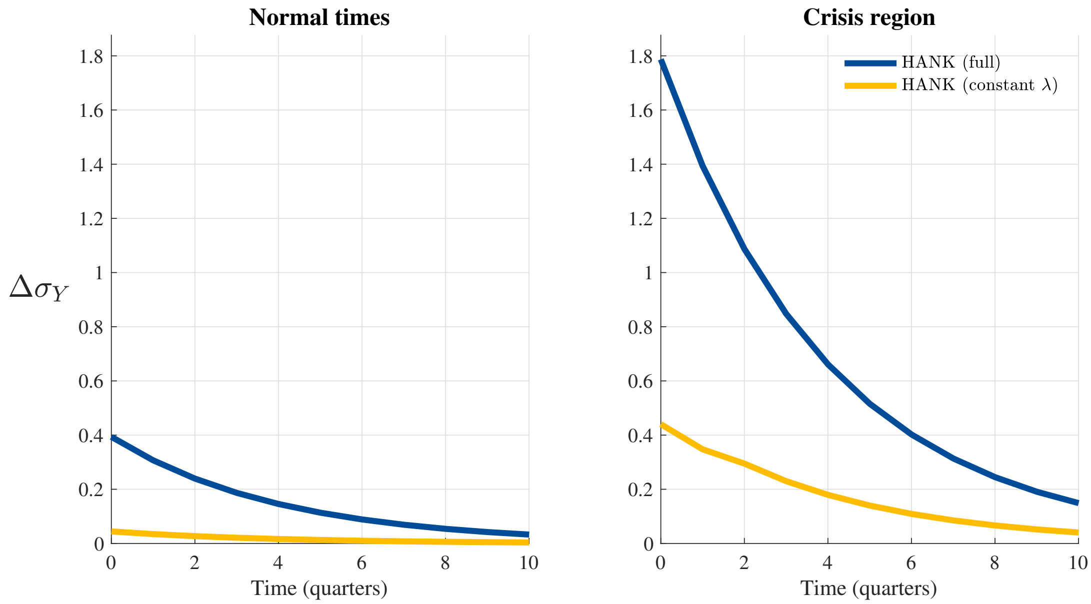

Notes. Comparison of impulse responses of endogenous macro uncertainty, $\sigma_{Y,t}$, to first-moment discount rate shock across different models and state space regions. The shock has a half-life of 1.4 quarters and is initialized at $0.2(\min(\rho_t) - \bar{\rho})$. The blue and yellow lines correspond, respectively, to the baseline model and the model that shuts off the interaction between micro and macro uncertainty. The left and right panels initialize the economy, respectively, in the risky steady state (normal times) and on the cusp of the ZLB (crisis region).

cusp of the crisis region rather than in the risky steady state. The spike in macro uncertainty after a contraction in economic activity during a crisis is over 3 times as large as during normal times. Figure 7 also highlights a quantitatively meaningful interaction between non-linearities at the micro and macro level: proximity to the crisis region makes macro uncertainty especially sensitive to economic activity precisely when we account for its interaction with micro uncertainty.

Result 4. *Uncertainty is not only a driver but also an endogenous byproduct of business cycle fluctuations in the full model. The endogenous response in macro uncertainty to a first-moment demand shock is largely driven by the interaction with micro uncertainty. Indeed, uncertainty hardly responds at all to a demand shock during normal times when we shut off the interaction between micro and macro uncertainty.*

In conclusion, macro uncertainty is most sensitive to aggregate demand and therefore most likely to exhibit endogenous spikes when we account for its interaction with micro uncertainty and when the economy is already in the crisis region.

Figure 8: General equilibrium interaction between micro, macro uncertainty and aggregate demand

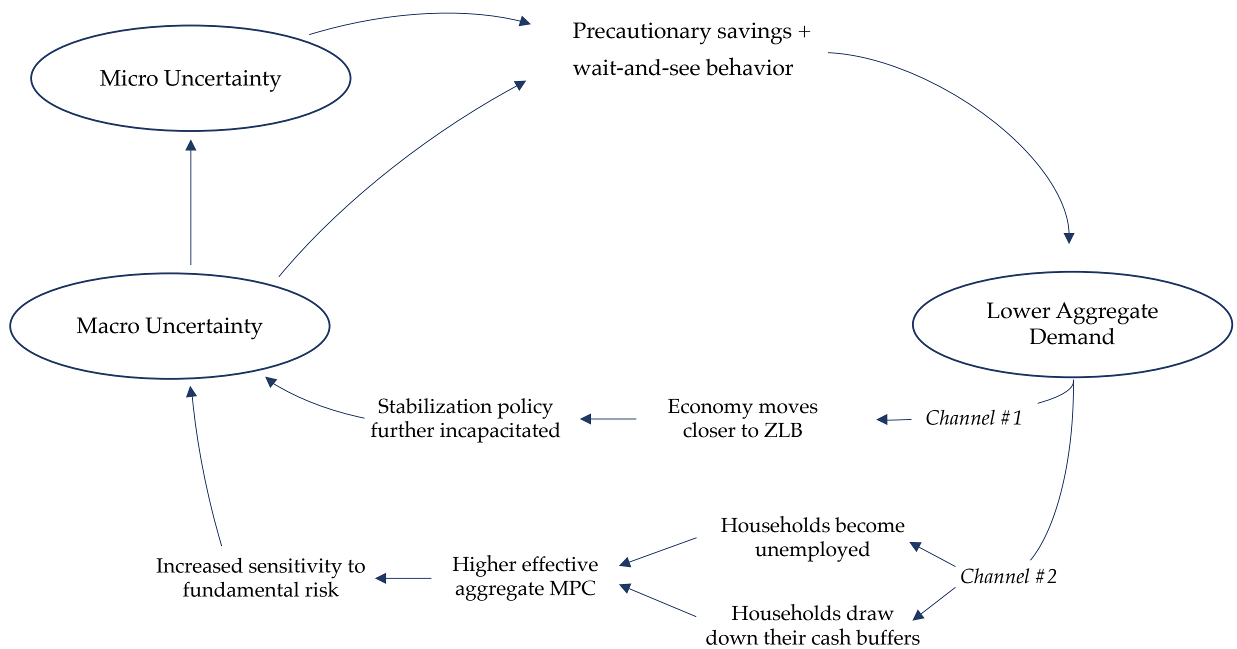

5.3 Uncertainty Multiplier

The discussion thus far highlights that a two-way interaction or feedback loop can emerge in general equilibrium between uncertainty and economic activity. Consistent with the discussion thus far, I illustrate the full feedback loop diagrammatically in Figure 8. Uncertainty thus emerges as both a driver and a byproduct of business cycle fluctuations.

I now define an Uncertainty Multiplier, which is a useful tool to characterize the general equilibrium interaction between micro and macro uncertainty. This multiplier measures how much endogenous amplification there is in macroeconomic uncertainty. In particular, it asks how much macroeconomic uncertainty is generated endogenously for a given change in exogenous fundamental risk. The Uncertainty Multiplier is therefore another lens through which we can study the features and implications of uncertainty. It is given by

$$\underbrace{\sigma_Y(\Gamma_t)}_{\substack{\text{Macro}\\\text{uncertainty}}} = \underbrace{G_Y(\Gamma_t)}_{\substack{\text{GE Uncertainty}\\\text{Multiplier}}} \cdot \underbrace{\sigma_\rho}_{\substack{\text{Exogenous}\\\text{fundamental risk}}} \tag{25}$$

In particular, this decomposition highlights that macro uncertainty is proportional to exogenous fundamental risk. This justifies my approach in Section 5.1 where I studied the fundamental risk shock $\Delta\sigma_\rho$ as a proxy for variation in endogenous macro uncertainty.

In Figure 9, I plot the impulse response of endogenous macro uncertainty to an exogenous fundamental risk shock. Consistent with previous discussion, the Uncertainty Multiplier is large when (a) we account for the interaction between micro and macro uncertainty (compare the two models) and (b) when the economy is in the crisis region (for each model, compare

Figure 9: Impulse response of endogenous macro uncertainty to $\Delta\sigma_\rho$ shock (Uncertainty Multiplier)

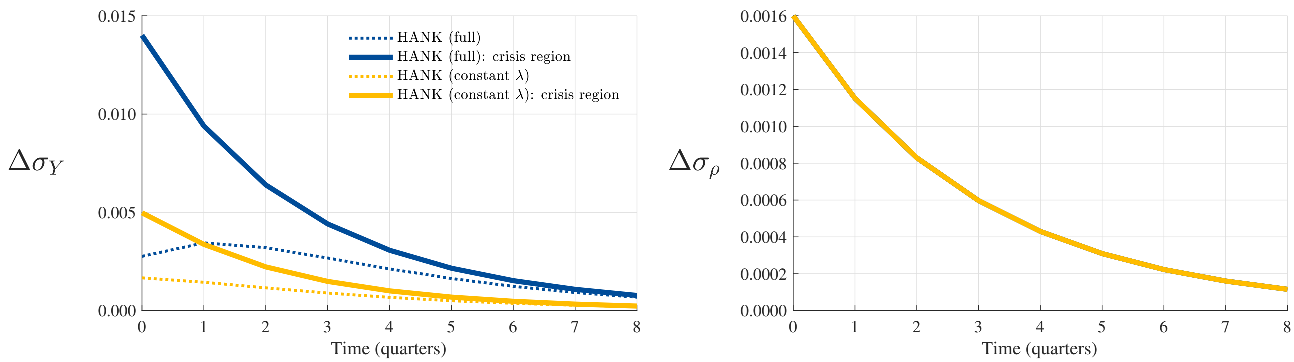

Notes. Comparison of impulse responses of endogenous macro uncertainty $\sigma_{Y,t}$ to fundamental risk shock across different models and state space regions. The shock, displayed in the right panel, has a half-life of 2 quarters and is initialized at half the size of the Basu and Bundick (2017) shock to facilitate comparison. The blue and yellow lines correspond, respectively, to the baseline model and the model that shuts off the interaction between micro and macro uncertainty. The dashed and solid lines initialize the economy, respectively, in the risky steady state (normal times) and on the cusp of the ZLB (crisis region).

the solid to the dashed line). As before, these two forces, non-linearities at the micro and at the macro level, interact.

Figure 10 offers another perspective on the Uncertainty Multiplier. Instead of plotting IRFs, I offer a state space representation of $G_Y(\Gamma_t)$ (y-axis) against different values of the discount rate shock ρ_t (x-axis). In both models, macroeconomic uncertainty increases endogenously as the economy approaches the crisis region (lower ρ_t). However, Figure 10 highlights starkly that the sensitivity of uncertainty to the demand shock is much stronger when we account for the interaction between micro and macro uncertainty in the full model: as we move left in the Figure, the two lines growth increasingly farther apart.

The Uncertainty Multiplier thus offers yet another perspective on the main insight of this paper: there is a two-way feedback loop between macro uncertainty on the one hand, and micro uncertainty and aggregate demand on the other. The model thus generates a natural covariance between uncertainty and economic activity.

5.4 Macro uncertainty in the data

This paper proposes a framework in which both micro and macro uncertainty are an endogenous byproduct of recessions but also have strong aggregate demand effects themselves. In general equilibrium, economic crises are thus associated with large, endogenous uncertainty spikes, especially when taking into account the interplay between uncertainty at the micro and macro level. This conclusion then naturally begs the question whether the uncertainty generated endogenously in my model is consistent with the data. I tackle this question by comparing the time series moments of macro uncertainty in my model simulations to those

Figure 10: State space representation of Uncertainty Multiplier for different values of discount rate ρ

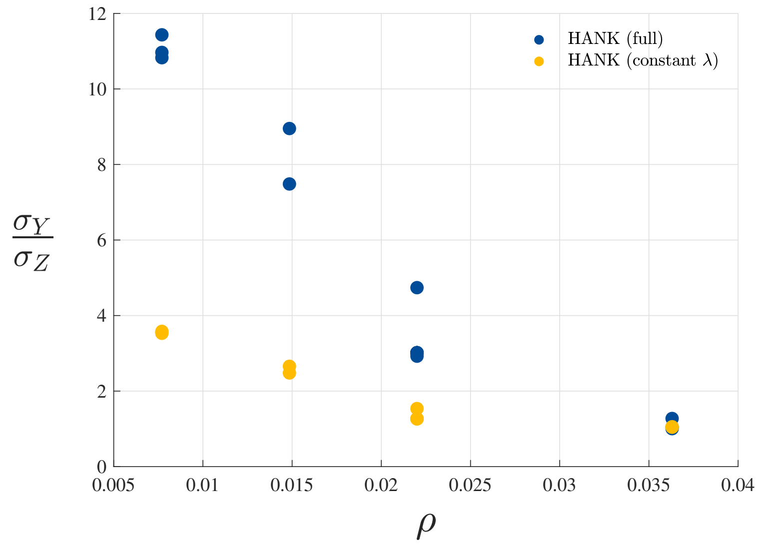

Notes. Comparison of Uncertainty Multiplier $G_Y(\rho,\alpha) = \frac{\sigma_Y(\rho,\alpha)}{\sigma_\rho}$ across different models via state space representation. The x-axis traces out different values in ρ, and the y-axis maps these into $G_Y(\rho,\alpha)$. For a given ρ, vertically aligned dots correspond to different values of α which are added for illustration. The blue and yellow dots correspond, respectively, to the baseline model and the model that shuts off the interaction between micro and macro uncertainty.

of empirical macro uncertainty proxies.

In Table 3, I confront the predictions of my model for the cyclicality, persistence, skewness and kurtosis of macro uncertainty with data. These moments are untargeted in my calibration. Furthermore, these simulations only feature first-moment discount rate shocks, and no exogenous second-moment shocks. Empirical work on uncertainty has proposed a range of indices and proxies to measure macro uncertainty in the data. The first row of Table 3 corresponds to simulated data from my model and subsequent rows correspond to empirical macro uncertainty indices. Along the columns, I compare four moments of uncertainty in my model to the data: its cyclicality (correlation with GDP growth), persistence, skewness and kurtosis.

The model matches the moments of macro uncertainty in the data surprisingly well: Uncertainty is strongly counter-cyclical, highly persistent, and has high skewness and kurtosis. It is especially noteworthy that I match the high skewness and kurtosis in the data, which bespeaks a degree of non-linearity at the macro level and justifies my emphasis on global solution methods.

Result 5. *Recessions are associated with large, endogenous spikes in uncertainty in the model. Indeed, the time series moments of macroeconomic uncertainty in simulations of the model closely match those in the data: macro uncertainty is strongly counter-cyclical, highly persistent, and features large*

Table 3: Time series moments of macro uncertainty: model simulation v. data

Macro uncertainty proxies	Corr w GDP growth	Persistence	Skewness	Kurtosis
Model simulation	-0.64	0.75	1.27	4.04
MacroH1	-0.62	0.92	1.84	7.08
MacroH3	-0.6	0.94	1.85	6.97
MacroH12	-0.6	0.96	1.67	6.49
RealH1	-0.46	0.81	1.21	4.77
FinancialH1	-0.42	0.89	0.88	4.02
VIX	-0.45	0.74	1.94	9.51
Policy Uncertainty	-0.62	0.8	3.13	21.07
World Uncertainty Index	-0.11	0.4	1.58	6.04

Notes. The top row computes time series moments of σ_Y using simulated model data. Subsequent rows report the same time series moments for different empirical macro uncertainty proxies. The indices MacroH*, RealH* and FinancialH* are from Jurado et al. (2015). The underlying data can be retrieved from Sydney Ludvigson's website. The VIX corresponds to the Chicago Board Options Exchange's Volatility Index. The Policy Uncertainty index is taken from Baker et al. (2016) and the World Uncertainty Index is from Ahir et al. (2018).

positive skewness and kurtosis.

6 ZLB Spells and the Paradox of Thrift

Due to the interplay between micro and macro uncertainty, life at the ZLB in my model is even worse than we thought.

6.1 ZLB spells with micro and macro uncertainty

The interplay between micro and macro uncertainty makes falling into the ZLB more likely and life at the ZLB more persistent.

In a HANK model with micro and macro uncertainty, at least three channels lead to a higher incidence of ZLB episodes. First, as illustrated in Section 5, negative demand shocks are substantially amplified as the economy approaches the ZLB. Second, uncertainty endogenously spikes at both the micro and macro level during economic crises, and an increase in uncertainty has itself a contractionary effect on economic activity. Finally, even during normal times, households anticipate that crisis episodes featuring significant micro and macro risk are more likely. As a result, households perceive the world to be riskier overall, inducing an increase in desired savings and a fall in the natural rate of interest in the risky steady state. The economy is therefore closer to the ZLB constraint even during normal times and monetary policy has less room for counter-cyclical stabilization policy.

6.2 Paradox of thrift revisited

> *"[T]he reactions of the amount of [an individual's] consumption on the incomes of others makes it impossible for all individuals simultaneously to save any given sums. Every such attempt to save more by reducing consumption will so affect incomes that the attempt necessarily defeats itself... [I]t makes [no sense] to neglect [this effect] when we come to aggregate demand."* Keynes (1936), p. 84-85.

The paradox of thrift has long been a central force in macroeconomic theories of business cycle fluctuations. It states that a prudent increase in desired savings at the household level can lead to a contraction of aggregate income in general equilibrium that actually lowers total savings. While the paradox of thrift is an important subject in its own right, its interaction with the ZLB will be the central focus of this subsection.[55]

What makes the ZLB so destabilizing in many models is that it induces a strong attraction of "self-fulfilling nature" as the economy approaches it. The paradox of thrift is at the heart of this dynamic. Consider the baseline representative-agent New Keynesian model. As the nominal interest rate falls towards zero, agents anticipate that, in future, monetary policy will no longer be able to offset negative shocks with interest rate cuts. Since the ZLB is asymmetric, however, monetary policy will continue to offset large positive shocks with interest rate hikes. As a result, households' conditional expectation of future real interest rates rises, inducing an increase in desired savings. Since it is not possible for all households to save at the same time, as posited by the paradox of thrift, aggregate income instead falls to offset the rise in aggregate savings demand. This contraction in activity pushes the economy even closer to the ZLB. In any RANK model, this dynamic relies on a strong motive for inter-temporal substitution among households.

The main result of this subsection is to show that, in a HANK economy with micro and macro uncertainty, the general equilibrium attraction exerted by the ZLB is amplified even though the partial equilibrium savings motives across households are dampened. In other words, a different set of channels gives rise to a strong GE interaction between the ZLB and the paradox of thrift, with the main channel operative in a RANK economy substantially dampened.

Figure 11 illustrates the amplification generated as the economy approaches the ZLB. The two panels plot output and the nominal interest rate in a state space representation for varying levels of the discount rate shock. As we move left along the x-axis, lowering the discount rate, aggregate demand and thus output fall. As does the nominal interest rate. By definition, the economy reaches the ZLB when the nominal interest rate falls to 0. In the left panel, I indicate this point, the "cusp of the ZLB", in the state space with a vertical dashed

[55] Much previous work has studied the interaction between the paradox of thrift and the ZLB in the context of representative-agent or two-agent models. See for example Eggertsson and Woodford (2003), Werning (2011) and Eggertsson and Krugman (2012).

Figure 11: Anticipation effects give rise to "paradox of prudence" near cusp of ZLB

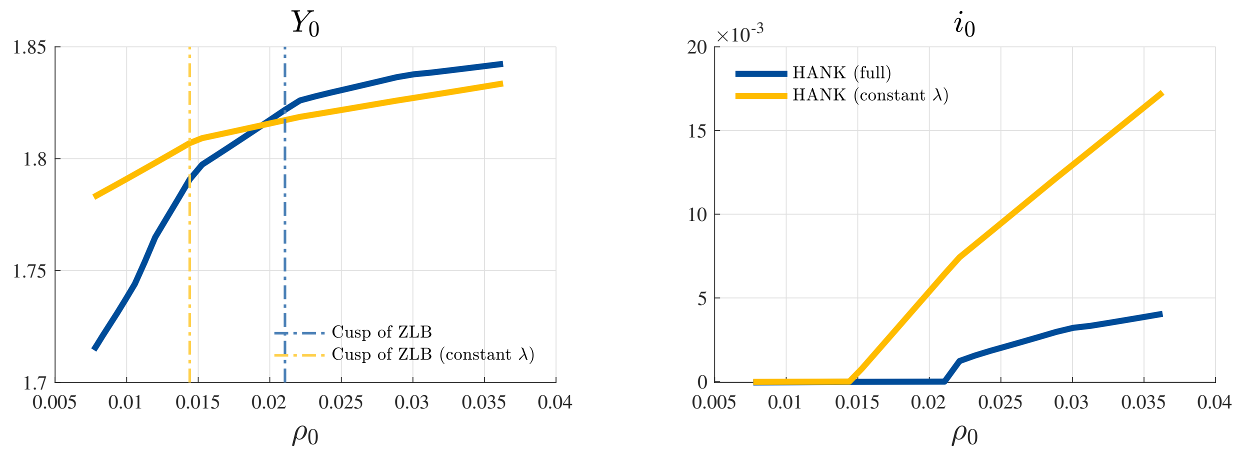

Notes. Panels trace out on-impact response in output and the nominal interest rate (y-axis) for simulations that initialize the economy at different shock realizations ρ_0 (x-axis). Blue and yellow lines correspond, respectively, to the baseline model and the model that shuts off the interaction between micro and macro uncertainty.

line. Notice that the nominal interest rate falls *linearly* with the discount rate. At a certain point, however, when the economy has come sufficiently close to the ZLB, the "self-fulfilling" amplification effects due to the paradox of thrift kick in as discussed above. In particular, these amplification effects are much weaker when we shut off the interaction betwen micro and macro uncertainty (yellow). In the full model (blue), output also falls linearly with the discount rate initially but becomes increasingly non-linear as the economy first approaches and then falls deeper into the ZLB region.

In the language of Kaplan et al. (2018), the set of direct and indirect channels that governs the GE attraction of the ZLB is very different in a HANK model with micro and macro uncertainty. My model not only features a large share of constrained households in normal times; the share of hand-to-mouth households is also strongly counter-cyclical, rising as the economy approaches the ZLB. Liquidity constraints therefore act to dampen the direct effect on desired savings at the micro level because, in the language of Keynes, not all individuals simultaneously try to save.

On the other hand, those households that are on their Euler equations exhibit a stronger rise in desired savings when micro and macro uncertainty interact. Not only do these households anticipate asymmetry in future aggregate states due to the ZLB, as they would in a RANK model, they also internalize the likely spike in unemployment risk. The interplay between micro and macro uncertainty therefore elicits a stronger precautionary motive among unconstrained households.

Finally, as is well known by now, households' consumption and savings decisions in a HANK economy are relatively more influenced by indirect channels in general equilibrium.

Taken together, these factors imply that the interaction between the ZLB and the paradox of thrift is amplified even though not all households increase their desired savings. The overall contraction in aggregate demand posited by the paradox of thrift can be thought of as the product of two effects: a partial equilibrium increase in households' desired savings at the micro level, and a general equilibrium multiplier which maps a given increase in households' precautionary motives into aggregate demand. In a HANK economy with micro and macro uncertainty, the partial equilibrium effect is significantly dampened as many households do not engage in forward-looking consumption smoothing, but the general equilibrium effect is amplified.

7 Cost of Business Cycles

The interplay between micro and macro uncertainty has ramifications for the welfare cost of business cycles. In a seminal contribution, Lucas (1987) made the case that the potential gains from stabilization policy are miniscule as he found the welfare cost of business cycle fluctuations, as perceived by a representative household, to be negligible. An active literature has since spawned in response.

The model developed in this paper offers two novel perspectives on this enduring debate. This paper is the first to emphasize that uncertainty spikes endogenously during crises at both the micro and the macro level. Previous work has focused on one or the other.

Second, business cycles are asymmetric. To my knowledge, this is the first paper to compute the cost of business cycles in a globally solved model with a zero lower bound constraint that endogenously generates negative skewness and positive kurtosis in aggregate output, consumption and investment. Beaudry and Pages (2001) and Krebs (2007) compute cost of recession metrics by assuming that, in the context of a model with two aggregate states, stabilization policy can maintain the economy in the good state indefinitely. This paper generates the asymmetry in business cycles that motivates such reasoning endogenously.

To characterize the welfare cost of business cycle fluctuations in my setting, I implement the following thought hypothetical. Consider households that live in the quantitative model of Section 2. These households would give up 3.9% of their average consumption to live in the RANK model instead (without ZLB).

8 Conclusion

This paper has developed a business cycle model in which micro and macro uncertainty interact. Macro uncertainty shocks are substantially amplified in such a setting, with the indirect effect *through* micro uncertainty emerging as the dominant transmission channel. But macro uncertainty is also endogenous to economic activity, especially during times of crisis. My paper therefore develops a perspective of uncertainty as both a driver and a byproduct of business cycles. I show that accounting for the interaction between uncertainty at the micro and macro level is crucial to understand its broader role in business cycle fluctuations.

References

Abowd, J. M. and A. Zellner (1985). Estimating gross labor-force flows. *Journal of Business & Economic Statistics 3*(3), 254–283.

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

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2015). Heterogeneous agent models in continuous time.

Ahir, H., N. Bloom, and D. Furceri (2018). The world uncertainty index. *Available at SSRN 3275033*.

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017). Micro Heterogeneity and Aggregate Consumption Dynamics. *NBER Macroeconomics Annual, forthcoming*.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Algan, Y., O. Allais, and W. J. Den Haan (2008). Solving heterogeneous-agent models with parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control 32*(3), 875–908.

Algan, Y., O. Allais, W. J. Den Haan, and P. Rendahl (2014). Solving and simulating models with heterogeneous agents and aggregate uncertainty. In *Handbook of computational economics*, Volume 3, pp. 277–324. Elsevier.

Altig, D., L. J. Christiano, M. Eichenbaum, and J. Linde (2011). Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic dynamics 14*(2), 225–247.

Arellano, C., Y. Bai, and P. J. Kehoe (2019). Financial frictions and fluctuations in volatility. *Journal of Political Economy 127*(5), 2049–2103.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2019). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. Technical report, National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model. Technical report, National Bureau of Economic Research.

Bachmann, R., G. Moscarini, and Others (2011). Business cycles and endogenous uncertainty. Citeseer.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The quarterly journal of economics 131*(4), 1593–1636.

Basu, S. and B. Bundick (2017). Uncertainty shocks in a model of effective demand. *Econometrica 85*(3), 937–958.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Beaudry, P. and C. Pages (2001). The cost of business cycles and the stabilization value of unemployment insurance. *European Economic Review 45*(8), 1545–1572.

Berger, D. and J. Vavra (2019). Shocks versus Responsiveness: What Drives Time-Varying Dispersion? *Journal of Political Economy 127*(5), 2104–2142.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The quarterly journal of economics 98*(1), 85–106.

Bilbiie, F. O. (2020). The new Keynesian cross. *Journal of Monetary Economics 114*, 90–108.

Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of political economy 93*(2), 223–247.

Blanchard, O. J., P. Diamond, R. E. Hall, and K. Murphy (1990). The cyclical behavior of the gross flows of US workers. *Brookings Papers on Economic Activity 1990*(2), 85–155.

Bloom, N. (2009). The impact of uncertainty shocks. *econometrica 77*(3), 623–685.

Bloom, N. (2014, may). Fluctuations in Uncertainty. *Journal of Economic Perspectives 28*(2), 153–176.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really uncertain business cycles. *Econometrica 86*(3), 1031–1065.

Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control 89*, 68–92.

Brumm, J. and S. Scheidegger (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica 85*(5), 1575–1612.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *The American Economic Review 104*(2), 379–421.

Bungartz, H.-J. and M. Griebel (2004). Sparse grids. *Acta numerica 13*(1), 147–269.

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions (2015). The master equation and the convergence problem in mean field games. *arXiv preprint arXiv:1509.02505*.

Challe, E., J. Matheron, X. Ragot, and J. F. Rubio-Ramirez (2017). Precautionary saving and aggregate demand. *Quantitative Economics 8*(2), 435–478.

Chodorow-Reich, G. and L. Karabarbounis (2016). The cyclicality of the opportunity cost of employment. *Journal of Political Economy 124*(6), 1563–1618.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*(1), 78–121.

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), 1–45.

Darby, M. R., J. C. Haltiwanger, and M. W. Plant (1985). Unemployment-rate dynamics and persistent unemployment under rational expectations. Technical report, National Bureau of Economic Research.

Darby, M. R., J. C. Haltiwanger, and M. W. Plant (1986). The ins and outs of unemployment: The ins win. Technical report, National Bureau of Economic Research.

Den Haan, W. J. (1996). Heterogeneity, aggregate uncertainty, and the short-term interest rate. *Journal of Business & Economic Statistics 14*(4), 399–411.

Den Haan, W. J. and Others (1997). Solving dynamic models with aggregate shocks and heterogeneous agents. *Macroeconomic dynamics 1*, 355–386.

Dou, W., X. Fang, A. W. Lo, and H. Uhlig (2020). Macro-finance models with nonlinear dynamics. *University of Chicago, Becker Friedman Institute for Economics Working Paper, The Rodney L. White Center Working Papers Series at the Wharton School*.

Duarte, V. (2018). Machine Learning for Continuous-Time Economics. *Available at SSRN 3012602*.

Eggertsson, G. B. and P. Krugman (2012). Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *The Quarterly Journal of Economics 127*(3), 1469–1513.

Eggertsson, G. B. and M. Woodford (2003). Optimal monetary policy in a liquidity trap. Technical report, National Bureau of Economic Research.

Elsby, M. W. L., R. Michaels, and G. Solon (2009, jan). The Ins and Outs of Cyclical Unemployment. *American Economic Journal: Macroeconomics 1*(1), 84–110.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics 46*(2), 281–313.

Fajgelbaum, P. D., E. Schaal, and M. Taschereau-Dumouchel (2017). Uncertainty traps. *The Quarterly Journal of Economics 132*(4), 1641–1692.

Fernández-Villaverde, J., G. Gordon, P. Guerrón-Quintana, and J. F. Rubio-Ramirez (2015). Nonlinear adventures at the zero lower bound. *Journal of Economic Dynamics and Control 57*, 182–204.

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester, and J. Rubio-Ramírez (2015, nov). Fiscal Volatility Shocks and Economic Activity. *American Economic Review 105*(11), 3352–3384.

Fernández-Villaverde, J., S. Hurtado, and G. Nuno (2019). Financial frictions and the wealth distribution. Technical report, National Bureau of Economic Research.

Fujita, S. and G. Ramey (2009). The cyclicality of separation and job finding rates. *International Economic Review 50*(2), 415–430.

Gali, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gertler, M. and P. Karadi (2011). A model of unconventional monetary policy. *Journal of monetary Economics 58*(1), 17–34.

Gertler, M. and N. Kiyotaki (2010). Financial intermediation and credit policy in business cycle analysis. In *Handbook of monetary economics*, Volume 3, pp. 547–599. Elsevier.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2014). Uncertainty, financial frictions, and investment dynamics. Technical report, National Bureau of Economic Research.

Guerrieri, V. and G. Lorenzoni (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics 132*(3), 1427–1467.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *The American Economic Review 103*(2), 732–770.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control 17*(5), 953–969.

Judd, K. L., L. Maliar, S. Maliar, and R. Valero (2014). Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. *Journal of Economic Dynamics and Control 44*, 92–123.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review 108*(3), 697–743.

Kehrig, M. (2015). The cyclical nature of the productivity distribution. *Earlier version: US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*.

Keynes, J. M. (1936). *The general theory of employment, interest, and money*. Springer.

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

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

Krueger, D. and F. Kubler (2004). Computing equilibrium in OLG models with stochastic production. *Journal of Economic Dynamics and Control 28*(7), 1411–1436.

Krusell, P. and A. A. J. Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Lasry, J.-M. and P.-L. Lions (2007). Mean field games. *Japanese journal of mathematics 2*(1), 229–260.

Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Laubach, T. and J. C. Williams (2016). Measuring the natural rate of interest redux. *Business Economics 51*(2), 57–67.

Leduc, S. and Z. Liu (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics 82*, 20–35.

Lions, P.-L. (2011). Mean field games and applications. In *Cours au College de France*. Springer.

Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.

Lucas, R. E. (2003). Macroeconomic priorities. *American economic review 93*(1), 1–14.

Ludvigson, S. C., S. Ma, and S. Ng (2015). Uncertainty and business cycles: exogenous impulse or endogenous response? Technical report, National Bureau of Economic Research.

Marston, S. T., M. Feldstein, and S. H. Hymans (1976). Employment instability and high unemployment rates. *Brookings Papers on Economic Activity 1976*(1), 169–210.

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

McKay, A., E. Nakamura, and J. Steinsson (2016). The power of forward guidance revisited. *American Economic Review 106*(10), 3133–3158.

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

Nakata, T. (2017). Uncertainty at the zero lower bound. *American Economic Journal: Macroeconomics 9*(3), 186–221.

Oh, H. and R. Reis (2012). Targeted transfers and the fiscal response to the great recession. *Journal of Monetary Economics 59*, S50—-S64.

Oksendal, B. (2013). *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media.

Ottonello, P. and T. Winberry (2017). Financial Heterogeneity and the Investment Channel of Monetary Policy.

Pástor, L. and P. Veronesi (2013). Political uncertainty and risk premia. *Journal of financial Economics 110*(3), 520–545.

Patterson, C. (2019). The matching multiplier and the amplification of recessions. In *2019 Meeting Papers*, Volume 95. Society for Economic Dynamics.

Plante, M., A. W. Richter, and N. A. Throckmorton (2018). The zero lower bound and endogenous uncertainty. *The Economic Journal 128*(611), 1730–1757.

Poterba, J. M. and L. H. Summers (1986). Reporting errors and labor market dynamics. *Econometrica: Journal of the Econometric Society*, 1319–1338.

Pröhl, E. (2019). Approximating equilibria with ex-post heterogeneity and aggregate risk. *Swiss Finance Institute Research Paper* (17-63).

Ravn, M. O. and V. Sterk (2016). Macroeconomic fluctuations with HANK & SAM: An analytical approach. *Journal of the European Economic Association*.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control 33*(3), 649–665.

Reiter, M. (2010). Solving the incomplete markets model with aggregate uncertainty by backward induction. *Journal of Economic Dynamics and Control 34*(1), 28–35.

Rotemberg, J. J. (1982). Sticky prices in the United States. *Journal of Political Economy 90*(6), 1187–1211.

Schmidt, L. (2016). Climbing and falling off the ladder: Asset pricing implications of labor market event risk. *Available at SSRN 2471342*.

Schmitt-Grohé, S. and M. Uribe (2005). Optimal fiscal and monetary policy in a medium-scale macroeconomic model. *NBER Macroeconomics Annual 20*, 383–425.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review 95*(1), 25–49.

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

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of political Economy 112*(3), 695–717.

Straub, L. and R. Ulbricht (2017). Endogenous uncertainty and credit crunches. *Available at SSRN 2668078*.

Van Nieuwerburgh, S. and L. Veldkamp (2006). Learning asymmetries in real business cycles. *Journal of monetary Economics 53*(4), 753–772.

Werning, I. (2011). Managing a liquidity trap: Monetary and fiscal policy. Technical report, National Bureau of Economic Research.

Werning, I. (2015). Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research.

Winberry, T. (2020). Lumpy Investment, Business Cycles, and Stimulus Policy.

Zariphopoulou, T. (1994). Consumption-investment models with constraints. *SIAM Journal on Control and Optimization 32*(1), 59–85.

Appendix for:

Micro and Macro Uncertainty

Andreas Schaab

November, 2020

Job Market Paper

Abstract

This appendix contains the proofs, additional material and extensions for my job market paper "Micro and Macro Uncertainty". Appendix Sections A through D feature the proofs and supplemental material for Sections 2 through 5, respectively, of the main text.

This appendix focuses exclusively on the analytical and empirical parts of my paper. A separate Numerical Appendix that can be found here provides additional details on the algorithmic and computational aspects of my paper's methodological contribution.

A Appendix for Section 2: Quantitative Model

The material in this section takes as given that the model admits a finite-dimensional aggregate state space representation. This will of course be true for each element in the sequence of approximating economies specified by my solution method. For additional details, see Section 3 in the main text or Appendix C.

I will formalize this assumption before I proceed. In particular, as in the main text, I denote the aggregate state of the economy at time t by $\Gamma_t \in \mathbb{R}^N$ for some N. In practice, for the baseline model we will have $\Gamma_t = (\rho_t, \boldsymbol{\alpha}_t)$, where $\boldsymbol{\alpha}_t$ parameterizes the finite-dimensional distribution representation of my solution method. However, for this section it will suffice to postulate an arbitrary aggregate state vector Γ_t which follows a (Markov) diffusion so that

$$d\Gamma_t = \mu_\Gamma(\Gamma_t)dt + \sigma_\Gamma(\Gamma_t)dB,$$

where B_t is a one-dimensional standard Brownian motion, representing the sole aggregate risk factor of my model, and

$$\mu_\Gamma : \mathbb{R}^N \to \mathbb{R}^N \quad , \qquad \sigma_\Gamma : \mathbb{R}^N \to \mathbb{R}^N.$$

Formally, Γ_t is a *time-homogeneous Ito diffusion* and I assume that it satisfies the usual regularity conditions (see e.g. Oksendal (2013)). In other words, the coefficients μ_Γ and σ_Γ depend on time t only *through* Γ_t, and the process Γ_t therefore satisfies the Markov property for Ito diffusions.

Before proceeding, I will introduce the notion of *infinitesimal generator* which is a composite functional operator associated with a particular diffusion process. This object and its properties are at the heart of Proposition 3 as I already anticipated in the main text.

Definition 4. The infinitesimal generator of a time-homogeneous Ito diffusion process X_t in $\mathbb{R}^N$ is defined as

$$\mathcal{A}_X f(x) = \lim_{t \to 0} \frac{\mathbb{E}_t\left[f(X_t)\right] - f(x)}{t}$$

where $X_0 = x \in \mathbb{R}^n$ and $f : \mathbb{R}^N \to \mathbb{R}$

In particular, the generator of the diffusion process Γ_t is then given by

$$\mathcal{A}_\Gamma f(x) = \mu_\Gamma(x) D_x f + \frac{1}{2}\sigma_\Gamma^T(x) D_{xx} f \sigma_\Gamma(x)$$

where the ^{T} superscript denotes a transpose. To arrive at this form, the assumption of some regularity conditions on f are necessary. The notation of (partial) derivative operators is such that $D_x f$ denotes the Jacobian of f and $D_{xx} f$ its Hessian matrix.

A.1 Household Problem

In the quantitative model, households consume, save, work, and make portfolio allocation decisions. The problem of a household is to maximize the objective function (1) in the main text subject to the set of budget, borrowing and short-sale constraints, and taking as given (i) all aggregate prices, (ii) the economy's aggregate law of motion, $d\Gamma_t$, and (iii) the idiosyncratic employment transition process.

It is easiest to work with a *recursive* representation of the household problem, which is possible under the assumption of a (time-homogeneous) aggregate state space representation (see above).[56] The appropriate state space on which we can formulate such a recursive representation is, of course, (a, k, z, Γ) where (a, k, z) describes the household's idiosyncratic state given by liquid and illiquid asset holdings, as well as employment status. As in the main text, I use superscript ^{j} notation to refer to a household with employment status z^j, for $j \in \{U, E\}$. In sum, the household consumption *policy function*, which I will characterize below, is then given by $c^j(a, k, \Gamma)$ for a household in employment state j, with asset holdings (a, k), and in the aggregate state Γ.

A.1.1 Derivation of the household budget constraint

In the main text, I work directly with *real* household asset holdings. In this subsection, I start from the nominal household budget constraint and derive its real analog, deflated by the household consumer price index. Recall that, in the *state space* representation of the problem, all aggregate prices are functions of the aggregate state of the economy, so that for example the real interest rate is given by $r_t = r(\Gamma_t)$. I will notationally suppress this dependence on Γ whenever there is no ambiguity. Similarly, I will suppress the dependence of functions on the idiosyncratic state (a, k, z) for ease of exposition.

Nominal consumption outlays of the household (in employment state j) are given by Pc^j, where P denotes the price index of the consumption bundle (CPI). I define a as the stock of "bond" holdings (liquid asset) and, similarly, k as the stock of "capital" holdings (illiquid asset). As such, the household budget constraint is given by an equation for the evolution of *liquid wealth*, Pa. We have

$$d(Pa) = \left[(i + \zeta)Pa + \left(i^k + \frac{P\Pi^Q}{K} \right) k + Pe - Q\iota^j - P\psi(\iota, k) - Pc \right] dt.$$

Several observations are in order. The portfolio returns are expressed in nominal terms here, where i is the nominal interest rate, Pa is the nominal bond value on which the nominal interest rate accrues (and on which annuity payments ζ are made). Similarly, i^k is the nominal rental rate per unit of capital k and $P\Pi^Q$ denotes the aggregate nominal profits of the capital producing sector. Q is the price of capital goods, where ι denotes the rate at which households

[56] I will formally discuss the recursive representation of the household problem when Γ_t is *infinite-dimensional* in Appendix C. The discussion here is understood to refer to an *approximate economy* under a finite-dimensional distribution representation.

purchase capital goods from the capital producer. I assume that the adjustment cost is paid in units of the consumption good, which is why it is valued at P. Finally, Pe denotes nominal household earnings.

CPI price inflation is defined as dP/P. The key question, then, is whether dP is stochastic, in the sense that it loads directly on the aggregate risk factor dB. As I will discuss in more detail in Appendix A.4, dP has no diffusion term as long as we assume price stickiness (under Rotemberg adjustment costs). In this case, I can define CPI inflation as

$$\frac{dP}{P} = \pi dt.$$

Returning to the nominal budget constraint and using the fact that dP has no diffusion term, we have $d(Pa) = Pda + adP$ and, dividing by P and rearranging, I arrive at

$$\frac{da}{dt} = (i - \pi + \zeta)a + \left(\frac{i^k}{P} + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota, k) - c,$$

where $r = i - \pi$ is the real riskfree rate, $r^k = i^k/P$ is the real rental rate, and $q = Q/P$ is the relative price of capital goods.

A.1.2 Household net worth

The definition of household net worth is conceptually complicated in this setting by the presence of adjustment costs. Similarly, due to these adjustment costs, it will generically not be possible to reduce the state space of the household from (a, k) to a single net worth variable. In other words, household behavior generally depends not only on wealth but rather on the portfolio split between liquid and illiquid assets.

Nonetheless, it is useful to consider how household net worth may be specified. I will define net worth n as though there were no adjustment costs, so that $n = a + qk$. In Appendix E, I study a version of the model where there are indeed no adjustment cost at the micro level, so that the household portfolio choice problem can be expressed in terms of the single state variable n, as well as a portfolio share θ of risky (but no longer illiquid) assets. Such setting lends itself straightforwardly to a study of consumption-based asset pricing with heterogeneous households. For further details, see Appendix E.

While such a state space reduction is not possible here, it is still useful to characterize the law of motion of household net worth as defined above. To do so, I specify a diffusion process for the relative price of capital given by

$$\frac{dq}{q} = \mu_Q dt + \sigma_Q dB,$$

where μ_Q and σ_Q are equilibrium objects. Also recall that the household's illiquid asset position evolves non-stochastically as in $dk = [(\zeta - \delta)k + \iota]dt$.[57] Net worth therefore evolves

[57] In the presence of capital quality shocks as in Brunnermeier and Sannikov (2014), for example, dk would instead be a stochastic differential equation, and a covariance term $(dk)(dq)$ would emerge.

as $dn = da + kdq + qdk$, or

$$\frac{1}{dt}dn = (r+\zeta)a + \left(r^k + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota,k) - c + kq\mu_Q + kq\sigma_Q\frac{dB}{dt} + q(\zeta-\delta)k + q\iota.$$

Simplifying and rearranging, we arrive at

$$dn = \Big[(r+\zeta)n + e - \psi(\iota,k) - c\Big]dt + \underbrace{\left(\frac{r^k}{q} + \frac{\Pi^Q}{qK}\right)}_{\text{Dividend yield}} qkdt + \Big[\underbrace{\Big(\mu_Q - \delta - r\Big)dt + \sigma_Q dB}_{\text{Excess capital gains rate}}\Big]qk$$

This derivation highlights that household net worth indeed loads on the aggregate risk factor dB via a capital gains term, and thus follows a stochastic differential equation. On the other hand, the *stocks* of liquid and illiquid assets held by the household, a and k, are not stochastic in this sense. In other words, it is the *value* of these positions that evolves stochastically but not their stocks, whose evolution is controlled by the household.

A.1.3 Household value function and FOCs

This subsection derives the optimality conditions for household behavior. Recognizing that labor supply is set by unions, each household is left with two active choice variables, consumption c and investment ι, i.e. the purchase of the illiquid asset.

The formal statement of this problem for a household that starts in state (a, k, z^j, Γ) at time $t = 0$ is

$$V^j(a,k,z,\Gamma) = \max_{c,h,\iota} \mathbb{E}_0\left[\int_0^\infty e^{-(\rho(\Gamma)+\zeta)t}u(c,h)dt\right]$$

subject to

$$\begin{aligned}
\dot{a} &= \Big(r(\Gamma)+\zeta\Big)a + \left(r^k(\Gamma) + \frac{\Pi^Q(\Gamma)}{K(\Gamma)}\right)k + e^j(a,k,\Gamma) - q(\Gamma)\iota - \psi(\iota,k) - c \\
\dot{k} &= (\zeta-\delta)k + \iota \\
h &= H(\Gamma) \\
dz^j &= (z^{-j} - z^j)\ dN\Big(\lambda^j(a,k,\Gamma)\Big) \\
d\Gamma &= \mu_\Gamma(\Gamma)dt + \sigma_\Gamma(\Gamma)dB \\
a &\geq \underline{a} \quad \text{and} \quad k \geq 0,
\end{aligned}$$

taking as given all aggregate prices (as functions of the aggregate state Γ). The constraint $h = H(\Gamma)$ represents the assumption that all labor supply is set symmetrically by labor unions, and households take their quota of work hours as given. Household earnings are given by

$$e^j(a,k,\Gamma) = e^j(\Gamma) = (1-\tau^{\text{lab}})z^j w(\Gamma)H(\Gamma) + \tau^{\text{lump}}(\Gamma) + \tau^{\text{UI}}(z^j).$$

Finally, as discussed in the main text, $dN(\cdot)$ represents a continuous-time Poisson process that is uncorrelated across households, and which takes as its argument the associated transition rate. In our case, this employment transition rate $\lambda^j(a, k, \Gamma)$ is given as an arbitrary function on the household state space.

Recursive representation. It is easiest to work with a recursive representation of the household problem. In Appendix A.2, I discuss formally how to derive such a recursive representation. It is cast in terms of a partial differential equation (PDE) that represents the continuous-time analog of the discrete-time Bellman equation, called the Hamilton-Jacobi-Bellman (HJB) equation.

The HJB equation associated with the household problem (under the assumption that Γ is finite-dimensional) is given by

$$
\begin{aligned}
(\rho + \zeta) V^j = \max_{c^j, \iota^j} \Bigg\{ & u(c^j, H) + V_k^j \Big[(\zeta - \delta) k + \iota^j \Big] \\
& + V_a^j \Big[(r + \zeta) a + \left(r^k + \frac{\Pi^Q}{K} \right) k + e^j - q\iota^j - \psi(\iota^j, k) - c^j \Big] \Bigg\} \\
+ \lambda^j \Big[V^{-j} - V^j \Big] + \mu_\Gamma V_\Gamma + \frac{1}{2} \sigma_\Gamma^T V_{\Gamma\Gamma} \sigma_\Gamma &
\end{aligned}
$$

where I again suppress the dependence of objects on state variables. I only explicitly account for the dependence on employment status using ^{j} superscripts in order to highlight that, really, this is a system of PDEs, one associated with each j. I use the shorthand notation $V_x^j = D_x V^j$ to denote the partial derivative of V^j with respect to x. Finally, notice that the terms in the last row, associated with the laws of motion for the household's employment status and the aggregate state of the economy, are outside the max operator because they are taken as given by the household.

This HJB equation holds everywhere in the interior of the household state space, i.e. for all $a > \underline{a}$ and $k > 0$. To characterize household behavior along the boundary of the state space, the HJB equation must be paired with a set of *state constraint boundary conditions*. For details, see e.g. Achdou et al. (2015). Intuitively, these constraints never bind in the interior of the state space. In continuous time, therefore, households that are at least an "ϵ" away from the constraints will never hit them in the next, infinitesimally small time step.

The household's optimization problem gives rise to the *policy functions* $c^j(a, k, \Gamma)$ and $\iota^j(a, k, \Gamma)$, which characterize household behavior at every possible point of the state space. I will characterize these policy functions next. In the following, I will frequently make use of the shorthand notation

$$
\begin{aligned}
s^j &= (r + \zeta) a + \left(r^k + \frac{\Pi^Q}{K} \right) k + e^j - q\iota^j - \psi(\iota^j, k) - c^j \\
m^j &= (\zeta - \delta) k + \iota^j,
\end{aligned}
$$

where s^j and m^j respectively correspond to the drift, or expected rate of change, in households' liquid and illiquid asset positions.

For later reference, I want to note that the household HJB equation can equivalently be written in terms of *generators* (see my discussion in the introduction to Appendix A). Denoting $\mathcal{A}^j = \mathcal{A}_{(a,k,z^j)}$, we have

$$(\rho + \zeta)V^j = \max_{c^j, \iota^j} \ u(c^j, H) + \mathcal{A}^j V + \mathcal{A}_\Gamma V,$$

where

$$\mathcal{A}^j f^j(a,k,\Gamma) = s^j f^j_a + m^j f^j_k + \lambda^j (f^{-j} - f^j)$$

$$\mathcal{A}_\Gamma f^j(a,k,\Gamma) = \mu_\Gamma f^j_\Gamma + \frac{1}{2}\sigma_\Gamma^T f^j_{\Gamma\Gamma}\sigma_\Gamma.$$

I will make use of this representation below.[58]

Optimality conditions. The optimality condition for consumption is simply given by

$$u_c(c^j, H(\Gamma)) = V^j_a(a,k,\Gamma).$$

Given the value function, we can invert this equation to solve for the policy function $c^j = c^j(a,k,\Gamma)$. Intuitively, households equate the marginal value of consuming with the marginal value of saving for the future, i.e. the marginal value of liquid wealth V^j_a.

The optimality condition for the household's investment decision is given by

$$V^j_k(a,k,\Gamma) = \Big(q + \psi_{\iota^j}(\iota^j, k)\Big) V^j_a(a,k,\Gamma).$$

When the adjustment cost function ψ is non-differentiable, as is the case with a kinked functional form for example, this FOC must be modified accordingly. Intuitively, households equate the marginal values of liquid and illiquid assets, accounting for the marginal rate of transformation.

A.1.4 Euler equation for diffusion processes

While the FOCs derived in the previous subsection fully characterize optimal household behavior, it is instructive to derive a set of Euler equations. Doing so in continuous time is less straightforward than in discrete time. For illustration, I will start with the case where household earnings follow a standard diffusion process without jumps, given by

$$dz = \mu_z dt + \sigma_z dW,$$

[58] Separating the generators as I do here is feasible because there is no interaction between the diffusion coefficients of the idiosyncratic state variables and the aggregate state variables. Indeed, the idiosyncratic state variables do not load on the aggregate risk factor at all in the baseline model. If there was an interaction, the appropriate generator associated with the vector of household state variables, $\mathcal{A}_{(a,k,z,\Gamma)}$, could not be decomposed into an *idiosyncratic* and an *aggregate* component as I do here.

where W is a standard Brownian motion that is uncorrelated across households. The drift and diffusion coefficients, μ_z and σ_z, are arbitrary functions on the household state space. Finally, when the earnings state variable z is continuous rather than discrete, the HJB equation that characterizes the household problem is no longer a *system* of equations. Similarly, I will drop ^{j} superscripts and denote the policy functions as $c(a,k,z,\Gamma)$ and $\iota(a,k,z,\Gamma)$ since z is now continuous.

Assumption. Crucially, I assume for now that $\mathbb{E}[dWdB]=0$ so that idiosyncratic risk and aggregate risk are uncorrelated. I will relax this assumption in Appendix Section XX.

Plugging the policy functions back into the HJB equation (with max operator), we get a characterization of households' *indirect* value function (without max operator since optimization already took place). This HJB equation for the diffusion-only case is given by

$$(\rho+\zeta)V = u(c,H) + mV_k + sV_a + \mu_z V_z + \frac{\sigma_z^2}{2}V_{zz} + \mu_\Gamma V_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma.$$

We can now derive what is known as the *HJB envelope condition* by differentiating with respect to a (and using the Envelope Theorem). This yields

$$(\rho-r)V_a = mV_{ka} + sV_{aa} + \mu_z V_{za} + \frac{\sigma_z^2}{2}V_{zza} + \mu_\Gamma V_{\Gamma a} + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma a}\sigma_\Gamma.$$

Next, notice that since the state variables, (a,k,z,Γ), follow time-homogeneous Ito diffusion processes, so does any function that takes (a,k,z,Γ) as its only arguments. Thus, applying Ito's lemma to $V_a = V_a(a,k,z,\Gamma)$ yields

$$dV_a = sV_{aa}dt + mV_{ak}dt + \mu_z V_{az}dt + \sigma_z V_{az}dW + \frac{\sigma_z^2}{2}V_{azz}dt + \mu_\Gamma V_{a\Gamma}dt + \sigma_\Gamma V_{a\Gamma}dB + \frac{1}{2}\sigma_\Gamma^T V_{a\Gamma\Gamma}\sigma_\Gamma dt.$$

Putting this equation together with the envelope condition yields

$$dV_a = (\rho-r)V_a dt + \sigma_z V_{az}dW + \sigma_\Gamma V_{a\Gamma}dB.$$

Using the household optimality condition for consumption, we have $du_c = dV_a$ and, in particular,

$$V_{az} = u_{cc}c_z$$
$$V_{a\Gamma} = u_{cc}c_\Gamma.$$

Thus, we arrive at the first of several characterizations of a continuous-time Euler equation for household consumption.

Lemma 5. *A continuous-time Euler equation for households' marginal utility of consumption is given by*

$$\frac{du_c}{u_c} = (\rho-r)dt + \sigma_z\frac{u_{cc}}{u_c}c_z dW + \sigma_\Gamma\frac{u_{cc}}{u_c}c_\Gamma dB.$$

Without taking a stance on the functional form $u(c,h)$, we cannot make much progress. Assuming CRRA preferences, however, with relative risk aversion coefficient γ implies

$$\frac{u_{cc}}{u_c} = -\frac{\gamma}{c}.$$

Similarly, notice that c is simply a function of the stochastic process u_c whose law of motion I derived above. We have $c = f(u_c) = u_c^{-\frac{1}{\gamma}}$, and so

$$\begin{aligned}
dc &= -\frac{1}{\gamma}c^{1+\gamma}\left(du_c\right) + \frac{1}{2}\left(\frac{1}{\gamma}\frac{1+\gamma}{\gamma}c^{1+2\gamma}\right)(du_c)^2 \\
&= -\frac{1}{\gamma}c^{1+\gamma}\left((\rho - r)u_c dt + u_{cc}c_z\sigma_z dW + u_{cc}c_\Gamma\sigma_\Gamma dB\right) + \frac{1}{2}\left(\frac{1}{\gamma}\frac{1+\gamma}{\gamma}c^{1+2\gamma}\right)\left(u_{cc}c_z\sigma_z dW + u_{cc}c_\Gamma\sigma_\Gamma dB\right)^2 \\
&= -\frac{1}{\gamma}c^{1+\gamma}\left((\rho - r)u_c dt + u_{cc}c_z\sigma_z dW + u_{cc}c_\Gamma\sigma_\Gamma dB\right) + \frac{1}{2}\left(\frac{1}{\gamma}\frac{1+\gamma}{\gamma}c^{1+2\gamma}\right)(-\gamma c^{-\gamma-1})^2\left((c_z\sigma_z)^2 + (c_\Gamma\sigma_\Gamma)\right. \\
&= -\frac{1}{\gamma}c^{1+\gamma}\left((\rho - r)u_c dt + u_{cc}c_z\sigma_z dW + u_{cc}c_\Gamma\sigma_\Gamma dB\right) + \frac{1}{2}\left((1+\gamma)c^{1+2\gamma}\right)c^{-2\gamma-2}\left((c_z\sigma_z)^2 + (c_\Gamma\sigma_\Gamma)^2\right) \\
&= \frac{r-\rho}{\gamma}cdt + \frac{1}{\gamma}\gamma c_z\sigma_z dW + c_\Gamma\sigma_\Gamma dB + \frac{1}{2}(1+\gamma)c^{-1}\left((c_z\sigma_z)^2 + (c_\Gamma\sigma_\Gamma)^2\right).
\end{aligned}$$

I summarize in the next Lemma.

Lemma 6. *Assuming CRRA preferences, a continuous-time Euler equation for household consumption is given by*

$$\frac{dc}{c} = \underbrace{\frac{r-\rho}{\gamma}dt}_{\text{Consumption Smoothing}} + \underbrace{\frac{1+\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2}_{\substack{\text{Precautionary Motive:}\\ \text{Earnings Risk}}} dt + \underbrace{\frac{1+\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2}_{\substack{\text{Precautionary Motive:}\\ \text{Aggregate Risk}}} dt + \frac{c_z\sigma_z}{c}dW + \frac{c_\Gamma\sigma_\Gamma}{c}dB$$

This Lemma characterizes a consumption Euler equation in *differential* form. From here, it is possible to derive an exact analog to the discrete-time Euler equation by integrating up over a given time horizon.

Let $Z = \log(c) = f(c)$. Then,

$$\begin{aligned}
dZ &= f'(c)dc + \frac{1}{2}f''(c)(dc)^2 = \frac{dc}{c} - \frac{1}{2}\left(\frac{dc}{c}\right)^2 \\
&= \frac{r-\rho}{\gamma}dt + \frac{c_z\sigma_z}{c}dW + \frac{c_\Gamma\sigma_\Gamma}{c}dB + \frac{1+\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 dt + \frac{1+\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2 dt - \frac{1}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 dt - \frac{1}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2 dt \\
&= \frac{r-\rho}{\gamma}dt + \frac{\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 dt + \frac{\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2 dt + \frac{c_z\sigma_z}{c}dW + \frac{c_\Gamma\sigma_\Gamma}{c}dB
\end{aligned}$$

Then, writing out the stochastic differential equation in integral form, we have

$$Z_t - Z_0 = \int_0^t \left[\frac{r-\rho}{\gamma} + \frac{\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 + \frac{\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2\right]dt + \int_0^t \frac{c_z\sigma_z}{c}dW + \int_0^t \frac{c_\Gamma\sigma_\Gamma}{c}dB$$

$$\log\left(\frac{c_t}{c_0}\right) = \int_0^t \left[\frac{r-\rho}{\gamma} + \frac{\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 + \frac{\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2\right]dt + \int_0^t \frac{c_z\sigma_z}{c}dW + \int_0^t \frac{c_\Gamma\sigma_\Gamma}{c}dB$$

$$c_t = c_0 \, \exp\left\{\int_0^t \left[\frac{r-\rho}{\gamma} + \frac{\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 + \frac{\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2\right]dt + \int_0^t \frac{c_z\sigma_z}{c}dW + \int_0^t \frac{c_\Gamma\sigma_\Gamma}{c}dB\right\}.$$

Lemma 7. *In the limit of small time steps, we have*

$$\mathbb{E}_0\left[c_{t+dt}\right] = \exp\left\{\frac{r-\rho}{\gamma}dt + \frac{1+\gamma}{2}\left(\frac{c_z\sigma_z}{c}\right)^2 dt + \frac{1+\gamma}{2}\left(\frac{c_\Gamma\sigma_\Gamma}{c}\right)^2 dt\right\} c_t,$$

where each of the terms on the RHS is evaluated at time t.

A.2 Formal Derivation of HJB Equation

This subsection contains no economic content and can be skipped accordingly. I will proceed in three steps: (1) I will start by presenting an informal derivation of the household's recursive problem to provide intuition. (2) I will then work through a formal derivation. (3) Finally, I relate the continuous-time recursive problem to its discrete-time analog for the benefit of those readers who are more familiar with discrete-time Bellman equations.

To derive a recursive representation of the household problem, I define the *value* of the household problem at time 0 as

$$V(a,k,z,\Gamma) = \max \mathbb{E}_0\left[\int_0^\infty u(c_t,h_t)dt\right]$$

subject to the constraints listed above, where the maximum is taken over all household control variables. As is standard, I assume that all controls are Markov.[59]

In progress, coming soon.

A.3 Nexus between the HJB and KF Equations

There is a deep link between the Hamilton-Jacobi-Bellman (HJB) equation that characterizes household behavior and the Kolmogorov forward (KF) equation that describes the evolution

[59]Formally, the household problem is defined over a time horizon $t \in [0,T]$. As is standard in macroeconomics, I define the *infinite-horizon* problem as the limit $T \to \infty$, and I suppress reference to T wherever possible to use notation that is familiar to economists.

of the cross-sectional household distribution. (For previous discussions of this link see, for example, Lasry and Lions (2007) or Achdou et al. (2015).)

As in the main text, I denote the cross-sectional household distribution by $g(a,k,z,\Gamma)$. The *marginal* distributions for employed and unemployed households are given by $g^j(a,k,\Gamma)$. It is well known that the law of motion of the distribution g is characterized by a KF equation, which is a particular kind of partial differential equation (PDE).

In this subsection, I discuss the nexus between the HJB and KF equations. It is intuitively quite easy to see why there would be a deep link between these equations. The HJB equation describes the behavior of households across the state space at time t. And it is through their behavior at time t, of course, that households transition to new positions in the state space at time $t+dt$. In other words, the household's decision to save or dissave at time t will determine its liquid asset position at time $t+dt$. In this sense, characterizing how the cross-sectional distribution evolves is simply a matter of aggregating up households' decisions and, thereby, the transition flows of households in the state space.

Mathematically, this relationship turns out to be characterized by *transposition*. The KF equation is given by

$$\frac{dg_t^j(a,k)}{dt} = -\partial_a\Big[s^j(a,k,\Gamma_t)g_t^j(a,k)\Big] - \partial_k\Big[m^j(a,k,\Gamma_t)g_t^j(a,k)\Big] - \lambda^j g_t^j(a,k) + \lambda^{-j} g_t^j(a,k).$$

In the spirit of my earlier discussion of *generators*, I will define the composite operator that represents this partial differential equation as

$$\mathcal{A}_j^* f^j(a,k) = -\partial_a\Big[s^j(a,k,\Gamma_t)f^j(a,k)\Big] - \partial_k\Big[m^j(a,k,\Gamma_t)f^j(a,k)\Big] - \lambda^j f^j(a,k) + \lambda^{-j} f^j(a,k),$$

so that

$$\frac{dg_t^j(a,k)}{dt} = \mathcal{A}_j^* g_t^j(a,k).$$

The following Lemma demonstrates the deep link between the KF and HJB equations.

Lemma 8. *(HJB-KF Nexus) The operator $\mathcal{A}_j^*$ which defines the KF equation of this economy is the adjoint of the idiosyncratic generator $\mathcal{A}_j$ that defines the HJB equation.*

An adjoint is to functional operators what the transpose is to matrices. Indeed, this result is of immense practical value: After discretizing the state space (a,k,z) on a grid, functional operators like ∂_a can be represented numerically by matrices. Then, finding $\mathcal{A}_j^*$ is as simple as transposing the matrix $\mathcal{A}_j$ that must be constructed to solve for the household value function.

A.4 Inflation Risk

Consumer prices are locally deterministic in my model due to price stickiness in the form of Rotemberg adjustment costs. Therefore, the baseline model features no *inflation risk*: CPI inflation at time t is t-adapted, in the sense that households have perfect foresight over the realization P_{t+dt} at time t.

This simple model feature has far-reaching implications. In particular, the evolution of the household's liquid asset position follows a non-stochastic differential equation. Since the illiquid asset position evolves similarly non-stochastically and employment status follows a Poisson rather than a diffusion process, the model's KF equation becomes a standard partial differential equation rather than a stochastic PDE. Indeed, in a model with inflation risk, the process a_t follows a stochastic differential equation and the associated KF equation becomes a stochastic PDE. This would be the case, for example, if we assume wage stickiness but allow for fully flexible consumer prices.

Recent work in the burgeoning HANK literature has become increasingly fond of assuming sticky wages instead of sticky prices (see e.g. Auclert et al. (2018)). In the presence of aggregate risk, such an assumption may considerably complicate the model of household behavior and the associated law of motion of the cross-sectional household distribution.

Assume that consumer prices follow a diffusion process with

$$\frac{dP}{P} = \mu_\pi dt + \sigma_\pi dB,$$

where $\sigma_\pi \neq 0$ represents inflation risk. (In the baseline model of this paper, $\sigma_\pi = 0$ of course and I denote $\pi = \mu_\pi$.) Following the steps presented in Appendix A.1.1, the associated evolution equation for the household's liquid asset position becomes

$$da = \left[(r+\zeta)a + \left(r^k + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota, k) - c\right]dt - a\sigma_\pi dB.$$

where I now define the real interest rate as $r = i - \mu_\pi$. The diffusion coefficient $a\sigma_\pi$ captures the household's exposure to inflation risk. It enters negatively since higher realized inflation implies a depreciation of real liquid wealth.

In the presence of inflation risk, the Kolmogorov forward equation for the evolution of the household cross-sectional distribution becomes a *stochastic* partial differential equation.

Lemma 9. *(KF with inflation risk) In the presence of inflation risk, the cross-sectional household distribution evolves according to*

$$\begin{aligned}\frac{dg_t^j(a,k)}{dt} =& -\partial_a\left[s^j(a,k,\Gamma_t)g_t^j(a,k)\right] - \partial_k\left[m^j(a,k,\Gamma_t)g_t^j(a,k)\right] \\ & - \lambda^j g_t^j(a,k) + \lambda^{-j} g_t^j(a,k) + \partial_a\left[a\sigma_\pi g_t^j(a,k)\right]dB.\end{aligned}$$

This equation is a special case of the general KF equation I derive in Appendix B. The proof can also be found there.

Inflation risk also has considerable implications for household behavior, of course. The HJB equation for household behavior becomes

$$(\rho+\zeta)V^j = \max_{c^j,\iota^j}\left\{u(c^j,H) + V_a^j\left[(r+\zeta)a + \left(r^k + \frac{\Pi^Q}{K}\right)k + e^j - q\iota^j - \psi(\iota^j,k) - c^j\right]\right.$$
$$\left. + \frac{1}{2}V_{aa}^j(a\sigma_\pi)^2 + V_{a\Gamma}^j(a\sigma_\pi\sigma_\Gamma) + V_k^j\left[(\zeta-\delta)k + \iota^j\right]\right\}$$
$$+ \lambda^j\left[V^{-j} - V^j\right] + \mu_\Gamma V_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma.$$

Two new terms emerge as a consequence of inflation risk. The first, proportional to V_{aa}, captures the uncertainty households now face over the evolution of real liquid wealth. The second, proportional to $V_{a\Gamma}$, emerges because inflation, and therefore the household's liquid wealth, covaries with the aggregate state of the economy, Γ, as both load on the aggregate risk factor dB.

A.5 Transfers and Rebates

In progress and coming soon.

A.6 Death Process and Blanchard (1985) Annuity Markets

How do we aggregate in this economy with death rates? I follow Blanchard (1985). Recall that the household budget constraints are given by

$$\dot{a} = \Big(r+\zeta\Big)a + \left(r^k + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota,k) - c \equiv s$$
$$\dot{k} = (\zeta-\delta)k + \iota \equiv m,$$

where, as before, I suppress that the policy functions c, e, ι, s and m take (a,k,z,Γ) as their arguments.

To aggregate in this economy, we aggregate over all household cohorts that are still alive. A cohort born at time 0 has remaining size at time t given by $\zeta e^{-\zeta t}$. This is by definition. Similarly, the size of the population at time t is given by

$$\underbrace{e^{-\zeta t}}_{\text{Surviving } t=0 \text{ cohort}} + \underbrace{\int_0^t \zeta e^{-\zeta(t-s)}ds}_{\text{Surviving subsequent cohorts}} = e^{-\zeta t} + \left[e^{\zeta(s-t)}\right]_0^t = e^{-\zeta t} + e^{\zeta(t-t)} - e^{-\zeta t} = 1.$$

It is important to realize here that the remaining size at time t of a cohort born at time 0 with *initial mass* μ is given by $\mu e^{-\zeta t}$.

I can now address the important question of aggregation. Denote by $c(s,t,a,k,z)$ the consumption at time t of a household cohort born at time s. If all households of a given cohort consumed the exact same amount, then we would have

$$C_t = c(0,t)e^{-\zeta t} + \int_0^t c(s,t)\zeta e^{\zeta(s-t)}ds$$

just like in the original Blanchard (1985) paper. But now we also have within-cohort heterogeneity. Therefore, the correct aggregation is given by

$$C_t = \int c(0,t,a,k,z)e^{-\zeta t}g(0,t,a,k,z)d(a,k,z) + \int \left[\int_0^t c(s,t,a,k,z)\zeta e^{\zeta(s-t)}ds \right] g(s,t,a,k,z)d(a,k,z).$$

Now, of course, I want to rewrite everything exploiting the underlying stationarity. In particular, we know that conditional on (a,k,z), households of different cohorts behave identically. And while $g(0,t)$ is very different from $g(s,t)$, we don't have to worry about this difference because, again, households are identical conditional on the state variables. So for consumption, it's as simple as always with

$$C_t = \int c_t(a,k,z)g_t(a,k,z)d(a,k,z).$$

Now what about aggregate savings and aggregate capital accumulation? Define S_t to be aggregate liquid savings in this economy. It must be that $S_t = 0$ at all t because of the bond market clearing condition. Define A_t to be total liquid (bond) wealth at time t. We have

$$\begin{aligned} A_t &\equiv \int a\left[e^{-\zeta t}g(0,t)\right]d(a,k,z) + \int a\left[\int_0^t \zeta e^{\zeta(s-t)}g(s,t)ds\right]d(a,k,z) \\ &= \int a\left[e^{-\zeta t}g(0,t) + \int_0^t \zeta e^{\zeta(s-t)}g(s,t)ds\right]d(a,k,z). \end{aligned}$$

Differentiating with respect to t, we have

$$\begin{aligned} S_t \equiv \dot{A}_t &= \int a\left[-\zeta e^{-\zeta t}g(0,t) + e^{-\zeta t}dg(0,t) + \zeta g(t,t) + \int_0^t \zeta \frac{\partial}{\partial t}\left(e^{\zeta(s-t)}g(s,t)\right)ds\right]d(a,k,z) \\ &= \int a\left[-\zeta e^{-\zeta t}g(0,t) + e^{-\zeta t}dg(0,t) + \theta g(t,t) + \int_0^t \zeta\left(-\theta e^{\zeta(s-t)}g(s,t) + e^{\zeta(s-t)}\frac{d}{dt}g(s,t)\right)ds\right]d(a,k,z). \end{aligned}$$

Splitting up these terms here, we have

$$\begin{aligned} S_t =& \zeta \int ag(t,t) - \zeta \int a\left[e^{-\zeta t}g(0,t) + \int_0^t \zeta e^{\zeta(s-t)}g(s,t)ds\right] \\ &+ \int a\left[e^{-\zeta t}\frac{d}{dt}g(0,t) + \int_0^t \zeta e^{\zeta(s-t)}\frac{d}{dt}g(s,t)ds\right]. \end{aligned}$$

A couple of observations are in order. The first term is simply the definition of the liquid wealth of newly born agents. By assumption, this is 0 since new household cohorts enter with zero wealth. But in principle it could be anything else (under a different birth assumption). The second term simplifies exactly to $-\zeta A_t$ as is immediately obvious. So what do we do with the third term? Let's swap the integrals so that

$$e^{-\zeta t}\left[\int a\frac{d}{dt}g(0,t)d(a,k,z)\right]+\int_0^t \zeta e^{\zeta(s-t)}\left[\int a\frac{d}{dt}g(s,t)d(a,k,z)\right]ds.$$

The Kolmogorov forward equation still holds for every cohort-specific density function, so that we have

$$e^{-\zeta t}\left[\int a\Big(\mathcal{A}_t^* g(0,t)\Big)d(a,k,z)\right]+\int_0^t \zeta e^{\zeta(s-t)}\left[\int a\Big(\mathcal{A}_t^* g(s,t)\Big)d(a,k,z)\right]ds.$$

Notice that the adjoint operator $\mathcal{A}_t^*$ is t-adapted here because the prices that show up in this operator are evaluated at time t (they don't vary across cohorts). And finally, I use the adjoint relationship to arrive at

$$e^{-\zeta t}\left[\int \Big(\mathcal{A}_t a\Big)g(0,t)d(a,k,z)\right]+\int_0^t \zeta e^{\zeta(s-t)}\left[\int \Big(\mathcal{A}_t a\Big)g(s,t)d(a,k,z)\right]ds,$$

where of course

$$\mathcal{A}a = s\partial_a a + m\partial_k a + \ldots = s$$

where all terms except the first are 0. Therefore, putting it all together, we have

$$S_t = -\zeta A_t + e^{-\zeta t}\left[\int s_t(a,k,z)g(0,t)d(a,k,z)\right]+\int_0^t \zeta e^{\zeta(s-t)}\left[\int s_t(a,k,z)g(s,t)d(a,k,z)\right]ds.$$

Finally, I now assume that we are in a stationary equilibrium. The assumption one makes about the initialization of newly born household cohorts imply that there are differences across cohort-specific densities. But, crucially, all cohorts face the same prices. Therefore, household policy functions, conditional on (a,k,z), no longer depend on time. We have

$$\begin{aligned}
S &= -\zeta A + \int s(a,k,z)\left[e^{-\zeta t}g(0,t,a,k,z)\right]d(a,k,z) + \int s(a,k,z)\left[\int_0^t \zeta e^{\zeta(s-t)}g(s,t,a,k,z)ds\right]d(a,k,z)\\
&= -\zeta A + \int s(a,k,z)\Big[\underbrace{e^{-\zeta t}g(0,t,a,k,z)+\int_0^t \zeta e^{\zeta(s-t)}g(s,t,a,k,z)ds}_{\text{Distribution of all surviving households}}\Big]d(a,k,z)
\end{aligned}$$

And what is the distribution of all surviving households? It is precisely the *stationary distribution* because that is the distribution of all households alive when in the stationary equilibrium. Therefore, finally, I have proven that

$$\begin{aligned}
S &= -\zeta A + \int s(a,k,z)g(a,k,z)d(a,k,z)\\
&= -\zeta A + \int \left[\Big(r+\zeta\Big)a + \left(r^k + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota,k) - c\right]g(a,k,z)d(a,k,z)\\
&= \int \left[ra + \left(r^k + \frac{\Pi^Q}{K}\right)k + e - q\iota - \psi(\iota,k) - c\right]g(a,k,z)d(a,k,z)
\end{aligned}$$

where g is the stationary distribution.

A parallel argument proves that the aggregate drift of the illiquid capital stock is given by

$$\begin{aligned} M &= -\zeta \int k g(a,k,z) d(a,k,z) + \int \Big[(\zeta - \delta) k + \iota \Big] g(a,k,z) d(a,k,z) \\ &= -\zeta K + (\zeta - \delta) K + \int \iota g(a,k,z) d(a,k,z) \\ &= I - \delta K, \end{aligned}$$

where the last line uses the market clearing condition for capital production.

Overall, positive death rates have implications for the aggregation of *flows* in portfolio positions but not their outstanding stocks. That is, we still have

$$\begin{aligned} A &= \int a g(a,k,z) d(a,k,z) \\ K &= \int k g(a,k,z) d(a,k,z). \end{aligned}$$

A.7 Illustration of Walras' Law

It is instructive and insightful to work through Walras' law. In particular, I start with households' budget constraints and will derive the goods market clearing condition.

Aggregating households' liquid asset evolution equation, I get

$$\begin{aligned} \sum_j \int s^j g^j(a,k) d(a,k) =& (r+\zeta) \sum_j \int a g^j(a,k) d(a,k) + \left(r^k + \frac{\Pi^Q}{K} \right) \sum_j \int k g^j(a,k) d(a,k) \\ &+ \sum_j \int e^j g^j(a,k) d(a,k) - q \sum_j \int \iota^j g^j(a,k) d(a,k) \\ &- \sum_j \int \psi(\iota^j, k) g^j(a,k) d(a,k) - \sum_j \int c^j g^j(a,k) d(a,k). \end{aligned}$$

Notice that for any variable $x^j(a,k)$, it is notationally equivalent to write

$$\sum_j \int x^j(a,k) g^j(a,k) d(a,k) = \int x(a,k,z) g(a,k,z) d(a,k,z).$$

Accounting for the death process as discussed in the previous subsection, aggregation yields

$$0 = -rB^G + \left(r^k + \frac{\Pi^Q}{K} \right) K + (1 - \tau^{\text{lab}})(1-U) w H + \tau^{\text{lump}} + \tau^{\text{UI}} U - q I^H - \Psi - C$$

where U is the unemployment rate. Notice that $L = (1-U)H$. Using the government budget constraint to solve out for τ^{lump} yields

$$0 = \left(r^k + \frac{\Pi^Q}{K} \right) K + wL + \Pi - q I^H - \Psi - C - G.$$

Since payoffs on the liquid asset, labor tax revenue and UI insurance represent direct transfers between the government and household sectors, they cancel out in this step.

Next, I substitute in for the profits of the capital and goods producing sectors, Π^Q and Π, respectively. This yields

$$
\begin{aligned}
0 &= r^k K + \left(qI - I - \Phi\left(\frac{I}{K}\right)K\right) + wL + \left(Y - wL - r^k K\right) - qI^H - \Psi - C - G \\
&= -I - \Phi\left(\frac{I}{K}\right)K + Y - \Psi - C - G
\end{aligned}
$$

where the second line also uses the market clearing condition for new capital, $I = I^H$. I abuse notation slightly and denote the aggregate adjustment cost paid by the capital producer simply by Φ. Rearranging yields

$$Y = C + I + \Phi + \Psi + G.$$

B Appendix for Section 3: Solution Method

This appendix provides additional details for as well as the proofs of the *analytical* results discussed in Section 3 of the main text. The numerical implementation of the solution method is discussed in a separate Numerical Appendix that can be found here.

B.1 Derivation of Stochastic KF Equation (Diffusion Processes)

The proof of Proposition 3 makes use of the Kolmogorov forward (KF) equation that characterizes the evolution of the cross-sectional household distribution. In this subsection, I prove the following auxiliary result. I adopt the notation introduced in Section 3.1 of the main text.

Lemma 10. *Assume that $\boldsymbol{x}_t$ follows a time-homogeneous Ito diffusion process given by*

$$d\boldsymbol{x} = \mu_{\boldsymbol{x}} dt + \sigma_{\boldsymbol{x}}^W dW + \sigma_{\boldsymbol{x}}^B dB$$

where $\mu_{\boldsymbol{x}}$, $\sigma_{\boldsymbol{x}}^W$ and $\sigma_{\boldsymbol{x}}^B$ are functions on the household state space, $(\boldsymbol{x}, \Gamma)$. W_t is a standard Brownian motion that is uncorrelated across households (idiosyncratic risk), and B_t is a standard Brownian motion that is perfectly correlated across households (aggregate risk). I assume that $\mathbb{E}[dWdB] = 0$. The law of motion of the cross-sectional household distribution $g_t(\boldsymbol{x})$ is given by the stochastic partial differential (Kolmogorov forward) equation

$$dg(t, \boldsymbol{x}) = (\mathcal{A}^* g)dt + (\mathcal{B}^* g)dB$$

where $\mathcal{A}^*$ is the adjoint of the generator $\mathcal{A}$ associated with the HJB equation and given by

$$(\mathcal{A}f)(\boldsymbol{x}) = f_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} + \frac{1}{2}(\sigma_{\boldsymbol{x}}^W)^T f_{\boldsymbol{x}\boldsymbol{x}}(\sigma_{\boldsymbol{x}}^W) + \frac{1}{2}(\sigma_{\boldsymbol{x}}^B)^T f_{\boldsymbol{x}\boldsymbol{x}}(\sigma_{\boldsymbol{x}}^B),$$

and

$$(\mathcal{B}^* f)(\boldsymbol{x}) = -\sum_i \partial_{x_i}\left[\sigma_{x_i}^B f\right].$$

Proof. We may equivalently write

$$d\boldsymbol{x} = \mu_{\boldsymbol{x}} dt + \sigma_{\boldsymbol{x}} \begin{pmatrix} dB \\ dW \end{pmatrix},$$

where

$$\sigma_{\boldsymbol{x}} = \begin{pmatrix} \sigma_{\boldsymbol{x}}^B & \sigma_{\boldsymbol{x}}^W \end{pmatrix}.$$

Consider any function $\phi(t, \boldsymbol{x})$. Then by Ito's lemma,

$$d\phi = \phi_t dt + \phi_{\boldsymbol{x}}^T d\boldsymbol{x} + \frac{1}{2}(d\boldsymbol{x})^T \phi_{\boldsymbol{x}\boldsymbol{x}}(d\boldsymbol{x}),$$

where $\phi_{\boldsymbol{x}}$ denotes the gradient of ϕ with respect to $\boldsymbol{x}$, and $\phi_{\boldsymbol{xx}}$ its Hessian. We may rewrite this as

$$d\phi = \phi_t dt + \phi_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} dt + \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}} d\mathbf{Z} + \frac{1}{2}\text{Trace}\left[\sigma_{\boldsymbol{x}}^T \phi_{\boldsymbol{xx}} \sigma_{\boldsymbol{x}}\right],$$

where I denote

$$d\mathbf{Z} = \begin{pmatrix} dB \\ dW \end{pmatrix}.$$

I now use the assumption $\mathbb{E}[dWdB] = 0$. This allows me to rewrite the above equation as

$$d\phi = \phi_t dt + \phi_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} dt + \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^W dW + \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B dB + \frac{1}{2}(\sigma_{\boldsymbol{x}}^W)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^W) dt + \frac{1}{2}(\sigma_{\boldsymbol{x}}^B)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^B) dt$$

To illustrate this step, imagine $\boldsymbol{x}$ was 2-dimensional. Then the term in question is

$$\begin{aligned}
\frac{1}{2}\text{Trace}\left[\sigma_{\boldsymbol{x}}^T \phi_{\boldsymbol{xx}} \sigma_{\boldsymbol{x}}\right] &= \frac{1}{2}\text{Trace}\left[\begin{pmatrix} \sigma_{x_1} & \sigma_{x_2} \end{pmatrix} \begin{pmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{pmatrix} \begin{pmatrix} \sigma_{x_1} \\ \sigma_{x_2} \end{pmatrix}\right] \\
&= \frac{1}{2}\text{Trace}\left[\begin{pmatrix} \sigma_{x1}\phi_{11} + \sigma_{x_2}\phi_{21} & \sigma_{x_1}\phi_{12} + \sigma_{x_2}\phi_{22} \end{pmatrix} \begin{pmatrix} \sigma_{x_1} \\ \sigma_{x_2} \end{pmatrix}\right] \\
&= \frac{1}{2}\text{Trace}\left[\sigma_{x1}\phi_{11}\sigma_{x_1} + \sigma_{x_2}\phi_{21}\sigma_{x_1} + \sigma_{x_1}\phi_{12}\sigma_{x_2} + \sigma_{x_2}\phi_{22}\sigma_{x_2}\right].
\end{aligned}$$

In integral representation, this becomes,

$$\begin{aligned}
\phi(T, \boldsymbol{x}_T) =& \phi(\tau, \boldsymbol{x}_\tau) + \int_\tau^T \left[\phi_t + \phi_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} + \frac{1}{2}(\sigma_{\boldsymbol{x}}^W)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^W) + \frac{1}{2}(\sigma_{\boldsymbol{x}}^B)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^B)\right] dt \\
&+ \int_\tau^T \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B dB_t + \int_\tau^T \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^W dW_t.
\end{aligned}$$

Consider the initial condition $\boldsymbol{x}_\tau = \bar{\boldsymbol{x}}$. Since I picked $\phi(t, \boldsymbol{x})$ arbitrarily, it is without loss to assume that $\phi(\tau, \bar{\boldsymbol{x}}) = 0$ and $\phi(t, \boldsymbol{x}) \to 0$ uniformly in $\boldsymbol{x}$ as $t \to T$. Then we have

$$\begin{aligned}
&\int \left[\phi(T, \boldsymbol{x}_T) - \phi(\tau, \boldsymbol{x}_\tau)\right] g_t(\boldsymbol{x}) d\boldsymbol{x} \\
&= \int_\tau^T \int \left[\phi_t + \phi_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} + \frac{1}{2}(\sigma_{\boldsymbol{x}}^W)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^W) + \frac{1}{2}(\sigma_{\boldsymbol{x}}^B)^T \phi_{\boldsymbol{xx}} (\sigma_{\boldsymbol{x}}^B)\right] g_t(\boldsymbol{x}) d\boldsymbol{x} dt \\
&\quad + \int_\tau^T \int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B g_t(\boldsymbol{x}) d\boldsymbol{x} dB_t + \int_\tau^T \int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^W g_t(\boldsymbol{x}) d\boldsymbol{x} dW_t,
\end{aligned}$$

where, slightly abusing notation, the integral is taken over the state space of $\boldsymbol{x}$. The key realization now is that, because dW_t is perfectly uncorrelated, we have

$$\int_\tau^T \int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^W g_t(\boldsymbol{x}) d\boldsymbol{x} dW_t = \mathbb{E}_t\left[\phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^W dW\right] = 0.$$

This follows because a law of large numbers holds with respect to idiosyncratic risk (but not with respect to aggregate risk).

I will now introduce an operator given by

$$(\mathcal{A}_t f)(\boldsymbol{x}) = f_{\boldsymbol{x}}^T \mu_{\boldsymbol{x}} + \frac{1}{2}(\sigma_{\boldsymbol{x}}^W)^T f_{\boldsymbol{x}\boldsymbol{x}}(\sigma_{\boldsymbol{x}}^W) + \frac{1}{2}(\sigma_{\boldsymbol{x}}^B)^T f_{\boldsymbol{x}\boldsymbol{x}}(\sigma_{\boldsymbol{x}}^B).$$

Therefore,

$$0 = \int_\tau^T \int \Big[\phi_t + (\mathcal{A}_t\phi)(\boldsymbol{x})\Big] g_t(\boldsymbol{x}) d\boldsymbol{x} dt \int_\tau^T \int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B g_t(\boldsymbol{x}) d\boldsymbol{x} dB_t$$

I will now drop time subscript notation. Going forward, every subscript will denote a partial derivative. And recall, the second integral is taken over the region of support for $\boldsymbol{x}$. Integrating by parts on the first term, I obtain

$$\int_\tau^T \int \phi_t g(t, \boldsymbol{x}) d\boldsymbol{x} dt = \int \Big[\phi g\Big]_\tau^T d\boldsymbol{x} - \int \int_\tau^T \phi g_t dt d\boldsymbol{x},$$

where now $g_t = \partial_t g(t, \boldsymbol{x})$. Since I constructed ϕ so that $\phi(\tau) = \phi(T) = 0$, the first term here drops out. Integrating over $\mathcal{A}$ by parts, and restricting attention to the class of functions ϕ with compact support, I obtain

$$\begin{aligned}
&\int_\tau^T \int (\mathcal{A}_t\phi)(\boldsymbol{x}) g_t(\boldsymbol{x}) d\boldsymbol{x} dt \\
&= -\int_t^T \int \phi \Big[\sum_i \partial_{x_i}(\mu_{x_i} g) - \frac{1}{2}\sum_{i,j} \partial_{x_i x_j}\Big(\sigma_{x_i}^W g \sigma_{x_j}^W + \sigma_{x_i}^B g \sigma_{x_j}^B\Big)\Big] d\boldsymbol{x} dt \\
&= \int_t^T \int \phi(\mathcal{A}^* g) d\boldsymbol{x} dt,
\end{aligned}$$

where $\mathcal{A}^*$ is the adjoint of $\mathcal{A}$.

What remains, therefore, is the aggregate risk term. Here, we have

$$\int_\tau^T \Big[\int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B g d\boldsymbol{x}\Big] dB = \int_\tau^T \Big[\phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B g\Big|_{\text{Boundary}} - \int \phi \sum_i \partial_{x_i}\Big(\sigma_{x_i}^B g\Big) d\boldsymbol{x}\Big] dB.$$

The first term on the RHS is again 0 due to our choice of ϕ. Therefore,

$$\int_\tau^T \Big[\int \phi_{\boldsymbol{x}}^T \sigma_{\boldsymbol{x}}^B g d\boldsymbol{x}\Big] dB = -\int_\tau^T \int \phi \sum_i \partial_{x_i}[\sigma_{x_i}^B g] d\boldsymbol{x} dB$$

In conclusion, I therefore obtain

$$0 = \int_\tau^T \int \phi\Big[g_t - (\mathcal{A}^* g)\Big] d\boldsymbol{x} dt + \int_\tau^T \int \phi\Big[\sum_i \partial_{x_i}[\sigma_{x_i}^B g]\Big] d\boldsymbol{x} dB.$$

Since this equation holds for arbitrary test functions $\phi(t, \boldsymbol{x})$, it must be that

$$0 = \int_{\tau}^{T} \phi \Big[g_t - (\mathcal{A}^* g) \Big] dt + \int_{\tau}^{T} \phi \Big[\sum_i \partial_{x_i} [\sigma^B_{x_i} g] \Big] dB \tag{26}$$

The analogous differential formulation is

$$dg(t, \boldsymbol{x}) = (\mathcal{A}^* g) dt - \sum_i \partial_{x_i} \Big[\sigma^B_{x_i} g \Big] dB.$$

This concludes the proof.

■

B.2 Derivation of Stochastic KF Equation (Jump Processes)

In progress, coming soon.

B.3 Proof of Proposition 3 (Diffusion Processes)

Proposition 3 applies to finite-dimensional approximations of the cross-sectional distribution that take the form

$$\hat{g}_t(\boldsymbol{x}) \equiv F(\boldsymbol{\alpha}_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x}),$$

for some $F(\cdot)$ and $\boldsymbol{\alpha}_t \in \mathbb{R}^N$. I conjecture and verify that, as long as $\boldsymbol{x}_t$ and $\boldsymbol{X}_t$ are time-homogeneous Ito diffusion processes, $\boldsymbol{\alpha}_t$ is similarly an N-dimensional Ito diffusion process, with

$$d\boldsymbol{\alpha}_t = \mu_{\boldsymbol{\alpha}} dt + \sigma_{\boldsymbol{\alpha}} dB,$$

where B_t is the standard, one-dimensional Brownian motion that represents the sole aggregate risk factor of the model.

By Ito's lemma,

$$\begin{aligned} d\hat{g}_t(\boldsymbol{x}) &= F_{\boldsymbol{\alpha}}^T d\boldsymbol{\alpha}_t + \frac{1}{2} \sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} dt \\ &= F_{\boldsymbol{\alpha}}^T \Big(\mu_{\boldsymbol{\alpha}} dt + \sigma_{\boldsymbol{\alpha}} dB \Big) + \frac{1}{2} \sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} dt, \end{aligned}$$

where $F_{\boldsymbol{\alpha}}$ and $F_{\boldsymbol{\alpha}\boldsymbol{\alpha}}$ are the gradient and Hessian of F with respect to $\boldsymbol{\alpha}$, respectively, evaluated at $\boldsymbol{\alpha}_t$ and $\boldsymbol{x}$. Using the approximation $d\hat{g}_t(\boldsymbol{x}) \approx dg_t(\boldsymbol{x})$ as well as the Kolmogorov forward equation which I derived in Appendix B.1, we have

$$\begin{aligned} F_{\boldsymbol{\alpha}}^T \mu_{\boldsymbol{\alpha}} dt + \frac{1}{2} \sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} dt + F_{\boldsymbol{\alpha}}^T \sigma_{\boldsymbol{\alpha}} dB &\approx dg_t(\boldsymbol{x}) \\ &= (\mathcal{A}^* g) dt + (\mathcal{B}^* g) dB. \end{aligned}$$

Matching coefficients, we have

$$F_{\boldsymbol{\alpha}}^T \sigma_{\boldsymbol{\alpha}} = (\mathcal{B}^* g)$$

$$F_{\boldsymbol{\alpha}}^T \mu_{\boldsymbol{\alpha}} + \frac{1}{2}\sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} = (\mathcal{A}^* g).$$

We can again make use of the approximation $F(\boldsymbol{\alpha}_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x})$ to finally arrive at

$$F_{\boldsymbol{\alpha}}^T \sigma_{\boldsymbol{\alpha}} = (\mathcal{B}^* F)$$

$$F_{\boldsymbol{\alpha}}^T \mu_{\boldsymbol{\alpha}} + \frac{1}{2}\sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} = (\mathcal{A}^* F)$$

where I suppress all function inputs for convenience.

To conclude the proof of Proposition 3, it remains to solve for the $N \times 1$ vectors $\mu_{\boldsymbol{\alpha}}$ and $\sigma_{\boldsymbol{\alpha}}$. In terms of the underlying economics, this step represents an estimation or forecasting problem from the perspective of the economy's optimizing agents. There are many equally valid approaches to this estimation problem, each corresponding to a different choice of norm. In the main text, I present Proposition 3 for a *least squares* norm. Here, I derive and discuss several alternatives.

Least squares. We seek to minimize

$$\mathbb{E}_0 \left[\sum_{t=0}^{T} \left(\int \left[F(\boldsymbol{\alpha}_t)(\boldsymbol{x}) - g_t(\boldsymbol{x}) \right]^2 d\boldsymbol{x} \right)^{\frac{1}{2}} \right].$$

The derivation is analogous to that of the standard least squares estimator. In particular, the estimates that minimize the $\mathbb{L}^2(\boldsymbol{x})$ norm are given by

$$\sigma_{\boldsymbol{\alpha}} = (F_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}})^{-1} F_{\boldsymbol{\alpha}}^T (\mathcal{B}^* F)$$

$$\mu_{\boldsymbol{\alpha}} = (F_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}})^{-1} F_{\boldsymbol{\alpha}}^T \left[(\mathcal{A}^* F) - \frac{1}{2}\sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} \right].$$

Collocation. Let $\mathcal{S}$ denote a set of collocation nodes denoted $\boldsymbol{s}$. If we want to pick $F(\boldsymbol{\alpha}_t)$ so that $F(\boldsymbol{\alpha}_t)(\boldsymbol{s}) = g_t(\boldsymbol{s})$ with *equality* at all collocation nodes $\boldsymbol{s}$, then we need as many $\boldsymbol{\alpha}_t$ as there are collocation nodes. I denote this number by N to remain consistent with the previous discussion, so that $\boldsymbol{\alpha}_t \in \mathbb{R}^N$.

At these collocation nodes, we thus have

$$\left(F_{\boldsymbol{\alpha}}^T \mu_{\boldsymbol{\alpha}} + \frac{1}{2}\sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}} \right) dt + F_{\boldsymbol{\alpha}}^T \sigma_{\boldsymbol{\alpha}} dB = (\mathcal{A}^* g) dt + (\mathcal{B}^* g) dB$$

with equality as well. Matching the diffusion coefficients leads to

$$\underbrace{F_{\boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_t)(\boldsymbol{s})}_{\text{N} \times \text{N}} \underbrace{\sigma_{\boldsymbol{\alpha},t}}_{\text{N} \times 1} = \underbrace{(\mathcal{B}_t^* g_t)(\boldsymbol{s})}_{\text{N} \times 1}$$

where I explicitly account for the function arguments to highlight that this is a system of N equations *at* the N collocation nodes. And since $F(\boldsymbol{\alpha}_t)(\boldsymbol{s}) = g_t(\boldsymbol{s})$ exactly at the collocation nodes, we can invert and simply rewrite this as

$$\underbrace{\sigma_{\boldsymbol{\alpha}}}_{\text{N} \times 1} = \underbrace{(F_{\boldsymbol{\alpha}}^T)^{-1}}_{\text{N} \times \text{N}} \underbrace{(\mathcal{B}^* F)}_{\text{N} \times 1}$$

and similarly

$$\underbrace{\mu_{\boldsymbol{\alpha}}}_{\text{N} \times 1} = \underbrace{(F_{\boldsymbol{\alpha}}^T)^{-1}}_{\text{N} \times \text{N}} \left(\underbrace{(\mathcal{A}^* F) - \frac{1}{2} \sigma_{\boldsymbol{\alpha}}^T F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} \sigma_{\boldsymbol{\alpha}}}_{\text{N} \times 1} \right),$$

where the RHS of each equation is evaluated *at* the N collocation nodes $\boldsymbol{s}$.

Example. It is easiest to illustrate the previous arguments by assuming $F(\boldsymbol{\alpha}_t)(\boldsymbol{x})$ takes the form of a linear basis function expansion. That is, consider the approximation

$$T(\boldsymbol{x})\boldsymbol{\alpha}(t) \equiv \sum_{n=1}^{N} \alpha^n(t) T^n(\boldsymbol{x}) \approx g(t, \boldsymbol{x}),$$

where N is the number of basis functions, and the α^n are the basis function coefficients. Any linear basis function can be chosen for T^n, such as Chebyshev polynomials, hat functions, splines, etc.

For these basis functions, the second-order Ito's term conveniently drops out. That is, relating this example back to the general setting, we have

$$F_{\boldsymbol{\alpha}} = \begin{pmatrix} T^1(\boldsymbol{x}) \\ \vdots \\ T^N(\boldsymbol{x}) \end{pmatrix}$$

and

$$F_{\boldsymbol{\alpha}\boldsymbol{\alpha}} = 0.$$

I start with the collocation approach first because it is particularly straightforward in this setting. Denoting the N collocation nodes by $s_1, \ldots, s^N$, we have the system of equations

$$\sum_{n=1}^{N} \alpha^n(t) T^n(s_1) = g(t, s_1)$$

$$\vdots$$

$$\sum_{n=1}^{N} \alpha^n(t) T^n(s_N) = g(t, s_N)$$

Taking the time differential yields

$$\sum_{n=1}^{N} d\alpha^n(t)T^n(s_1) = dg(t, s_1) = (\mathcal{A}^* g)(s_1)dt + (\mathcal{B}^* g)(s_1)dB$$
$$\vdots$$
$$\sum_{n=1}^{N} d\alpha^n(t)T^n(s_N) = dg(t, s_N) = (\mathcal{A}^* g)(s_1)dt + (\mathcal{B}^* g)(s_N)dB$$

Inverting this system, we have

$$d\boldsymbol{\alpha} = T(\boldsymbol{s})^{-1}\Big[(\mathcal{A}^* T)(\boldsymbol{s})\ \boldsymbol{\alpha}\ dt + (\mathcal{B}^* T)(\boldsymbol{s})\ \boldsymbol{\alpha}\ dB\Big].$$

Matching coefficients then yields

$$\mu_{\boldsymbol{\alpha}} = T(\boldsymbol{s})^{-1}(\mathcal{A}^* T)(\boldsymbol{s})\ \boldsymbol{\alpha}$$
$$\sigma_{\boldsymbol{\alpha}} = T(\boldsymbol{s})^{-1}(\mathcal{B}^* T)(\boldsymbol{s})\ \boldsymbol{\alpha}.$$

The second-order Ito term has dropped out as anticipated. In this form, it is particularly easy to see that the formulas for $\mu_{\boldsymbol{\alpha}}$ and $\sigma_{\boldsymbol{\alpha}}$ only depend on terms that are readily available during the value function iteration step. In particular, there are three types of objects: First, the basis function matrix $T(\boldsymbol{x})$, which must be evaluated at $\boldsymbol{s}$, is chosen ex ante. Therefore, both $T(\boldsymbol{s})$ and $T(\boldsymbol{s})^{-1}$ can easily be computed before even starting the value function iteration. Second, the vector $\boldsymbol{\alpha}$ is part of the aggregate state of the approximate economy and, therefore, part of the grid on which the value function is computed. Finally, we have the functional operators $\mathcal{A}^*$ and $\mathcal{B}^*$. When the household's state space is discretized on a grid, these operators become *matrices*. And as I have repeatedly demonstrated throughout this paper, both matrices *only depend* on the household's policy functions.

B.4 Proof of Proposition 3 (Jump Processes)

In progress, coming soon.

B.5 Choosing $F(\cdot)$ from a parametric family

I will now present several parametric families that work well in practice. For simplicity, I will present the functional forms when $\boldsymbol{x} = x$ is one-dimensional. In most cases, the higher-dimensional generalizations are straightforward.

Example. (*Nodal basis functions*) The most basic class of basis functions that can be employed to great effect are nodal basis functions.[60] Let

$$F^n(\alpha_t)(x) = \sum_{i=1}^{n} \alpha_{i,t} T_i(x),$$

where $T_i(x)$ is the (possibly asymmetric) hat function which takes on the value 1 on the grid point on which it is centered, and 0 on all other *nodes*. For off-node grid points, interpolation is used.

Example. (*Chebyshev polynomials*) The representation with Chebyshev polynomials takes on the same form as the nodal basis representation, except that $T_i(x)$ is now the ith Chebyshev polynomial.

Example. (*Radial basis functions*) Let $\alpha_{i,t} = \{\gamma_{i,t}, \mu_{i,t}, \sigma_{i,t}\}$, then

$$F^n(\alpha_t)(x) = \sum_{i=1}^{n} \gamma_{i,t} \exp\left(- \frac{(x - \mu_{i,t})^2}{2(\sigma_{i,t})^2} \right).$$

Example. (*Generalized beta density functions*) Let $\mathrm{B}(\alpha, \beta)$ denote the beta function. Then the standard beta distribution PDF can be used with $\alpha_t \in \mathbb{R}^3$,

$$F(\alpha_t)(x) = \alpha_{1,t} \frac{x^{\alpha_{2,t}-1}(1-x)^{\alpha_{3,t}-1}}{\mathrm{B}(\alpha_{2,t}, \alpha_{3,t})}.$$

Other functional forms from the five-parameter generalized beta family can be used as well. The beta family can be particularly efficient when modeling the household income distribution.

B.6 Non-parametric algorithm

I will present the main argument for the special case where $F(\alpha_t)(\boldsymbol{x})$ is affine. Abusing notation slightly, the approximation mapping can be rendered as $g_t(\boldsymbol{x}) \approx C(\boldsymbol{x}) + \alpha_t F(\boldsymbol{x})$. Assuming we have simulated functional data for g_t at hand, the estimation problem may then be set up as

$$\min_{F(\boldsymbol{x}), C(\boldsymbol{x}), \alpha_t} \left|\left| g_t(\boldsymbol{x}) - C(\boldsymbol{x}) - \alpha_t F(\boldsymbol{x}) \right|\right|_{\mathbb{L}^2(t \times \boldsymbol{x})}.$$

Lemma 11. *The efficient affine basis function can be implemented using* $C(\boldsymbol{x}) = \mathbb{E}_0(g_t(\boldsymbol{x}))$ *and*

$$F(\boldsymbol{x}) = \left(\sum_t \alpha_t' \alpha_t \right)^{-1} \sum_t \alpha_t' \Big[g_t(\boldsymbol{x}) - C(\boldsymbol{x}) \Big],$$

where $\alpha_t' = (F(\boldsymbol{x})F(\boldsymbol{x})')^{-1} F(\boldsymbol{x})(g(\boldsymbol{x}) - C(\boldsymbol{x}))'$.

[60]This is precisely the basis used in the class of linear perturbation methods recently popularized by Reiter (2009).

The goal of Lemma 11 is to estimate an efficient basis function representation

$$F(\boldsymbol{\alpha}_t)(\boldsymbol{x}) \approx g_t(\boldsymbol{x})$$

after simulating the model to obtain function data for $g_t(\boldsymbol{x})$. Lemma 3 tackles this problem by restricting attention, for clearer illustration, to the affine class of basis functions given by

$$C(\boldsymbol{x}) + T(\boldsymbol{x})\boldsymbol{\alpha}_t.$$

The idea is to express the simulated data $g_t(\boldsymbol{x})$ as

$$g_t(\boldsymbol{x}) = C(\boldsymbol{x}) + T(\boldsymbol{x})\boldsymbol{\alpha}_t + \epsilon_t(\boldsymbol{x}),$$

where ϵ_t is the residual error in the approximate distribution representation. It is important to note that $\epsilon_t(\boldsymbol{x})$ is itself a *functional* residual since we are considering an estimation problem on functional data.

Consider the problem

$$\min_{T(\boldsymbol{x}),C(\boldsymbol{x}),\boldsymbol{\alpha}_t} \left|\left|g_t(\boldsymbol{x}) - C(\boldsymbol{x}) - T(\boldsymbol{x})\boldsymbol{\alpha}_t\right|\right|_{\mathbb{L}^2(t\times\boldsymbol{x})}$$

where α_t is again a $N \times 1$ vector. The matrix of basis functions T can be thought of as a $J \times N$ matrix in the context of a discretized grid of J nodes (i.e. J grid points for household state $\boldsymbol{x}$). Similarly, $C(\boldsymbol{x})$ can be thought of as a $J \times 1$ vector. The $\mathbb{L}^2(t \times \boldsymbol{x})$ norm under consideration could correspondingly be recast as

$$\min_{T_{nj},C_j,\alpha_{tn}} \sum_{t,j} \left(g_{tj} - C_j - \sum_n \alpha_{tn} T_{nj}\right)^2.$$

In terms of the residual simulation error mentioned in the previous paragraph, this loss criterion is equivalent to $\sum_i \epsilon_{ti}^2 = \sum_i \epsilon_t(x_i)^2 = \epsilon_t'(\boldsymbol{x})\epsilon_t(\boldsymbol{x})$ for a given t.

I will set the constant equal to the unconditional mean of the simulated distribution.

$$C = \mathbb{E}_T(g_t)$$

Let $\hat{g}_t = g_t - C$. Then we have

$$\min \sum_t \left(\hat{g}_t - T\boldsymbol{\alpha}_t\right)'\left(\hat{g}_t - T\boldsymbol{\alpha}_t\right).$$

Factoring out,

$$\min \sum_t \left(\hat{g}_t'\hat{g}_t - \hat{g}_t'(T\boldsymbol{\alpha}_t) - (T\boldsymbol{\alpha}_t)'\hat{g}_t + (T\boldsymbol{\alpha}_t)'(T\boldsymbol{\alpha}_t)\right)$$

$$\sum_t \left(\hat{g}_t'\hat{g}_t - 2(T\boldsymbol{\alpha}_t)'\hat{g}_t + (T\boldsymbol{\alpha}_t)'(T\boldsymbol{\alpha}_t)\right)$$

Differentiating with respect to α_t yields

$$0 = -2\frac{\partial}{\partial \boldsymbol{\alpha}_t}(T\alpha_t)'\hat{g}_t + \frac{\partial}{\partial \boldsymbol{\alpha}_t}(T\boldsymbol{\alpha}_t)'(T\boldsymbol{\alpha}_t)$$
$$= -2\underbrace{T'\hat{g}_t}_{N\times 1} + \underbrace{(T'T + (T'T)')\boldsymbol{\alpha}_t}_{N\times 1}.$$

Rearranging, we have

$$(T'T + (T'T)')\boldsymbol{\alpha}_t = 2T\hat{g}_t$$

Finally, note that $T'T = (T'T)'$ so that we have

$$\boldsymbol{\alpha}_t = (T'T)^{-1}T'\hat{g}_t.$$

This is, of course, precisely consistent with the estimation problem for the law of motion $d\boldsymbol{\alpha}_t$, which is the subject of Proposition 3.

Next, consider the optimality condition for $T(\boldsymbol{x})$. Using

$$\frac{\partial}{\partial X}b'X'DXc = D'Xbc' + DXcb'.$$

we have

$$0 = \sum_t \frac{\partial}{\partial T}\left[\hat{g}_t\hat{g}_t' - 2(T\boldsymbol{\alpha}_t)'\hat{g}_t + (T\boldsymbol{\alpha}_t)'(T\boldsymbol{\alpha}_t)\right]$$

which becomes

$$0 = \sum_t \left[-2\hat{g}_t\boldsymbol{\alpha}_t' + 2T(\boldsymbol{\alpha}_t\boldsymbol{\alpha}_t')\right],$$

or simply

$$0 = \sum_t \left[\Big(\hat{g}_t - T\alpha_t\Big)\boldsymbol{\alpha}_t'\right].$$

Rearranging, we can solve for $T(\boldsymbol{x})$ as

$$T(\boldsymbol{x}) = \left(\sum_t \boldsymbol{\alpha}_t\boldsymbol{\alpha}_t'\right)^{-1} \sum_t \hat{g}_t\boldsymbol{\alpha}_t'$$

This concludes the proof.

C Appendix for Section 4: Data and Empirics

C.1 Employment Transitions in CPS Micro Data

In this subsection, I provide details and additional material for my estimation of employment transition rates using Current Population Survey (CPS) data.

C.1.1 Raw CPS data

The Current Population Survey (CPS) is a household survey conducted by the Bureau of Labor Statistics (BLS) at a monthly frequency since 1948. The survey features a rotating panel of households. A household is typically in the survey for four consecutive months. It is therefore possible to use a household identifier to match records across months and create a panel of household employment transitions. There is a long tradition of constructing gross employment flows from matched CPS micro data. See for example Poterba and Summers (1986), Blanchard and Diamond (1990) and Shimer (2012).

I obtain the raw CPS data from the NBER.[61] I restrict my sample to the period 1996 through 2019 throughout the analysis. A key question I use to construct employment categories was only introduced to the CPS in 1994. Furthermore, the years 1994 and 1995 saw several changes in, for example, the definition of household identifiers. This is a well-known issue that has made it impossible to match data across several months in 1994 and 1995. I therefore start my sample in 1996. Similarly, I end the sample in December 2019, before the onset of the Covid pandemic in the U.S.

C.1.2 Employment status definitions

The conventional definition of unemployment used by the BLS counts those above age 16 who are currently unemployed but actively seeking employment, as a percent of the civilian labor force. This notion is also called headline unemployment or U3 unemployment.

U3 has long been a contentious definition of unemployment and is oftentimes criticized as being too narrowly defined. Indeed, it is well known that groups such as the *marginally attached* exhibit transition behavior that is significantly distinct from a more narrowly defined *not in the labor force* (NILF) group and more similar to the behavior of the unemployed. Specifically, the marginally attached (and some other groups) are much more likely to transition into employment or unemployment than discouraged workers, the disabled or the retired, all of whom are lumped together in the NILF category.

My model calls for a definition of employment and unemployment that captures as large a share as possible of those households who exhibit positive employment transition rates. Since the model speaks to questions about *aggregate* consumer behavior, it is only fitting to include as large a share of the total (prime-age) population as is reasonable when calibrating its parameters.

[61] See https://data.nber.org/data/cps-basic2/.

On the other hand, there is also a strong argument to *exclude* households that are not currently in the labor force and are sufficiently unlikely to transition back. To include such households, for example the disabled or retired, it would be more consistent to define a distinct employment state, in which households are not exposed to employment transitions at all. As I want to restrict the model to two employment states, the most reasonable compromise is to define a non-employment group that includes all those households that currently face a sufficiently high probability of transitioning back into employment.

Concretely, I let employment state E correspond to the conventional definition of *employment*. On the other hand, I define the model's unemployment state U as the union of the conventionally unemployed, the marginally attached, and those employed part-time involuntarily for economic reasons. I furthermore restrict the sample to prime-age workers. Therefore, my overall definition of unemployment is more similar to the BLS' U6 measure.

I define the *marginally attached* as those households currently not in the labor force that (i) want a job, (ii) have looked for employment in the last 12 months, and (iii) would be able to take a job next week if offered.[62] Similarly, I categorize households as unemployed part-time for economic reasons if they are currently working fewer than 35 hours per week but would like to work more if given the chance.[63]

C.1.3 Creating monthly panel data

To match households across their months in the survey and create an associated panel, I follow the steps in Shimer (2012).[64] In particular, to match households I use the household identifier variables provided in the CPS, variables for race, sex and age, as well as the household's line item number and months in the survey.[65]

C.1.4 From gross flows to continuous-time transition rates

After taking the steps discussed thus far, one can compute the monthly *gross flows* across employment states. I now discuss how to map discrete-time flow data at a monthly frequency to continuous-time transition rates at a quarterly frequency. The following derivations are closely related to those in Shimer (2012) who argues that using continuous-time transition rates can circumvent a time aggregation bias inherent in discrete-time data.

Let j and k denote employment states, $j, k \in \{U, E\}$. The time horizon is $t \in [0, \infty)$ at a quarterly frequency. Denote by $N_t^{jk}(\tau)$ the number of households in employment state j at

[62] The corresponding CPS variables are PRWNTJOB, PEDWLKO and PEDWAVL.

[63] The corresponding CPS variable is PEHRRSN1.

[64] My code builds on the replication code files provided by Robert Shimer, https://sites.google.com/site/robertshimer/research/flows. For additional details, please see Shimer (2012).

[65] Households are in the CPS for four consecutive months, then leave the survey and subsequently come back for a second rotation of four consecutive months. I use data from both rounds.

time t and state k at time $t + \tau$. Define

$$n_t^{jk}(\tau) = \frac{N_t^{jk}(\tau)}{N_t^{jj}(\tau) + N_t^{jk}(\tau)}$$

as the share of all those households in employment state j at time t that have transitioned into state k at time $t + \tau$. Naturally, $n_t^{jk}(0) = N_t^{jk}(0) = 0$ for $k \neq j$, and $n_t^{jj}(\tau) = 1 - n_t^{jk}(\tau)$. Finally, I denote by λ_t^j the continuous-time transition rate *out of* state j at time t. This is the same transition rate that appears in the quantitative model. In particular, λ_t^E is the job separation rate and λ_t^U the job finding rate.

Lemma 12. *The relationship between flow shares and transition rates is given by*

$$n_t^{jk}(\tau) = \lambda_t^j \frac{1 - e^{-(\lambda_t^j + \lambda_t^k)\tau}}{\lambda_t^j + \lambda_t^k}.$$

Lemma 13. *Data on discrete-time gross flows at monthly frequency (from the CPS) can be mapped into continuous-time Poisson transition rates at quarterly frequency using the formulas*

$$\lambda^U = -\frac{n^{UE}}{n^{UE} + n^{EU}} \cdot 3 \log\left(1 - n^{EU} - n^{UE}\right)$$
$$\lambda^E = -\frac{n^{EU}}{n^{UE} + n^{EU}} \cdot 3 \log\left(1 - n^{EU} - n^{UE}\right).$$

C.1.5 Estimating the cyclicality of employment transition rates

Coming soon.

C.2 Business Cycle Moments

The quantitative model generates naturally asymmetric business cycles and closely matches prominent features of the business cycles in U.S. postwar history. In this section, I discuss the data I use to compute these business cycle moments.

C.2.1 Raw data

The primary series I consider are output, investment, consumption and hours worked. Investment is defined as non-residential fixed investment and durable goods consumption. Consumption is defined as expenditures on non-durable goods and services. Output, consumption and investment are from the NIPA accounts. For hours, I use hours of all persons in the nonfarm business sector from the BLS.

These series are expressed in real terms after deflating by the CPI index. Similarly, all variables are expressed per capita, dividing by the civilian non-institutionalized population aged 16 and over, and in logs.

I consider a sample from the first quarter of 1953 until the fourth quarter of 2019. I exclude the immediate postwar years as they featured a degree of volatility in, for example, GDP growth that is uncharacteristic for the postwar period as a whole.

Finally, I detrend all series using an HP filter, as is standard. I use a smoothing parameter of $\lambda = 1600$, which is appropriate for quarterly data.

C.3 Macroeconomic Uncertainty Indices

Since the seminal work by Bloom (2009), there has been a surge in interest in identifying and measuring macro uncertainty in the data. Numerous proxies and indices have been proposed. To confront my model's predictions about uncertainty with data, I use a whole range of empirical macro uncertainty proxies.

- Jurado et al. (2015) provide several direct econometric estimates of macro uncertainty. They interpret uncertainty as conditional volatility in a series' unforecastable component and macro uncertainty in particular as the *common* variation in a host of aggregate time series. Using close to 300 macroeconomic and financial time series, they estimate factor-based forecasting models and compute a macroeconomic uncertainty index (as well as financial and real uncertainty indices) at 1, 3 and 12-month horizons.[66]

- The Chicago Board Options Exchange (CBOE)'s Volatility Index (VIX) is a commonly used proxy for expected volatility in financial markets. Using S&P 500 index options, the index computes a one-month forward looking market expectation of volatility.

- The Economic Policy Uncertainty Index I use is from Baker et al. (2016).[67] Based on newspaper coverage, this index tallies the frequency of articles in leading U.S. newspapers that "*contain the following triple: 'economic' or 'economy'; 'uncertain' or 'uncertainty'; and one or more of 'congress', 'deficit', 'Federal Reserve', 'legislation', 'regulation' or 'White House'*."

- Finally, the World Uncertainty Index for the United States is from Ahir et al. (2018).[68] The index is based on the frequency with which the word "uncertainty" appears in the Economist Intelligent Unit's quarterly country reports.

[66] The raw time series can be found on Sydney Ludvigson's website, https://www.sydneyludvigson.com/.
[67] I access the data using FRED.
[68] I also access this time series using FRED.

D Appendix for Section 5: Details, Robustness and Sensitivity Analysis

In progress and coming soon.

E Extension 1: A Model with Stocks

In this appendix section, I study a variant of the quantitative model with an alternative asset market structure. I assume that households can trade bonds and *stocks*. Unlike capital in the baseline model, stocks are liquid and risky. Households do not incur transaction costs when trading stocks. This setting is therefore closer to the traditional consumption-based asset pricing framework. I focus on studying the implications of *household heterogeneity* for asset prices in this environment.

As in Section A, I maintain throughout that the economy admits a finite-dimensional state space representation where, in particular, the aggregate state Γ_t follows a time-homogeneous Ito diffusion process. This is precisely the representation that my solution method implies. See the introduction to Section A for further details.

The supply side of this model variant is unchanged. In particular, I allow for nominal rigidities in both intermediate goods prices (firm problem) and wages (union problem). The main focus of this section is on the household problem, which I discuss next.

E.1 Household Problem

Household preferences are again given by

$$\mathbb{E}_0 \int_0^\infty e^{-\int_0^t (\rho_s + \zeta)s \, ds} u(c_t, h_t) dt,$$

where labor supply $h_t = H_t$ is set symmetrically by labor unions.

Households consume and save, investing their wealth in bonds and stocks. As in the baseline model, I denote by a the household's outstanding stock of bonds or liquid assets. To emphasize that households now trade a different second asset, I denote the household's stock of risky (and liquid) assets by s. In equilibrium, the aggregate stock of shares held by the household sector will, of course, be equal to the outstanding stock of capital that is used by firms in production.

The household's liquid asset position evolves according to

$$\dot{a} = \Big(r + \zeta\Big) a + Ds + e - c - \iota s,$$

which is unchanged from the baseline model except that the cost of stock purchases is now given by ιs, where ι is the investment rate. I also denote by D the dividend cash flow paid to the risky asset. The household's position in the risky asset evolves according to

$$\dot{s} = \Big(\Phi(\iota) + \zeta - \delta\Big) s.$$

In other words, households are endowed with a technology that transforms ιs units of the consumption good into $\Phi(\iota) s$ units of the risky asset. While concavity in $\Phi(\cdot)$ still captures a kind of adjustment cost, I restrict my attention only to a particular kind of adjustment cost as I discuss next.

As before, household earnings are given by

$$e = (1-\tau^{\text{lab}})zwH + \tau^{\text{lump}} + \tau^{\text{UI}}(z).$$

For expositional convenience, I assume that z_t follows a diffusion rather than a Poisson process. That is, I assume that

$$dz = \mu_z dt + \sigma_z dW,$$

where W_t is a standard Brownian motion that is uncorrelated across households.

Recasting household problem in terms of net worth. I will now discuss this section's main departure from the baseline model of the main text. In particular, I assume that the adjustment cost implicit in $\Phi(\cdot)$ is of such a form that the state space of the household problem can be reduced from (a, s, z, Γ) to (n, z, Γ) where n is the household's net worth. This state space reduction is a standard approach in many asset pricing settings.

I now derive a recursive representation of the household problem in terms of *liquid net worth*, which is defined by the equations $\theta n = Qs$ and $(1-\theta)n = a$. That is, total liquid net worth is $n = Qs + a$. By Ito's lemma, we have $dn = sdQ + Qds + (ds)(dQ) + da$. Since the evolution of the household's risky asset position is non-stochastic, however, the term $(ds)(dQ)$ vanishes. That is, there is no capital quality risk. Plugging in and simplifying yields

$$dn = rndt + \theta n \left[\frac{D-\iota}{Q} + \frac{dQ}{Q} + \Phi(\iota) - \delta - r \right] dt + \Big(e - c \Big) dt$$

which is the key equation of this section. Two observations are in order. First, the choice of ι is entirely static in this setting, yielding the optimality condition

$$\Phi'(\iota) = \frac{1}{Q},$$

so that optimal household investment is only a function of the price of capital, $\iota = \iota(Q)$. Second, conditional on household net worth its evolution is independent from previous portfolio allocation: Since households can directly trade the risky asset at price Q on the secondary market without incurring transaction costs, households can instantaneously rebalance their portfolio as desired. In other words, the risky asset share θ is a *direct* choice variable for the household.

These two features of the household problem allow me to work with the following, simplified representation. Define

$$dR = \underbrace{\frac{D-\iota(Q)}{Q}}_{\text{Dividend yield}} dt + \underbrace{\Big[\Phi(\iota(Q)) - \delta\Big] dt + \frac{dQ}{Q}}_{\text{Capital gains}} \equiv \mu_R dt + \sigma_R dB$$

to be the effective rate of return on the risky asset. I use

$$\mu_R = \frac{D - \iota(Q)}{Q} + \Phi(\iota(Q)) - \delta + \mu_Q$$

$$\sigma_R = \sigma_Q.$$

After internalizing the optimal internal rate of capital investment, $\iota = \iota(Q)$, this return is exogenous from the perspective of the household: it depends on macro conditions and prices, but not on the particular portfolio composition of the household. I can therefore rewrite the law of motion of the household's liquid net worth as

$$dn = rndt + \theta n(dR - rdt) + (e - c)dt$$

or simply

$$dn = rndt + \theta n(\mu_R - r)dt + (e - c)dt + \theta n\sigma_R dB,$$

where $\theta n\sigma_R$ represents the household's total exposure to aggregate risk *via* the risky asset and potential capital gains.

The strategy will now be to write the household problem recursively using as state variables (n, z, Γ). In the baseline model, I used the asset positions (a, k) as state variables for the household problem. Since they evolve non-stochastically, the associated portfolio equations in that setting do not load on the aggregate risk factor. An important implication of this structure is, of course, that the KF equation describing the law of motion of the cross-sectional household distribution is an ordinary rather than a stochastic partial differential equation. Analogously, an application of Proposition 3 to the baseline model implied $\sigma_{\boldsymbol{\alpha}} = 0$.

Using household net worth n directly as state variable implies a direct exposure to the aggregate risk factor. As a result, the associated HJB will feature a new set of terms related to this aggregate risk exposure. Similarly, I will show that the KF equation in this setting becomes a *stochastic* partial differential equation, with $\sigma_{\boldsymbol{\alpha}} \neq 0$ under the representation of Proposition 3.

Borrowing constraint. I assume that households are subject to a borrowing constraint on net worth

$$n \geq \underline{n}.$$

Recursive representation. The household problem can be written in terms of the household state variables (n, z) as well as the aggregate state space Γ. After internalizing the optimal choice of internal capital investment, $\iota(Q)$, we have

$$(\rho + \zeta)V = \max_{c,\theta} \Big\{ u(c, h) + V_n \Big[rn + \theta n(\mu_R - r) + e - c \Big] + \frac{1}{2} V_{nn}(\theta n\sigma_R)^2$$
$$+ V_z\mu_z + \frac{1}{2}V_{zz}\sigma_z^2 + V_{n\Gamma}\theta n\sigma_R\sigma_\Gamma + V_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma \Big\},$$

where I assume that $\mathbb{E}[dWdB] = 0$. As I have anticipated, two new terms emerge in this setting because the household's net worth evolution equation loads directly on the aggregate risk factor. First, there is the direct risk exposure, which is scaled by V_{nn}. Second, households are now exposed both directly through net worth and indirectly through Γ, so that a covariance term emerges that is scaled by $V_{n\Gamma}$.

The first-order conditions for consumption and portfolio choice are given by

$$u_c = V_n$$
$$\theta = -\left(\frac{V_n}{nV_{nn}}\frac{\mu_R - r}{\sigma_R^2} + \frac{V_{n\Gamma}}{nV_{nn}}\frac{\sigma_\Gamma}{\sigma_R}\right).$$

The first-order condition for labor is

$$-u_h = u_c(1-\tau^{\text{lab}})wz.$$

The household HJB and the corresponding optimality condition hold everywhere in the interior of the household state space.

Lemma 14. *The household Euler equations for marginal utility and consumption are given by*

$$\frac{du_c}{u_c} = (\rho - r)dt - \frac{\mu_R - r}{\sigma_R}dB + \sigma_z\frac{u_{cc}}{u_c}c_z dW$$

and, assuming CRRA preferences,

$$\frac{dc}{c} = \frac{r-\rho}{\gamma}dt + \frac{1}{2}(1+\gamma)\left[\left(\frac{\mu_R - r}{\gamma\sigma_R}\right)^2 + \left(\frac{c_z\sigma_z}{c}\right)^2\right]dt + \frac{\mu_R - r}{\gamma\sigma_R}dB + \frac{c_z\sigma_z}{c}dW.$$

It is instructive to compare these Euler equations to the analogous set of equations derived in Appendix A.1. For convenience, I restate here the Euler equation for marginal utility

$$\frac{du_c}{u_c} = (\rho - r)dt + \sigma_\Gamma\frac{u_{cc}}{u_c}c_\Gamma dB + \sigma_z\frac{u_{cc}}{u_c}c_z dW.$$

One can clearly see that the volatility-adjusted expected excess return on the risk asset, $(\mu_R - r)/\sigma_R$ becomes a sufficient statistic of sorts for the household's exposure to aggregate risk. In the baseline model, this exposure is captured by the scaling factor

$$\sigma_\Gamma\frac{u_{cc}}{u_c}c_\Gamma$$

instead.

Comparison to Brunnermeier and Sannikov (2014). The model presented in this section is much like that in Brunnermeier and Sannikov (2014) except that it features a continuum of heterogeneous households instead of a two-agent structure. In my setting,

$$\frac{dV_n}{V_n} = (\rho - r) - \frac{\mu_R - r}{\sigma_R}dB + \frac{V_{nz}}{V_n}\sigma_z dW.$$

This equation can be directly mapped into Proposition II.2 of Brunnermeier and Sannikov (2014), using the change in notation $V_n = \theta$. For the sake of comparison, let $V_n = \theta^{\text{BS}}$. Then,

$$\frac{d\theta^{\text{BS}}}{\theta^{\text{BS}}} = \mu_\theta dt + \sigma_\theta^B dB + \sigma_\theta^W dW,$$

where

$$\mu_\theta = \rho - r$$
$$\underbrace{-\sigma_Q \sigma_\theta^B}_{\text{Risk premium}} = \underbrace{\frac{D - \iota(Q)}{Q} + \Phi(\iota(Q)) - \delta + \mu_Q - r}_{\text{Expected excess return on capital}}$$
$$\sigma_\theta^W = \frac{V_{nz}}{V_n}\sigma_z.$$

Of course, Brunnermeier and Sannikov (2014) only have aggregate risk and do not have the earnings risk term.

E.2 Aggregation, Government and Market Clearing

The household-level state space is given by $(n, z) \in [\underline{n}, \infty) \times [\underline{z}, \bar{z}]$. I will denote the joint household income and wealth distribution by $g_t(n, z)$. Aggregate consumption, for example, is then defined as

$$C(\Gamma) = \int c(n, z, \Gamma) g(n, z, \Gamma) d(n, z).$$

When convenient, I will drop the notational dependence on Γ. Aggregate net worth is defined as

$$N = \int n g(n, z) d(n, z).$$

Finally, aggregate capital investment by households is defined as

$$I = \int \iota(Q) k g(n, z) d(n, z) = \frac{\iota(Q)}{Q} \int \theta(n, z) n g(n, z) d(n, z).$$

The evolution of the household income and wealth distribution is given by the following Proposition, where I use a shorthand for household saving, $s = rn + \theta n(\mu_R - r) + e - c$.

Lemma 15. *The evolution of $g(n, z)$ is characterized by the stochastic Kolmogorov forward equation*

$$dg = -\partial_n \Big[sg \Big] + \frac{1}{2} \partial_{nn} \Big[(\theta n \sigma_R)^2 g \Big] - \partial_z \Big[\mu_z g \Big] + \frac{\sigma_z^2}{2} \partial_{zz} g - \partial_n \Big[\theta n \sigma_R g \Big] dB.$$

Proof. This is a special case of the stochastic KF equation derived in Appendix B. ■

Government. The description of the government sector is unchanged from the baseline model, with monetary policy following a Taylor rule.

Market clearing. There are four markets in this economy that must clear: the markets for goods, labor, stocks and bonds. The labor market clearing condition simply asserts that, as in the baseline model, labor unions mandate symmetric supply of work hours.

Households' aggregate bond holdings must be exactly equal to the government's outstanding debt position, so that

$$0 = B^G + \int [1 - \theta(n,z)] n g(n,z) d(n,z).$$

Similarly, aggregate holdings of the risky asset by households must in equilibrium be equal to the value (or market capitalization) of the outstanding stock of capital, so that

$$QK = \int \theta(n,z) n g(n,z) d(n,z).$$

Rewriting the bond market clearing condition, we have

$$N = \int n g(n,z) d(n,z) = \int \theta(n,z) n g(n,z) d(n,z)$$

Putting the two together yields

$$N = QK,$$

so that aggregate household net worth must be equal to the value of outstanding capital. Finally, the goods market clearing condition is

$$Y = C + G + I$$

which can be readily verified by working through Walras' law.

Finally, all this implies that the law of motion for aggregate net worth in this economy is given by

$$\frac{dN}{N} = \Big(\mu_Q + \Phi(\iota) - \delta\Big) dt + \sigma_Q dB$$

and that for capital by

$$dK = \Big(\Phi(\iota) - \delta\Big) K dt.$$

Just for illustration and better intuition, the law of motion for aggregate net worth can then be written as

$$\begin{aligned} dN &= K\Big(Q\mu_Q dt + Q\sigma_Q dB\Big) + QK\frac{dK}{K} dt \\ &= KdQ + QdK, \end{aligned}$$

which of course follows directly from Ito's lemma as well, noting that $N = QK$, and therefore confirms the derivation.

E.3 Summary of Equilibrium Conditions

The equations comprising the equilibrium of this economy can be collected into three blocks.

Macro block. The macro relationships are given by

$$
\begin{aligned}
Y &= e^Z K^\alpha L^{1-\alpha} \\
Y &= C + \iota K \\
dK &= \Phi(\iota)K - \delta K \\
mc &= \frac{1}{e^Z}\frac{D^\alpha w^{1-\alpha}}{\alpha^\alpha (1-\alpha)^{1-\alpha}} \\
\Pi &= (1 - mc)Y \\
D &= \frac{\alpha}{1-\alpha} w \frac{L}{K} \\
i &= r^* + \lambda\pi \\
r &= i - \pi \\
\mu_R &= \frac{D-\iota}{Q} + \Phi(\iota) - \delta + \mu_Q \\
\sigma_R &= \sigma_Q \\
\Phi'(\iota) &= \frac{1}{Q} \\
N &= QK \\
\tau &= \Pi + \tau^{\text{lab}} wL \\
\pi_t &= \mathcal{W}_t(\pi_t)
\end{aligned}
$$

together with the law of motion for aggregate productivity given by

$$dZ = -\theta Z dt + \sigma dB.$$

Aggregation block. The aggregation block was already described earlier.

Micro block. The micro block solves for the policy functions $c(n, z, \Gamma)$, $h(n, z, \Gamma)$ and $\theta(n, z, \Gamma)$ for every aggregate state, given aggregates

$$\{r, w, \tau, D, Q, \mu_R, \sigma_R\}.$$

F Extension 2: Financial Frictions and Intermediary-Based Asset Pricing

Introduction:

- In the wake of the Great Recession, we saw a resurgence in emphasis on *financial frictions*. See for example Gertler and Kiyotaki (2010), Brunnermeier and Sannikov (2014) and He and Krishnamurthy (2013).
- Since this initial wave of interest, the study of cross-sectional heterogeneity among households and firms has emerged as the new frontier of macroeconomic business cycle analysis. The burgeoning heterogeneous-agent New Keynesian (HANK) literature is, at present, entirely divorced from work on financial frictions. In large part, this is because financial frictions represent a non-linearity at the macro level that requires the kind of global solution method that has remained elusive in the context of heterogenoeus-agent macro models. In this appendix section, I try to bridge this gap by solving a HANK model with financial frictions at the macro level globally.
- While the baseline model presented in the main text focuses on demand (discount rate) shocks in the context of a zero lower bound (ZLB) constraint, the empirical relevance of the ZLB has only been a recent phenomenon. This section, on the other hand, focuses on supply (TFP) shocks in the context of financial frictions. I demonstrate that my main insights on the interaction between micro and macro uncertainty remain robust in this setting.
- This section focuses on three main questions:
  1. What is the transmission of uncertainty in the presence of financial frictions? Are my results on the interaction between micro and macro uncertainty robust in a setting with supply shocks and financial frictions?
  2. What is the transmission mechanism of *financial shocks* in a heterogeneous-agent New Keynesian model?
  3. How does household debt relief compare against bank recapitalization as a stabilization tool in this setting?

F.1 Model

The model I present in this section can be thought of as a continuous-time version of Gertler and Karadi (2011) with a continuum of heterogeneous households. I start by discussing the problem of financial intermediaries and the financial friction they face.

F.1.1 Financial intermediary

Instead of setting up the classic structure used in Gertler and Karadi (2011) where households consist of workers and bankers, I assume that bankers die and are born at rate ζ outright. I also assume outright that bankers find it optimal not to pay out net worth to households until they exit.

Balance sheet. The representative intermediary's balance sheet is given by

$$\underbrace{Q_t s_t}_{\text{Assets}} = \underbrace{n_t + b_t}_{\text{Liabilities}},$$

where Q_t is the prevailing market price of equity, s_t is the stock of equity shares held by the intermediary, b_t is the outstanding debt of the intermediary, and n_t is its net worth (or equity capital).

I can anticipate at this point that I will assume Modigliani-Miller and focus on the case where firms are exclusively equity-financed. For now, however, the details of firms' capital structures are not relevant for the financial intermediary's portfolio allocation problem. In particular, the intermediary earns a return dR_t (in units of numeraire) on its equity portfolio, which it takes as given, and pays the riskfree rate r_t on its debt. In the absence of payouts, the intermediary's net worth position then evolves according to

$$\begin{aligned} dn_t &= Q_t s_t dR_t - r_t b_t \\ &= r_t n_t + Q_t s_t (dR_t - r_t). \end{aligned}$$

It will be convenient to work with the portfolio share θ defined as $\theta_t n_t = Q_t s_t$ and $(1-\theta_t)n_t = -b_t$. This implies

$$dn_t = r_t n_t + \theta_t n_t (dR_t - r_t).$$

Optimal portfolio allocation. I assume that bankers exit at rate ζ and find it optimal not to pay out any funds to their shareholders until they exit, at which point they rebate their entire net worth position. Denote by the random variable τ the stopping time of exiting. Then the value of a small intermediary franchise at time 0 is given by

$$V_0 = \max_{\{s_t, b_t\}} \mathbb{E}_0 \left[\frac{\Lambda_\tau}{\Lambda_0} n_\tau \right],$$

where Λ_t is the discount factor applied by the intermediary. I will further elaborate on this SDF in later subsections.

The following Lemma provides a recursive representation of the intermediary's problem. I first present the intermediary problem in the absence of a financial constraint for illustration. Afterwards, I introduce a net worth constraint and showcase how it affects the intermediary's recursive problem.

Lemma 16. *When financial intermediaries are unconstrained, their continuous-time value function solves the Hamilton-Jacobi-Bellman equation*

$$(\rho+\zeta)V_t = \max\ \zeta\Lambda_t n_t + \mathbb{E}_t\left[\frac{dV_t}{dt}\right].$$

Proof. We have

$$\begin{aligned} V_0 &= \max\ \mathbb{E}_0\left[\frac{\Lambda_\tau}{\Lambda_0}n_\tau\right] \\ &= \max\ \mathbb{E}_0\int_0^\infty \mathbb{P}(t=\tau)\left[\frac{\Lambda_t}{\Lambda_0}n_t\right]dt \\ &= \max\ \int_0^\infty \zeta e^{-\zeta t}\mathbb{E}_0\left[\frac{\Lambda_t}{\Lambda_0}n_t\right]dt \end{aligned}$$

In discrete time, it would read as follows

$$V_0 = \max\ \sum_{t=1}^{\infty}\mathbb{P}(t=\tau)\mathbb{E}_0\left[\frac{\Lambda_t}{\Lambda_0}n_t\right]$$

Notice that the summation here starts at $t=1$. Next, recalling that ζ is the *rate* at which bankers exit,

$$\begin{aligned} V_0 &= \max\ \sum_{t=1}^{\infty}\theta(1-\theta)^{t-1}\mathbb{E}_0\left[\frac{\Lambda_t}{\Lambda_0}n_t\right] \\ &= \max\ \mathbb{E}_0\left[\theta\frac{\Lambda_1}{\Lambda_0}n_1 + \sum_{t=2}^{\infty}\theta(1-\theta)^{t-1}\left(\frac{\Lambda_t}{\Lambda_0}n_t\right)\right] \end{aligned}$$

We now want to use

$$\begin{aligned} V_t &= \max\ \sum_{s=1}^{\infty}\theta(1-\theta)^{s-1}\mathbb{E}_t\left[\frac{\Lambda_{t+s}}{\Lambda_t}n_{t+s}\right] \\ V_{t+1} &= \max\ \sum_{s=1}^{\infty}\theta(1-\zeta)^{s-1}\mathbb{E}_{t+1}\left[\frac{\Lambda_{t+1+s}}{\Lambda_{t+1}}n_{t+1+s}\right] \end{aligned}$$

Thus,

$$\begin{aligned}
V_0 &= \max\ \mathbb{E}_0\left[\theta\frac{\Lambda_1}{\Lambda_0}n_1 + \frac{\Lambda_1}{\Lambda_0}(1-\theta)\sum_{t=2}^{\infty}\theta(1-\theta)^{t-1}\frac{\Lambda_0}{\Lambda_1}\frac{1}{1-\zeta}\left(\frac{\Lambda_t}{\Lambda_0}n_t\right)\right] \\
&= \max\ \mathbb{E}_0\left[\theta\frac{\Lambda_1}{\Lambda_0}n_1 + \frac{\Lambda_1}{\Lambda_0}(1-\theta)\sum_{t=2}^{\infty}\theta(1-\theta)^{t-2}\frac{\Lambda_t}{\Lambda_1}n_t\right] \\
&= \max\ \mathbb{E}_0\left[\theta\frac{\Lambda_1}{\Lambda_0}n_1 + \frac{\Lambda_1}{\Lambda_0}(1-\theta)\sum_{t=1}^{\infty}\theta(1-\theta)^{t-1}\frac{\Lambda_{t+1}}{\Lambda_1}n_{t+1}\right] \\
&= \max\ \mathbb{E}_0\left[\theta\frac{\Lambda_1}{\Lambda_0}n_1 + (1-\theta)\frac{\Lambda_1}{\Lambda_0}V_1\right].
\end{aligned}$$

Now let's move this into continuous time. There are two objects we have to convert, the death probability (into a death rate) and the discount factor. In discrete time with CRRA, we have

$$\Lambda_t = \beta^t c_t^{-\gamma},$$

where c_t is the level of consumption spending per period. Furthermore, θ is a death probability per period. Intermediaries exit with probability θ and remain with probability $(1-\theta)$. We therefore have

$$V_t = \max \mathbb{E}_t\left\{\beta(\Delta)\frac{\tilde{\Lambda}_{t+\Delta}}{\tilde{\Lambda}_t}\left[\theta(\Delta)n_{t+\Delta} + (1-\theta(\Delta))V_{t+\Delta}\right]\right\}$$

We have

$$\beta(\Delta) = e^{-\rho\Delta} \approx 1-\rho\Delta$$

$$1-\theta(\Delta) = e^{-\zeta\Delta} \approx 1-\zeta\Delta.$$

Recall that $(1-\theta)$ is the probability of *remaining*. Thus, we have

$$V_t = \max \mathbb{E}_t\left\{(1-\rho\Delta)\frac{\tilde{\Lambda}_{t+\Delta}}{\tilde{\Lambda}_t}\left[\zeta\Delta n_{t+\Delta} + (1-\zeta\Delta)V_{t+\Delta}\right]\right\}.$$

I will now renormalize the intermediary's value function, which is a typical step in this class of problems. Define the new effective value function as

$$V_t^{\text{new}} = \tilde{\Lambda}_t V_t^{\text{old}}$$

and I will abuse notation and just continue using V notationally. So multiplying the previous equation by $\tilde{\Lambda}_t$ I arrive at

$$\tilde{\Lambda}_t V_t = \max \mathbb{E}_t\left\{(1-\rho\Delta)\tilde{\Lambda}_{t+\Delta}\left[\zeta\Delta n_{t+\Delta} + (1-\zeta\Delta)V_{t+\Delta}\right]\right\}$$

$$V_t = \max \mathbb{E}_t\left\{(1-\rho\Delta)\left[\zeta\Delta\tilde{\Lambda}_{t+\Delta}n_{t+\Delta} + (1-\zeta\Delta)V_{t+\Delta}\right]\right\}.$$

Now we subtract $(1-\rho\Delta)V_t$ from both sides to arrive at

$$V_t-(1-\rho\Delta)V_t=\max\mathbb{E}_t\Big\{(1-\rho\Delta)\zeta\Delta\tilde{\Lambda}_{t+\Delta}n_{t+\Delta}+(1-\rho\Delta)V_{t+\Delta}-(1-\rho\Delta)\zeta\Delta V_{t+\Delta}-(1-\rho\Delta)V_t\Big\}$$

or simply

$$\rho\Delta V_t=\max\mathbb{E}_t\Big\{(1-\rho\Delta)\zeta\Delta\tilde{\Lambda}_{t+\Delta}n_{t+\Delta}+(1-\rho\Delta)(V_{t+\Delta}-V_t)-(1-\rho\Delta)\zeta\Delta V_{t+\Delta}\Big\}.$$

Now dividing by Δ,

$$\rho V_t=\max\mathbb{E}_t\Big\{(1-\rho\Delta)\zeta\tilde{\Lambda}_{t+\Delta}n_{t+\Delta}+(1-\rho\Delta)\frac{V_{t+\Delta}-V_t}{\Delta}-(1-\rho\Delta)\zeta V_{t+\Delta}\Big\}.$$

And finally, taking the limit $\Delta\to 0$, we arrive at

$$\begin{aligned}\rho V_t&=\max\mathbb{E}_t\Big\{\zeta\tilde{\Lambda}_t n_t+\lim_{\Delta\to 0}\frac{V_{t+\Delta}-V_t}{\Delta}-\zeta V_t\Big\}\\&=\max\quad\zeta\tilde{\Lambda}_t n_t-\zeta V_t+\mathbb{E}_t\left[\frac{dV_t}{dt}\right].\end{aligned}$$

■

The state variables of the intermediary problem are (n,Γ), where n is a single intermediary's net worth position and Γ is an as yet to be determined set of aggregate state variables. As in previous sections, I will focus on the case where the aggregate state Γ is finite-dimensional and follows a time-homogeneous Ito diffusion process, which will be consistent with my solution method. In that case, the return process for the risky asset also evolves as a diffusion process,

$$dR=\mu_R dt+\sigma_R dB,$$

where μ_R and σ_R are equilibrium objects. B_t is a standard Brownian motion that represents the sole aggregate risk factor of the model and will be specified below. I can thus rewrite the intermediary's HJB as

$$\begin{aligned}(\rho+\zeta)V(n,\Gamma)=\zeta\Lambda n+V_n\Big[rn+\theta n(\mu_R-r)\Big]+\frac{1}{2}V_{nn}(\theta n\sigma_R)^2\\+V_{n\Gamma}\theta n\sigma_R\sigma_\Gamma+V_\Gamma\mu_\Gamma+\frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma.\end{aligned}$$

Importantly, the derivation thus far assumes that the portfolio allocation of financial intermediary's is unconstrained across the interior of their state space. I will next introduce the financial constraint that plays a key role in this section.

Financial constraint. Following Gertler and Karadi (2011), I assume a balance sheet constraint of the form

$$V(n,\Gamma)\geq\lambda(\Gamma)Q(\Gamma)s(n,\Gamma),$$

where $\lambda(\Gamma)$ is a time-varying financial shock and $s(n,\Gamma)$ is the intermediary's policy function for equity allocation.

Lemma 17. *In the presence of the financial constraint on intermediary net worth, the continuous-time value function of an intermediary solves the Hamilton-Jacobi-Bellman equation*

$$(\rho+\zeta)V(n,\Gamma) = \zeta\Lambda n + V_\Gamma \mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma + \max_\theta \left\{ V_n\Big[rn + \theta n(\mu_R - r)\Big] + \frac{1}{2}V_{nn}(\theta n\sigma_R)^2 + V_{n\Gamma}\theta n\sigma_R\sigma_\Gamma + \mu\Big[V - \lambda\Lambda\theta n\Big] \right\},$$

where μ is the multiplier on the net worth constraint.

Proof. For illustration, I start again with the discrete-time Bellman equation and introduce a Kuhn-Tucker multiplier μ_t, yielding

$$\Lambda_t V_t = \max\ \mathbb{E}_t\left\{\beta\Lambda_{t+1}\Big[\theta n_{t+1} + (1-\theta)V_{t+1}\Big]\right\} + \mu_t\Lambda_t\Big[V_t - \lambda_t Q_t s_t\Big].$$

Using the same steps as before, with the normalization $\Lambda_t V_t \mapsto V_t$, I get

$$V_t = \max\ \mathbb{E}_t\left\{(1-\rho\Delta)\Big[\zeta\Delta\Lambda_{t+\Delta}n_{t+\Delta} + (1-\zeta\Delta)V_{t+\Delta}\Big]\right\} + \mu_t\Big[V_t - \lambda_t\Lambda_t Q_t s_t\Big].$$

Subtracting $(1-\rho\Delta)V_t$ and dividing by Δ, we have

$$\rho V_t = \max\ \mathbb{E}_t\left\{(1-\rho\Delta)\zeta\Lambda_{t+\Delta}n_{t+\Delta} + (1-\rho\Delta)\frac{V_{t+\Delta}-V_t}{\Delta} - (1-\rho\Delta)\zeta V_{t+\Delta}\right\} + \frac{1}{\Delta}\mu_t\Big[V_t - \lambda_t\Lambda_t Q_t s_t\Big],$$

with complementary slackness conditions

$$\text{either}\ \ \mu_t = 0, \qquad \text{or}\ \ V_t - \lambda_t\Lambda_t Q_t s_t = 0.$$

We can now reinterpret the multiplier as a *rate* scaled by Δ. This derivation is informal, of course.

More formally, I follow the approach in Zariphopoulou (1994) who proves the existence and uniqueness of a viscosity solution for a similar but simpler problem. The analogous formulation in my case, using the notion of *admissible solutions*, would be

$$(\rho+\zeta)V(n,\Gamma) = \zeta\Lambda n + \max_{\theta\in\mathcal{A}}\left\{V_n\Big[rn + \theta n(\mu_R - r)\Big] + \frac{1}{2}V_{nn}(\theta n\sigma_R)^2 + V_{n\Gamma}\theta n\sigma_R\sigma_\Gamma\right\} + V_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma.$$

where

$$\mathcal{A} = \Big\{\theta : V \geq \lambda\theta n\Big\}.$$

It is now of illustrative value to recast this problem in terms of the Kuhn-Tucker machinery with complementary slackness conditions.

$$(\rho+\zeta)V(n,\Gamma) = \zeta\Lambda n + V_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma$$
$$+ \max_\theta \left\{ V_n\Big[rn + \theta n(\mu_R - r)\Big] + \frac{1}{2}V_{nn}(\theta n\sigma_R)^2 + V_{n\Gamma}\theta n\sigma_R\sigma_\Gamma + \mu\Big[V - \lambda\Lambda\theta n\Big]\right\}.$$

It's important to note here that Λ enters the constraint now because we have normalized V. ■

The first-order condition for θ is given by

$$0 = V_n n(\mu_R - r) + V_{nn}\theta(n\sigma_R)^2 + V_{n\Gamma}n\sigma_R\sigma_\Gamma - \mu\lambda\Lambda n,$$

or simply

$$\theta = -\frac{V_n}{nV_{nn}}\frac{\mu_R - r}{\sigma_R^2} - \frac{V_{n\Gamma}}{nV_{nn}}\frac{\sigma_\Gamma}{\sigma_R} + \frac{\mu}{nV_{nn}}\frac{\lambda\Lambda}{\sigma_R^2}. \tag{27}$$

The complementary slackness conditions, given by $\mu \geq 0$, $V - \lambda\Lambda\theta n \geq 0$, and $\mu(V - \lambda\Lambda\theta n) = 0$ imply the following procedure for determining θ, V and μ: Start with a guess for the value function. For a given state (n,Γ), solve for the implied θ and check whether it satisfies the constraint. If so, this $\theta(n,\Gamma)$ is a possible solution, for which $\mu(n,\Gamma) = 0$. Now let's consider the other case with $\mu \neq 0$. Importantly, equation (27) is still one of the Kuhn-Tucker conditions and therefore has to hold always. As such, we now obtain θ from the complementary slackness condition with $\mu \neq 0$, implying

$$\theta = \frac{V}{\lambda\Lambda n}$$

and we obtain μ from equation (27), so that

$$\mu = \frac{1}{\lambda\Lambda}\left(V_n(\mu_R - r) + V_{nn}\theta n\sigma_R^2 + V_{n\Gamma}\sigma_R\sigma_\Gamma\right)$$
$$= \frac{1}{\lambda\Lambda}\left(V_n(\mu_R - r) + V_{nn}\frac{V}{\lambda\Lambda}\sigma_R^2 + V_{n\Gamma}\sigma_R\sigma_\Gamma\right).$$

Aggregate intermediary net worth. Intermediaries exit at rate ζ, rebating their net worth positions to households. At the same time, a new cohort enters, also at rate ζ, with startup funds provided by households. Denote the startup funds per intermediary by η, so that the rate at which households provide funds for new intermediaries is given by $\eta\zeta$.

To maintain the assumption of a representative intermediary, we must set $\eta = N$, where N is aggregate net worth. In this case, households never receive any net funds from intermediaries: all net worth rebated from exiting bankers is used as startup money by new entrants. If this was not the case, the new cohort would no longer be symmetric to the incumbents.

However, I will show in the results section that the intermediary problem is, in fact, linear in n. Therefore, intermediary behavior can easily be aggregated and the potentially non-degenerate net worth distribution of the intermediary sector does not become part of the aggregate state space.

F.1.2 Firms

My model of the corporate sector follows closely that of Dou et al. (2020).[69] There is a representative firm that performs two productive functions. On the one hand, the firm uses capital and labor in the production of the consumption good. On the other hand, the firm invests in the production of new capital, which it can then trade on a secondary market at price Q_t. Firms issue equity shares to fund these operations.[70] In equilibrium, the financial intermediary sector will hold these shares. Since the marginal value of retained earnings is always weakly larger in the financial sector, firms are assumed to pay out all net income they generate as dividends to shareholders, rather than accumulate retained earnings. In the absence of retained earnings, firms then issue new equity whenever they want to invest in new capital. As in Gertler and Karadi (2010), for example, I assume that firms issue shares equal to the units of capital they acquire, so that, by arbitrage, corporate equity shares and capital trade at the same price, Q_t.

The representative firm produces the final consumption good using the production technology

$$Y_t = e^{Z_t} K_t^{\alpha} L_t^{1-\alpha},$$

where Z_t denotes aggregate TFP and follows a continuous-time AR(1) process. Time variation in TFP is the source of aggregate risk in this model. In particular, I assume that

$$dZ = -\theta_Z Z dt + \sigma_Z dB.$$

Goods production is perfectly competitive. As a result, we have

$$w_t = (1-\alpha)\frac{Y_t}{L_t}$$

since firms statically set the marginal cost of labor equal to its marginal product.

The representative firm therefore generates net income at rate

$$\text{net income}_t = Y_t - w_t L_t - \Phi(I_t, K_t) + d(Q_t K_t)$$

[69]Unlike in their paper, I set the price of investment goods to 1.

[70]Since the Modigliani-Miller conditions hold for the non-financial corporate sector in this setting, it is without loss to assume equity financing. In other words, while there is a financial friction between intermediaries and households, there is no friction in funding for non-financial firms.

all of which is paid out to shareholders. I assume that capital adjustment costs are quadratic, given by

$$\Phi(I_t, K_t) = I_t + \frac{\kappa}{2}\left(\frac{I_t}{K_t}\right)^2 K_t.$$

The firm's investment decision is static, and we have

$$I_t = \arg\max \left\{ Q_t I_t - I_t - \frac{\kappa}{2}\left(\frac{I_t}{K_t}\right)^2 K_t \right\}.$$

This gives rise to the classic Tobin's-Q expression for the firm's investment choice as a function of the capital price

$$I_t = \frac{Q_t - 1}{\kappa} K_t.$$

In summary, then, the return on corporate equity is given by

$$dR_t = \left[\frac{Y_t - w_t L_t}{Q_t K_t} + \frac{1}{Q_t}\frac{(Q_t - 1)^2}{2\kappa} - \delta + \mu_{Q,t}\right] dt + \sigma_{Q,t} dB$$

where I have conjectured that the price of capital follows a diffusion process in equilibrium

$$\frac{dQ_t}{Q_t} = \mu_{Q,t} dt + \sigma_{Q,t} dB.$$

F.1.3 Households

The model of household behavior is simplified in this section. In particular, I reduce the savings problem of households to just one asset, a. When $a > 0$, a household holds deposits with the financial intermediary. When $a < 0$, the household has taken on a consumer loan.

Household preferences are given by

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} \left[\frac{c_t^{1-\gamma}}{1-\gamma} - \frac{h_t^{1+\eta}}{1+\eta}\right] dt,$$

and the evolution of household savings is given by

$$\dot{a}_t = r_t a_t + (1 - \tau^{\text{lab}}) w_t z_t h_t + \tau_t^{\text{debt relief}}(a) + \tau_t^{\text{lump}} - c_t.$$

Except for the absence of the second asset and the new *debt relief* term, this budget constraint is identical to the one in the baseline model. Depending on their indebtedness, I here allow for a new transfer from the government, $\tau_t^{\text{debt relief}}(a)$, that is meant to capture household debt relief policies. Households again face an overall constraint on borrowing, so that $a_t \geq \underline{a}$.

The household problem can be rewritten recursively in the form of a Hamilton-Jacobi-Bellman equation, given by

$$\rho v(a, k, \Gamma) = \max_{c,h} \left\{ \frac{c^{1-\gamma}}{1-\gamma} - \frac{h^{1+\eta}}{1+\eta} + v_a \left[r_t a_t + (1 - \tau^{\text{lab}}) w z h + \tau^{\text{debt relief}} + \tau^{\text{lump}} - c \right] \right\}$$
$$+ v_z \mu_z + \frac{\sigma_z^2}{2} v_{zz} + v_\Gamma \mu_\Gamma + \frac{1}{2} \sigma_\Gamma^T v_{\Gamma\Gamma} \sigma_\Gamma.$$

The associated optimality condition for consumption is again given by

$$c^{-\gamma} = v_a.$$

The household's labor-leisure condition is given by

$$c^{\gamma} h^{\eta} = (1 - \tau^{\text{lab}}) w z.$$

These conditions characterize the household policy functions $c(a, z, \Gamma)$ and $h(a, z, \Gamma)$.

F.1.4 Aggregation and Market Clearing

The cross-sectional household distribution is again denoted $g_t(a, z)$, so that aggregate consumption and labor supply are given by

$$\begin{aligned} C(\Gamma) &= \int c(a, z, \Gamma) g(a, z, \Gamma) d(a, z) \\ L(\Gamma) &= \int z h(a, z, \Gamma) g(a, z, \Gamma) d(a, z), \end{aligned}$$

where L is *effective* aggregate labor supply. As always, I will suppress the dependence of macroeconomic aggregates on the aggregate state Γ.

Consider the aggregate net worth evolution equation, given by

$$\begin{aligned} dN &= rN + \theta N (dR - r) \\ &= rN + Y - wL + \frac{(Q-1)^2}{2\kappa} K - \delta Q K + QK(\mu_Q dt + \sigma_Q dB) - r\theta N \end{aligned}$$

where $\theta N = QS = QK$. Simplifying this further,

$$\begin{aligned} dN &= r(1-\theta)N + \alpha Y + \Pi^Q - \delta Q K + K dQ \\ K dQ + Q dK - dB &= -rB + \alpha Y + QI - \Phi(I, K) - \delta Q K + K dQ \\ K dQ + Q dK - dB &= -rB + \alpha Y - \Phi(I, K) + Q dK + K dQ \\ -dB &= -rB + \alpha Y - \Phi(I, K) \end{aligned}$$

Finally, I can use the household budget constraint as well as the bond market clearing condition

$$A = B$$

so that

$$dB - rB = dA - rA = w - C = -\alpha Y + \Phi(I, K)$$

and using the fact that labor is paid its marginal product

$$w - C = -Y + wL + \Phi(I, K)$$

or

$$Y = C + \Phi(I, K) = C + I + \frac{\kappa}{2}\left(\frac{I}{K}\right)^2 K.$$

So to summarize, the market clearing conditions in this model are given by

$$\begin{aligned} Y &= C + I + \frac{\kappa}{2}\left(\frac{I}{K}\right)^2 K \\ A &= B \\ S &= K, \end{aligned}$$

where the last equation simply states that the value of aggregate capital must be equal to the aggregate market capitalization of the corporate sector.

F.1.5 Summary of Equilibrium Conditions

Let's collect the main equations. Also, I will commit to the convention that $b < 0$ corresponds to households making deposits at the bank and $b > 0$ to the bank holding bonds.

Financial intermediary. We have

$$\begin{aligned} n &= Qs + b \\ \theta n &= Qs \\ (1-\theta)n &= b \\ dn &= rn + \theta n(\mu_R - r) + \theta n \sigma_R dB \\ V &= \Omega n \end{aligned}$$

The HJB of the financial intermediary in terms of the value function is given by

$$\begin{aligned} (\rho + \zeta)V(n, \Gamma) = {} & \zeta \Lambda n + V_\Gamma \mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T V_{\Gamma\Gamma}\sigma_\Gamma \\ & + \max_\theta \left\{ V_n\left[rn + \theta n(\mu_R - r)\right] + \frac{1}{2}V_{nn}(\theta n \sigma_R)^2 + V_{n\Gamma}\theta n \sigma_R \sigma_\Gamma + \mu\left[V - \lambda\Lambda\theta n\right] \right\}. \end{aligned}$$

Using $V = \Omega n$, this becomes

$$(\rho - r + \zeta)\Omega = \ \zeta\Lambda + \Omega_\Gamma \mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T \Omega_{\Gamma\Gamma}\sigma_\Gamma + \Omega\mu.$$

Unconstrained region. In this region, $V > \lambda\Lambda\theta n$. We have

$$\begin{aligned} \mu &= 0 \\ 0 &= \Omega\frac{\mu_R - r}{\sigma_R} + \Omega_\Gamma \sigma_\Gamma. \end{aligned}$$

From the intermediary's PE perspective, the portfolio share θ remains indeterminate. The the above equation exactly holds, I believe the intermediary is indifferent about any θ. And when this condition is even slightly off, the intermediary would assume infinitely large arbitrage positions. So θ will have to be pinned down from the market clearing conditions in general equilibrium.

Constrained region. In this region, $V = \lambda\Lambda\theta n$. We have

$$\theta = \frac{\Omega}{\lambda\Lambda}$$
$$\mu = \frac{\sigma_R}{\lambda\Lambda}\left[\Omega\frac{\mu_R - r}{\sigma_R} + \Omega_\Gamma\sigma_\Gamma\right].$$

Households. The micro block describing household behavior can be summarized by the household policy functions $c(a, z, \Gamma)$ and $h(a, z, \Gamma)$. These policy functions are derived from households' Hamilton-Jacobi-Bellman equation. To obtain aggregate household behavior, the cross-sectional distribution $g_t(a, z)$ is used.

Firms. We have

$$Y = e^Z K^\alpha L^{1-\alpha}$$
$$w = (1-\alpha)\frac{Y}{L}$$
$$dK = I - \delta K$$
$$I = \frac{Q-1}{\kappa}K$$
$$\Pi^Q = \frac{(Q-1)^2}{2\kappa}K$$

and finally

$$\mu_R = \frac{Y - wL}{QK} + \frac{1}{Q}\frac{(Q-1)^2}{2\kappa} - \delta + \mu_Q$$
$$\sigma_R = \sigma_Q.$$

Market clearing. We have

$$A + B = 0$$
$$Y = C + I + \frac{\kappa}{2}\left(\frac{I}{K}\right)^2 K$$
$$\theta N = QK,$$

where the last equation characterizes the aggregate market for equity shares.

F.2 Analytical Results

PDE for Ω. Consider the Ansatz

$$V(n,\Gamma) = \Omega(\Gamma)n.$$

Plugging into the HJB, we have

$$\begin{aligned}(\rho+\zeta)\Omega n = \zeta\Lambda n + \Omega_\Gamma n\mu_\Gamma + \frac{1}{2}n\sigma_\Gamma^T\Omega_{\Gamma\Gamma}\sigma_\Gamma \\ + \Omega\Big[rn + \theta n(\mu_R - r)\Big] + \Omega_\Gamma\theta n\sigma_R\sigma_\Gamma + \mu\Big[\Omega n - \lambda\Lambda\theta n\Big].\end{aligned}$$

The quadratic term in θ dropped out. That means the Kuhn-Tucker conditions for θ are now given by

$$\begin{aligned}\Omega n(\mu_R - r) + \Omega_\Gamma n\sigma_R\sigma_\Gamma - \mu\lambda\Lambda n = 0 \\ \mu \geq 0 \\ \Omega n - \lambda\Lambda\theta n \geq 0 \\ \mu(\Omega n - \lambda\Lambda\theta n) = 0.\end{aligned}$$

The first case, with $\mu = 0$, implies

$$0 = \Omega\frac{\mu_R - r}{\sigma_R} + \Omega_\Gamma\sigma_\Gamma.$$

And otherwise, for $\mu \neq 0$, we have

$$\begin{aligned}\theta &= \frac{\Omega}{\lambda\Lambda} \\ \mu &= \frac{1}{\lambda\Lambda}\Big[\Omega(\mu_R - r) + \Omega_\Gamma\sigma_R\sigma_\Gamma\Big].\end{aligned}$$

To verify the Ansatz, we can divide by n, yielding

$$\begin{aligned}(\rho+\zeta)\Omega = \zeta\Lambda + \Omega_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T\Omega_{\Gamma\Gamma}\sigma_\Gamma \\ + \Omega\Big[r + \theta(\mu_R - r)\Big] + \Omega_\Gamma\theta\sigma_R\sigma_\Gamma + \mu\Big[\Omega - \lambda\Lambda\theta\Big],\end{aligned}$$

and this equation must hold for all states (n,Γ). The first observation here is that n drops out everywhere, and so indeed we are left with a PDE over Γ only. So as long as there exists a function $\Omega(\Gamma)$ satisfying the above equation across (n,Γ), then the Ansatz works.

Let's consider the two regions of the state space. In the region of (n,Γ) where the intermediary is unconstrained, with $\mu = 0$, we have

$$\begin{aligned}(\rho+\zeta)\Omega &= \zeta\Lambda + \Omega_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T\Omega_{\Gamma\Gamma}\sigma_\Gamma + \Omega\Big[r + \theta(\mu_R - r)\Big] - \theta\Omega(\mu_R - r) \\ (\rho - r + \zeta)\Omega &= \zeta\Lambda + \Omega_\Gamma\mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T\Omega_{\Gamma\Gamma}\sigma_\Gamma.\end{aligned}$$

Over the constrained region, on the other hand, we have

$$
\begin{aligned}
(\rho - r + \zeta)\Omega &= \zeta\Lambda + \Omega_\Gamma \mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T \Omega_{\Gamma\Gamma}\sigma_\Gamma + \frac{\Omega}{\lambda\Lambda}\Big[\Omega(\mu_R - r) + \Omega_\Gamma \sigma_R \sigma_\Gamma\Big] \\
&= \zeta\Lambda + \Omega_\Gamma \mu_\Gamma + \frac{1}{2}\sigma_\Gamma^T \Omega_{\Gamma\Gamma}\sigma_\Gamma + \Omega\mu.
\end{aligned}
$$

Together, these equations constitute a PDE for Ω.

F.3 Quantitative Results

In progress and coming soon.