Journal of Economic Dynamics & Control 91 (2018) 349–368

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Solving an incomplete markets model with a large cross-section of agents☆

Thomas M. Mertens[a,*], Kenneth L. Judd[b]

[a] Federal Reserve Bank of San Francisco 101 Market Street, Mail stop 1130, San Francisco, CA 94705, USA
[b] Hoover Institution, Stanford University 434 Galvez Mall, Stanford, CA 94305, USA

ARTICLE INFO

Article history:
Received 13 April 2017
Revised 1 December 2017
Accepted 22 January 2018
Available online 5 February 2018

JEL classification:
C60
C02
E44
G12

Keywords:
Perturbation methods
Incomplete markets
Heterogeneity

ABSTRACT

This paper shows that perturbation methods can be applied to a DSGE model with incomplete markets and a finite but arbitrarily large number of heterogeneous agents. We develop a simple but general solution technique that handles many state and choice variables for each agent and thus has an extremely high-dimensional state space. The method is based on perturbations around a point at which the solution is known. The novel idea is to exploit the symmetry of the problem to overcome the curse of dimensionality. We use the analysis to demonstrate the impact of heterogeneity on macroeconomic quantities and the pricing of risk. Furthermore, we set our technique apart from standard methods used in the literature.

1. Introduction

A large body of literature in finance and macroeconomics makes the simplifying assumption that aggregate variables are determined by the behavior of a representative agent. In reality, different people earn different incomes, have different talents, and hold different expectations. For this heterogeneity to be reflected in aggregate outcomes, incompleteness of asset markets is essential. In reality, substantial amounts of idiosyncratic risk can only be partially insured. Labor income risk serves as a prime example. Modeling this type of idiosyncratic risk permits a more stringent test of our current economic theory since we can use information about the entire distribution of economic outcomes across the population.

This paper proposes a numerical method and solves an incomplete markets model with a finite but arbitrarily large number of households. The algorithm is based on perturbation methods and thus is simple to apply and particularly well suited for economies in which the state space is large. We demonstrate that the solution around the deterministic steady state, the standard point of approximation for perturbation methods, is highly symmetric. Therefore, despite the state space consisting of distributions of state variables across households, computing the solution remains manageable.

☆ We thank Dave Backus, Fernando Duarte, Xavier Gabaix, Per Krusell, Sydney Ludvigson, Hanno Lustig, Semyon Malamud, Michael Reiter, Tony Smith, Stijn van Nieuwerburgh, Kjetil Storesletten, and the audiences at Yale, New York Fed, NYU Stern, Wharton, Duke, 2011 SCE conference, SITE 2011, the Chicago Booth-Deutsche Bank symposium, the Cowles GE conference, the SED, and the EFA for helpful comments. The views expressed in this paper are solely those of the authors and do not necessarily represent those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

* Corresponding author.
E-mail addresses: thomas.mertens@sf.frb.org (T.M. Mertens), judd@hoover.stanford.edu (K.L. Judd).

https://doi.org/10.1016/j.jedc.2018.01.025
0165-1889/

We analyze a dynamic stochastic general equilibrium (DSGE) model with aggregate risk in production and an endogenous capital stock. A firm employs a Cobb-Douglas production technology to produce a single consumption good. Future total factor productivity is uncertain due to aggregate risk. Households maximize expected discounted utility given by an additively separable utility function with constant relative risk aversion in consumption.

We add idiosyncratic shocks to labor income that households cannot insure against. Households only trade claims to capital, which renders markets incomplete. As a result, equilibrium outcomes respond to idiosyncratic conditions. Households hold different levels of capital, which translates into inequality of wealth and consumption.

The analysis of this model presents a difficult problem. Ultimately, we want to be able to study the interaction of households' choices and asset returns across states of the economy as well as the distribution of capital holdings and consumption across the population. Therefore, we need a solution method that solves for individual behavior and aggregate variables including asset prices as a function of the entire distribution of economic conditions. But, in turn, this distribution is affected by all individuals' behavior. In other words, the state space might contain several *distributions* of variables across households.

We lay out the mathematical structure of equilibrium conditions. We scale all standard deviations by a perturbation parameter σ such that a value of $\sigma = 1$ corresponds to our model of interest. Setting this parameter to zero allows us to study an auxiliary economy where we eliminate all uncertainty, the deterministic economy.

To compute the equilibrium of our economy, we develop a solution technique for models with many heterogeneous agents and incomplete markets based on perturbation methods. Perturbation methods build an approximation of the optimal policies as functions of the state variables based on Taylor expansions. The first step is to find a special case of the model in which the solution is known. Our model possesses a well-defined deterministic steady state available in closed form around which we expand optimal policies with respect to all state variables. At the point of expansion, all households are identical in all respects and thus the distribution of capital is degenerate. Having pinned down the deterministic steady state, we build a Taylor expansion with respect to all state variables. We know that equilibrium outcomes are functions of the state space. Thus we expand the deterministic economy in all state variables. But since we allow for an arbitrarily large number of households, we also have an arbitrarily large number of state variables.

The novel idea lies in exploiting the symmetry of decision rules across households. If two households are identical in their objectives, they respond identically to the same economic conditions. For example, starting out from a case where all households have identical state variables, a marginal increase in household one's wealth will impact the decision of household two the same way that a marginal increase in household two's wealth would impact household one's decision. Exploiting this symmetry, we solve for the decision rules of all households as a function of the entire distribution of individual states. As a result, we only need to expand the optimality condition for one household around the deterministic steady state.

A second symmetry arises from the fact that all other households are "anonymous" to a household in the following sense. At the deterministic steady state, an increase in wealth of household two impacts household one the same way as an equal increase in household three's wealth would. As a result, many coefficients in the Taylor series are identical. In an economy where the state space consists of the distribution of capital, only two coefficients need to be computed in the first-order approximation, the response to a marginal increase in one's own capital and the response to a marginal increase in some other household's capital.

Recognizing these symmetries simplifies the problem substantially. First, we only need to expand the optimality conditions of one household. Second, we only need to compute few coefficients to reconstruct the entire Taylor series. We show that this is not an approximation but arises from the structure of the problem. This is true independent of the number of households and we can thus deal with an arbitrarily large cross-section of households.

The last step of the algorithm makes the transition from the deterministic to the stochastic economy. Since shocks are part of the state space, the previous expansion delivers equilibrium reactions to known, deterministic changes in these state variables. For example, the previous expansion would compute the asset price reaction if the next period's productivity was above its steady-state level. To move to the stochastic economy, we integrate over all possible realizations of the shocks and weight them by their probability. From this logic it follows immediately that we need a higher-order expansion. If we were to resort to a first-order approximation, integrating over the first-order approximation would not affect equilibrium behavior since a linear solution is certainty-equivalent. Higher-order expansions bring in the effects of uncertainty. A second-order approximation reflects the effect of the variance of shocks, a cubic approximation additionally takes the third moment into account, and so on.

Our solution method is asymptotically valid and converges to the true solution within the radius of convergence. By adding higher moments, we can construct better approximations to the true policy function. In practice, of course, convergence is not complete. Therefore, we discuss a means of testing the accuracy of our solution. We plug our approximation into the equilibrium conditions to check its optimality.

We discuss the generality of the solution method. It applies whenever equilibrium or optimality conditions for a competitive equilibrium or dynamic programming problems imply that the choice variables are smooth functions of state variables. The dynamic programming problem or competitive equilibrium can feature arbitrarily many state variables and is thus interesting for a large set of economic applications. We also discuss implementation of constraints as well as the addition of portfolio choice to our economy.

Next, we demonstrate the results from our solution method. First, we confirm previous research in finding that heterogeneity has an effect on the steady-state level of capital. Since households face idiosyncratic risk, they respond by building up a buffer stock of precautionary savings. With aggregate risk this channel is enforced. Due to the utility specification featuring constant relative risk aversion, households increase their capital holdings further to guard against uncertain returns to capital.

Furthermore, we show the risk factors for this incomplete markets economy. We present the expansion of a stochastic discount factor in closed form up to a given order. The expansion consists of standard risk factors such as total factor productivity and, if stochastic, its variance but also risk due to incompleteness of markets such as the variance of idiosyncratic labor income risk. This last factor only appears because there are missing markets that prevent households from insuring against their idiosyncratic conditions.

Finally, we state another economy taken from Panageas (2011) where we can apply perturbation methods. Since this example allows for an explicit solution, we can compare our solution method to a standard technique that replaces the actual law of motion with a function of moments of the distribution of state variables. We solve this asset pricing economy explicitly using the linearized law of motion and the solution method of this paper. We find our technique shows superior performance.

We make two main caveats to our analysis. First, the model is not a state-of-the-art calibrated DSGE model but rather serves as an illustration of the method. We chose this model because it allows our algorithm to be compared directly to alternatives discussed in the literature (see Den Haan et al., 2010). However, we discuss generalizations in Section 6. Second, we do not prove the existence of an ergodic distribution. Our method introduces an arbitrarily small penalty for deviating from an average level of capital for the auxiliary deterministic economy. In this case, the long-run distribution of capital is determined but we do not extend the proof to the stochastic economy.

This paper contributes to a growing literature on introducing heterogeneity into economic models and therefore relates to several strands of research. After the seminal works of Bewley (1977) and Aiyagari (1994), the literature has focused on idiosyncratic risk with aggregate shocks. In finance, a literature on asset pricing under heterogeneous expectations has evolved that takes changes in the wealth distribution into account (e.g. see Chiarella and He (2001) or Chiarella et al., 2006). First, in special cases one might be able to find closed-form solutions as in Heathcote et al. (2014) and Moll (2009). Another promising idea is to use a multiplier approach to characterize features of the distribution of state variables across the population as in Chien et al. (2010) and Chien and Lustig (2010). Other papers make simplifying assumptions on the number of agents and the number of possible shocks, as in Dumas and Lyasoff (2012). Special cases with closed-form solutions can be used as a starting point for the expansion.

Most of the literature, however, is concerned with approximations. One idea is to replace the distribution of wealth by aggregate wealth only when calculating the equation of motion for aggregate variables. The most prominent method was developed in Krusell and Smith (1998) and inspired methods in the subsequent literature, for example in Storesletten et al. (2007), Gomes and Michaelides (2008), Gomes and Schmid (2010), and Favilukis et al. (2017) where aggregate states and prices might influence the equation of motion. Alternatively, one might work with a limited history of shocks as in Lustig and van Nieuwerburgh (2010). Since we are particularly interested in the effect of distributions on equity prices and the effect on new financial securities, this approximation method is not appropriate for this research project.

Recently, alternative solution methods for models with heterogeneous agents have been developed in Den Haan et al. (2010), Den Haan (2010), Judd et al. (2009), and Judd et al. (2010). Den Haan et al. (2010) and Den Haan (2010) parameterize the distribution of state variables. Feng, Miao, Peralta-Alva and Santos (2014) approximate the equilibrium on a lower-dimensional space. This paper develops a technique that does not require the specification of a class of distributions. Compared with Judd et al. (2009) and Judd et al. (2010), the method in this paper allows us to study as many agents as desired whereas the number of agents is limited in their method. Reiter (2009) combines projection with a perturbation. Our method applies to models with many state variables and choices for each individual and is asymptotically valid. Furthermore, the usual differences between perturbation and projection apply: while projection methods were designed as global solution methods able to handle non-differentiabilities, our method returns quasi-analytical expressions, can handle many state variables, and is easy to implement.

Our method builds on perturbation methods as used in Jin and Judd (2002), Judd and Guu (2000), Judd (2002), Hassan and Mertens (2017), Mertens (2009), Fernández-Villaverde and Rubio-Ramírez (2006), and Garlappi and Skoulakis (2010). This paper, however, is not the first that attempts to use perturbation methods to analyze general equilibrium models with a large cross-section of agents. Alternative ideas have been explored by Preston and Roca (2007) and Kim et al. (2010). This line of work starts by restricting the state space for prices from the outset. More specifically, Kim et al. (2010) use the process for prices from the representative agent economy, which amounts to assuming the state space depends on aggregate capital only. Preston and Roca (2007) allow for expansions in moments of the distribution. By contrast, this paper is the first to recognize the symmetry of the problem and build a solution method that exploits it without limiting the state space.

2. The stochastic neoclassical growth model with heterogeneous households

This section describes the stochastic neoclassical growth model with aggregate risk in productivity and idiosyncratic labor income risk.

2.1. Households

A finite number of I households lives for an infinite number of periods indexed by t. Households are each endowed with one unit of time that they devote towards labor inelastically.[1] While they are identical in their preferences, households differ in their productivity. Each period, a household receives an idiosyncratic shock to its productivity and thus its labor income. Markets are incomplete in that there is no asset available that lets households insure against their individual productivity. Therefore, households can only partially insure against this shock by holding savings to buffer the impact of shocks on consumption. The single tradable contract consists of claims to capital, which is risky due to aggregate productivity shocks.

Each household i has rational expectations and chooses streams of consumption and capital holdings to maximize expected discounted utility

$$\max_{c_t^i, k_{t+1}^i} \sum_{t=0}^{\infty} \beta^t E_0\left[u_c(c_t^i)\right] \qquad i = 1, \ldots, I \tag{1}$$

where $u(\cdot)$ is the utility function, β is the time discount factor, and c_t^i and k_{t+1}^i are household i's choices of consumption and capital holdings in period t.

Households maximize utility subject to their budget constraint

$$c_t^i + k_{t+1}^i = (1 - \delta + r_t^k)k_t^i + w_t \psi_t^i. \tag{2}$$

The rate of return on capital before depreciation is denoted by r_t^k, the rate of depreciation by δ, and wages by w_t. The shock to individual productivity is denoted by ψ_t^i, which is independent and identically distributed across households. It follows a stochastic process that we specify in our calibration.

To keep a concise notation, we introduce capital case letters for aggregate quantities of consumption and capital

$$C_t = \sum_{i=1}^{I} c_t^i \qquad K_t = \sum_{i=1}^{I} k_t^i.$$

2.2. Technology

Aggregate capital and labor enter the production process for the single consumption good that is produced with Cobb-Douglas technology. Output is given by $Y = f(K, L, z) = zK^{\alpha}L^{1-\alpha}$ where z denotes the shock to total factor productivity, K aggregate capital, L aggregate labor demand, and α the capital share of output.

Firms maximize output net of costs for capital and labor. Given constant returns to scale in the production of the consumption good, capital and labor pay their marginal products reflected in returns to capital and wages

$$\begin{aligned} r_t^k &= \alpha z_t K_t^{\alpha-1} L_t^{1-\alpha} \\ w_t &= (1-\alpha) z_t K_t^{\alpha} L_t^{-\alpha} \end{aligned} \tag{3}$$

Due to the shocks to total factor productivity, the returns to capital and wages are risky.

2.3. Definition of equilibrium

Households' first-order conditions determine their optimal choices of consumption

$$u_c'(c_t^i) = \beta E_t\left[(1 - \delta + r_{t+1}^k) u_c'(c_{t+1}^i)\right]. \tag{4}$$

Their budget constraint pins down the amount of capital holdings. The aggregate resource constraint

$$C_t + K_{t+1} - (1-\delta)K_t = Y_t \tag{5}$$

shows how current output and depreciated capital can be used for consumption or next period's capital stock. The derivation follows from households' budget constraints and the market clearing conditions for capital by which demand k_{t+1}^i aggregates to supply K_{t+1}.

The state space of this economy consists of the set of individual capital holdings of each of the I households, the level of their individual productivity, and aggregate shocks. In other words, the distribution of capital and labor income shocks across households is part of the state space. Furthermore, we need to keep track of two aggregate state variables: the state of total factor productivity and preference shocks. As a result, the state space is extremely high dimensional if we allow for a large cross-section of households.

To reflect the distinction between individual and aggregate state variables, we introduce two separate variables. $\mathbf{X}$ denotes the matrix of individual state variables k^i and ψ^i. $\mathbf{z}$ is the vector of aggregate state variables. We also specify whether the state space belongs to the DSGE model (which corresponds to $\sigma = 1$) or an auxiliary economy in which all shocks

[1] We discuss a variety of extensions in Section 6.

have a proportionately smaller standard deviation, scaled by the factor σ. In particular, the auxiliary deterministic economy, characterized by $\sigma = 0$, is of interest and will be discussed separately below. We denote an element of the state space by $\mathbb{S}_t = (\mathbf{X}_t, \mathbf{z}_t, \sigma) \in \mathbb{R}^{2I+1+1}$. Lowercase bold variables stand for vectors and uppercase bold font letters for matrices.

An equilibrium consists of price, choice, and transition functions over the state space such that all equilibrium conditions are satisfied. To simplify notation, we stack all individual choices for all households in the optimal choice functions $\mathbf{C}$ in a matrix, all price functions into a vector-valued function $\mathbf{p}$, and all transition functions into $\boldsymbol{\tau}$. Appendix A.1 shows the structure of these functions in detail. Furthermore, we denote a collection of these functions $(\mathbf{C}, \mathbf{p}, \boldsymbol{\tau})$ by $\mathbf{B}$. In order to write down the law of motion, we collect all innovations to stochastic processes in a variable $\mathbf{e}_{t+1} = (\{\theta^i_{t+1}\}^I_{i=1}, \varepsilon_{t+1})$, where θ^i are the innovations to idiosyncratic labor income shocks ψ^i and ε innovations to total factor productivity z.

We collect all equilibrium conditions in a single operator G

$$G\mathbf{B}(\mathbb{S}_t) = \begin{pmatrix} E_t[g_1(\mathbb{S}_t, \mathbf{C}, \mathbf{p}, \boldsymbol{\tau}, \mathbf{e}_{t+1})] \\ g_2(\mathbb{S}_t, \mathbf{C}, \mathbf{p}, \boldsymbol{\tau}, \mathbf{e}_{t+1}) \end{pmatrix}. \tag{6}$$

The operator takes the state variables as its inputs along with the collection of choice variables $\mathbf{C}$, price functions $\mathbf{p}$, and transition functions for state variables and shocks $\boldsymbol{\tau}$ as functions of the state space. The operator g^1 consists of all I Euler equations stacked in one vector. g^2 stacks individual budget constraints, equations of motion for total factor and idiosyncratic labor productivity, as well as market clearing conditions and definitions of aggregate variables. Appendix A.2 shows the structure of these equilibrium conditions.

Definition 1 (Definition of equilibrium for economy)**.** An equilibrium of the economy is a collection of choice ($\mathbf{C}$), price ($\mathbf{p}$), and transition functions ($\boldsymbol{\tau}$) such that $G\mathbf{B}(\mathbb{S}_t) = 0$.

These equilibrium relationships ensure that all households make their choices optimally and markets clear.

By varying the parameter σ, which scales the standard deviation of all shocks proportionately, we can produce a range of different economies that vary by their amount of uncertainty. For the numerical method, this comparative statics exercise turns out to be convenient. In particular, for the case where the standard deviation is zero, there is one point in the state space at which we can solve the model explicitly.

Definition 2 (Deterministic steady state)**.** A deterministic steady-state is a point in the state space $\mathbb{S}_0 = (\mathbf{X}_0, \mathbf{z}_0, 0)$ such that each household's first-order conditions are satisfied and all prices and quantities remain constant over time.

3. Numerical method

This section proposes a numerical technique based on perturbation methods to solve the model of the previous section. Perturbation methods derive a higher-order approximation to the solution around the deterministic steady state. The key insight in this paper is that they are very well suited for incomplete market models with a large cross-section of households in which heterogeneity arises from idiosyncratic shocks. We show that the model has symmetry properties at the deterministic steady state that perturbation methods can exploit. Hence the solution of the model remains tractable despite the high dimensionality of the state space.

3.1. Deterministic steady state

To solve for the deterministic steady state, we set the standard deviation of all shocks to zero. Since households are heterogeneous only with respect to their idiosyncratic labor income shocks in our model, the deterministic steady state features identical households and no heterogeneity, conditional on having the same initial capital. The distribution of capital, however, is not pinned down by equilibrium conditions. To see this, we look at the steady-state condition

$$1 - \delta + f_K(K_t, L_t, 0) = \frac{1}{\beta} \tag{7}$$

which only depends on aggregate capital and not on its distribution across households.

To pin down the distribution of capital, we modify the auxiliary economies with less uncertainty by imposing a small penalty for deviations from average capital holdings. Specifically, we adjust the utility function

$$u(c^i_t, k^i_t) = u_c(c^i_t) - \frac{\nu}{2}(1 - \sigma)(k^i - \bar{k})^2 \tag{8}$$

where $\bar{k}$ is the average level of individual capital holdings. As a result, the deterministic steady state now ensures that all households have the same amount of capital $k^i = \bar{k}$.

Importantly, the parameter ν in the penalty function can be arbitrarily small such that it has a negligible effect on equilibrium quantities in the model of the previous section. Furthermore, we scale the penalty term by a factor $1 - \sigma$ so that it appears in the deterministic steady state (when $\sigma = 0$) but not in the model of interest where $\sigma = 1$.[2]

[2] When using computer algebra systems such as Mathematica, one has an alternative option of leaving ν unspecified as a free parameter. Once all derivatives have been computed as a function of ν, we can set ν to zero.

3.2. Symmetry in higher-order expansions

Computing a higher-order Taylor series for the equilibrium policy functions, equilibrium quantities, and prices is essential to our solution method for two reasons. First, heterogeneity manifests its impact only in higher-order terms and second, so does stochasticity. To compute high-order derivatives, a high-precision arithmetic might be necessary as shown in Swanson et al. (2005).

A Taylor series expansion of high order serves as a good approximation to equilibrium outcomes. If the conditions of the implicit function theorem are met, as in our economy of the previous section, the approximation converges within the radius of convergence when we increase the order of the expansion. In practice, of course, we have to truncate the Taylor series at a finite level. But the stage at which we stop can be endogenous to the accuracy of the solution.

Taylor expansions are at the heart of perturbation methods and we state them using the standard multi-index notation (see, for example, Taylor (2010) (p. 3–16) for a reference). We denote a C-tuple of integers by $\boldsymbol{\iota}_i = (\iota_{i1}, \iota_{i2}, \ldots, \iota_{iC})$ to index individual states for household i. Let $\mathbf{I} = \{\boldsymbol{\iota}_1, \ldots, \boldsymbol{\iota}_I\}$ be the collection of such indices for all households. Furthermore, $\mathbf{j}$ is a Z-tuple of integers to index all aggregate shocks. The order of differentiation is then given by $\|\mathbf{I}\| + |\mathbf{j}| + k$ where $\|\mathbf{I}\| = \sum_{i=1}^{I}\sum_{\chi=1}^{C}\iota_{i\chi}$ and $|\mathbf{j}| = \sum_{\zeta=1}^{Z}\mathbf{j}_\zeta$. We also define the product of all entries $\mathbf{I}! = \prod_{i=1}^{I}\prod_{\chi=1}^{C}\iota_{i\chi}$ and $\mathbf{j}! = \prod_{\zeta=1}^{Z}\mathbf{j}_\zeta$. A concise notation for a derivative of choice $\mathbf{C}$ at the deterministic steady state reads

$$\mathbf{C}^{(\mathbf{I},\mathbf{j},k)}(\mathbb{S}_0) = \left(\prod_{i=1}^{I}\prod_{\chi=1}^{C}\partial_{\iota_{i\chi}}\right)\left(\prod_{\zeta=1}^{Z}\partial_\zeta\right)\partial_k\mathbf{C}(\mathbb{S}_0) \tag{9}$$

where $\partial_{\iota_{i\chi}} = \partial^{\iota_{i\chi}}/\partial\mathbf{X}_{i\chi}^{\iota_{i\chi}}$, $\partial_\zeta = \partial^{j_\zeta}/\partial\mathbf{z}_\zeta^{j_\zeta}$, and $\partial_k = \partial^k/\partial\sigma^k$. Finally, we define the monomials in the Taylor series accordingly. Let $(\mathbf{X}-\mathbf{X}^0)^{\mathbf{I}} = \prod_{i=1}^{I}\prod_{\chi=1}^{C}(\mathbf{X}_{i,\chi}-\mathbf{X}_{i\chi}^0)^{\iota_{i\chi}}$ and analogously $(\mathbf{z}-\mathbf{z}^0)^{\mathbf{j}} = \prod_{\zeta=1}^{Z}(\mathbf{z}_\zeta-\mathbf{z}_\zeta^0)^{\mathbf{j}_\zeta}$.

Once we know the derivatives at a specific point, we can recover the choice variable of the ϖ-th choice of household i from the Taylor series

$$\mathbf{C}_{\varpi}^{i}(\mathbf{X},\mathbf{z},\sigma) = \sum_{o=0}^{\infty}\sum_{\|\mathbf{I}\|+|\mathbf{j}|+k=o}\frac{1}{\mathbf{I}!\cdot\mathbf{j}!\cdot k!}\mathbf{C}^{(\mathbf{I},\mathbf{j},k)}(\mathbb{S}_0)(\mathbf{X}-\mathbf{X}^0)^{\mathbf{I}}(\mathbf{z}-\mathbf{z}^0)^{\mathbf{j}}\sigma^k. \tag{10}$$

Note that Eq. (10) can serve as an approximating function over the state space.

To obtain derivatives at the deterministic steady state, we employ perturbation methods. Ultimately, we are interested in a solution to Eq. (6). Perturbation methods tell us to take derivatives of each equation with respect to all state variables successively and evaluate the resulting equations at the deterministic steady state. By the chain rule of differentiating Eq. (6), we obtain equations for the derivatives of the policy function at the deterministic steady-state which we can then solve for. Plugging them into Eq. (10) results in an approximation of the policy function. To give an example, take a first-order condition of the form (4) of household $i \neq 1$. We take the derivative with respect to the first individual state variable x_1^1

$$u_c''\frac{\partial c_t^i}{\partial x_1^1} - \beta\frac{\partial r_t^k}{\partial x_1^1}u_c' - \beta(1-\delta+r_t^k)u_c''\frac{dc_{t+1}^i}{dx_1^1} = 0 \tag{11}$$

where arguments of the utility function are suppressed. For clearer exposition, we evaluated the expression at $\nu = 0$, and did not expand the derivative dc_{t+1}^i/dx_1^1 using the chain rule.

More generally, perturbation methods require us to take derivatives of every equilibrium equation g_i^1 or g_i^2

$$\begin{aligned}\frac{dg_i}{dx_1^1} = {} & \frac{\partial g_i}{\partial x_1^1} + \frac{\partial g_i}{\partial\mathbf{C}_t}\frac{\partial\mathbf{C}_t}{\partial x_1^1} + \frac{\partial g_i}{\partial\mathbf{p}_t}\frac{\partial\mathbf{p}_t}{\partial x_1^1} + \frac{\partial g_i}{\partial\boldsymbol{\tau}_t}\frac{\partial\boldsymbol{\tau}_t}{\partial x_1^1} + \frac{\partial g_i}{\partial\mathbf{X}_{t+1}}\frac{\partial\mathbf{X}_{t+1}}{\partial x_1^1} + \frac{\partial g_i}{\partial\mathbf{z}_{t+1}}\frac{\partial\mathbf{z}_{t+1}}{\partial x_1^1} + \frac{\partial g_i}{\partial\mathbf{C}_{t+1}}\frac{d\mathbf{C}_{t+1}}{dx_1^1}\\ & + \frac{\partial g_i}{\partial\mathbf{p}_{t+1}}\frac{d\mathbf{p}_{t+1}}{dx_1^1} + \frac{\partial g_i}{\partial\boldsymbol{\tau}_{t+1}}\frac{d\boldsymbol{\tau}_{t+1}}{dx_1^1}\end{aligned} \tag{12}$$

and plug in steady-state values.

Now we use the fact that the functional form for g_i^1 is known and so are all partial derivatives. They can be obtained by taking the derivatives of the equilibrium conditions (in the example of the Euler equation, these derivatives entail differentiating marginal utilities). When evaluating them at the deterministic steady state, the only remaining variables in the differentiated equilibrium conditions are the derivatives of the optimal policies $\mathbf{C}$, prices $\mathbf{p}$, and transition functions $\boldsymbol{\tau}$ evaluated at the deterministic steady state. Solving for these derivatives delivers the coefficients in the Taylor series of the optimal policies. So far, we described the standard use of perturbation methods.

The key innovation of this paper lies in recognizing the symmetry of the problem. In principle, we would have to start with household one, differentiate its first-order conditions with respect to each household's state variables, move to household two and so on. However, we show that the actual computation of these derivatives can be much simplified.

Mathematically, we can state the symmetry of the model as follows. If $\mathbf{X}^{i\leftrightarrow j}$ denotes the matrix of state variables where we exchange the state variables of household i with household j and vice versa and do the same for policy functions $\mathbf{C}^{i\leftrightarrow j}$, then we can express the symmetry requirement as

$$g_k^i(\mathbf{X},\mathbf{z},\sigma,\mathbf{C},\mathbf{p},\boldsymbol{\tau},\mathbf{e}_{+1}) = g_k^j(\mathbf{X}^{i\leftrightarrow j},\mathbf{z},\mathbf{C}^{i\leftrightarrow j},\mathbf{p},\boldsymbol{\tau},\mathbf{e}_{+1}) \qquad k=1,2. \tag{13}$$

Specifically, we explore symmetry along two dimensions.

First, all derivatives with respect to state variables of other households than the one whose policy we approximate are identical, i.e.,

$$\frac{dc^i}{dx_1} = \frac{dc^1}{dx_2} \quad \forall i > 1.$$

For example, when differentiating household one's first-order condition, there are only two different coefficients in the first-order expansion:[3] The derivative with respect to the household's own state variables and those of any other household. These two numbers are sufficient to build the entire Taylor series since coefficients on other households' state variables have to be the identical given the same fundamentals.

Second, we only have to take derivatives of first-order conditions for household one. Household two's first-order conditions look identical and thus lead to identical coefficients

$$\frac{dc^i}{dx_j} = \frac{dc^j}{dx^i}.$$

The symmetry here is that household one's response to a marginal increase in household two's state variable is the same as household two's response to a marginal increase in household one's state variable. This carries over to all derivatives.

Exploiting this symmetry, a first-order approximation requires two coefficients to be computed for each state variable. The first coefficient returns the change in policy of household one in response to a change in her own capital holdings. The second coefficient asks for household one's reaction in response to a change in the state variable by an arbitrary household two. For the second-order term, the system becomes slightly more complex. For each state variable, we compute how the two coefficients from the linear system change in response to a change in state variables. More specifically, we need to compute four values: (1) a change in the first coefficient in response to a change in household one's state variable, (2) a change in the first coefficient in response to a change in another household's state variable, (3) a change in the second coefficient in response to a change in household two's state variable, and (4) a change in the second coefficient in response to a change in a third household's state variable. Increasing the order, we have quadratic growth in the number of coefficients. Solving the system of equations is straightforward: From the second order on, the system of equations for the unknown coefficients is linear.

A first-order approximation implements standard linearization which is not sufficient for our purposes. Due to linearity, heterogeneity does not affect aggregate equilibrium outcomes because, under these rules, the average choice is the choice of the average household. Heterogeneity only enters through higher-order terms, starting with a second-order approximation. For the same reason, stochasticity impacts equilibrium only through higher-order terms. The first-order approximation is certainty equivalent while higher-order terms add the effects of variance, third and higher moments.

3.3. Uncertainty

Having obtained a higher-order approximation for the deterministic economy, we move towards the model of interest by introducing uncertainty. We accomplish the transition by varying the perturbation parameter σ.

Taking a first-order expansion with respect to the perturbation parameter produces coefficients that are all zero. The reason lies in the fact that the first-order expansion of the standard deviation introduces shocks only into the linearized, and thus certainty equivalent, economy. Hence, uncertainty does not come into play apart from the realization of shocks. Only through second- and higher-order terms do we recover the solution to the stochastic system. The second-order term introduces shocks into the quadratic economy. This approximation is no longer certainty equivalent and uncertainty takes effect. To be more precise, the second-order term introduces a constant effect due to the variance of shocks stemming from Jensen's inequality, the third-order term recovers the reaction to skewness and time- or state-dependent variation in the variance of shocks, and so on.

We can interpret the way uncertainty enters the equilibrium as effectively altering the coefficients in the Taylor series. Building the expansion with respect to the standard deviation of shocks effectively changes the coefficients of the Taylor series for the deterministic system. To see this, we rewrite Eq. (10) in the form

$$\mathbf{C}^i_{\varpi}(\mathbb{S}) = \sum_{o=0}^{\infty} \sum_{\|\mathbf{I}\|+|\mathbf{j}|=o} \frac{1}{\mathbf{I}!\cdot\mathbf{j}!} \left(\sum_{k=0}^{\infty} \frac{\mathbf{I}!\cdot\mathbf{j}!}{\mathbf{I}!\cdot\mathbf{j}!\cdot k!} \mathbf{C}^{(\mathbf{I},\mathbf{j},k)}(\mathbb{S}_0)\sigma^k \right) (\mathbf{X}-\mathbf{X}^0)^{\mathbf{I}} (\mathbf{z}-\mathbf{z}^0)^{\mathbf{j}}.$$

The rearrangement demonstrates that the expansion of the stochastic system looks just like the deterministic system except that the coefficients (in brackets) contain a "correction term" for the stochasticity of the function.

We can see this term graphically as depicted in Fig. 1. In the second order, the function shifts while the third-order term would also tilt the function, and even higher orders change its curvature.

[3] To make the exposition clearer, the following logic assumes a single distribution of individual state variables as the state space.

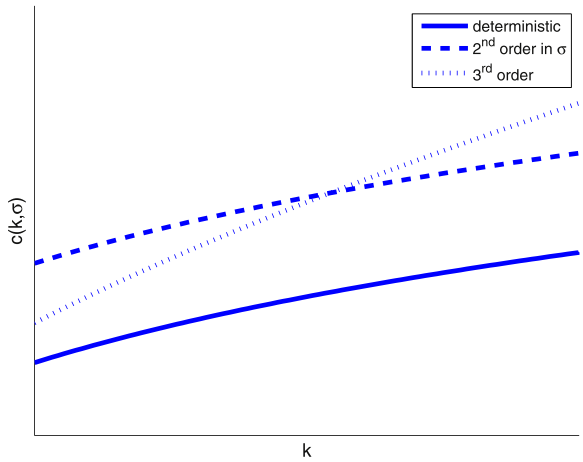

Fig. 1. Perturbation methods build an approximation in state variables around the deterministic steady state (thick solid line). The expansion with respect to the standard deviation shifts (second order) and tilts this line (third order).

3.4. The law of motion

Perturbation methods deliver a polynomial representation of the approximation. The law of motion is no exception to this rule. In our approximate solution, the equation of motion is not only a function of an aggregate statistic of state variables but the entire distribution. With every increase in the order of approximation, our solution method includes the corresponding moments from the distribution of state variables. In this sense, the solution method proposes a set of approximating statistics with which to approximate policy functions. As an additional feature, a better approximation adds moments to the previous approximation without needing to recompute previous approximations.

The first-order expansion results in a law of motion for aggregate capital of the form

$$\begin{aligned}\tfrac{1}{I}K_{t+1} \approx k_0 &+ (k^1_{\mathbf{X}_{11}} - k^1_{\mathbf{X}_{21}} + Ik^i_{\mathbf{X}_{21}})(\mathbf{X}_1 - \bar{\mathbf{X}}_{11}) + (k^1_{\mathbf{X}_{12}} - k^1_{\mathbf{X}_{22}} + Ik^i_{\mathbf{X}_{22}})(\mathbf{X}_2 - \bar{\mathbf{X}}_{12}) \\ &+ Ik^1_{\mathbf{z}_1}(\mathbf{z}_1 - \bar{\mathbf{z}}_1) + Ik^1_{\mathbf{z}_2}(\mathbf{z}_2 - \bar{\mathbf{z}}_2) + Ik^1_{\mathbf{z}_3}(\mathbf{z}_3 - \bar{\mathbf{z}}_3) + \ldots\end{aligned}$$

where $\mathbf{X}_j = \frac{1}{I}\sum_{i=1}^{I}\mathbf{X}_{ij}$, $\bar{\mathbf{X}}_{1.}$ is the value of the state variable at the deterministic steady state, and $k^1_{\mathbf{X}_{11}} = \frac{\partial k^1}{\partial \mathbf{X}_{11}}$ (see Appendix B for a derivation).

In our model in the previous section, aggregate capital in the next period depends, to a first order, on aggregate capital and aggregate shocks today. As discussed before, heterogeneity does not impact aggregate variables in the first-order term.

The second-order approximation depends both on the cross-sectional variance of individual state variables as well as a quadratic term in the aggregate state variable. It thus depends not only on the cross-section but also nonlinearly on the time-series variation of aggregate quantities.

The above expression is not particular to the law of motion. Any function that depends identically on household choices, will be approximated in this fashion. Appendix B contains details.

3.5. Distribution of equilibrium variables

Given our approximation method, we can compute the distribution of any equilibrium outcome or nonlinear function thereof. Therefore, we combine perturbation methods with a nonlinear change of variables.[4] For example, from capital holdings and idiosyncratic labor income shocks, we can compute the distribution of wealth or human wealth.

Suppose we have some economic variable of interest that is a nonlinear function $h(\mathbf{X}, \mathbf{z}, \mathbf{C}, \mathbf{p}, \boldsymbol{\tau})$ of the state variables and choices, which we approximate by a Taylor expansion. The coefficients can be computed as follows

$$\frac{dh}{dx^i_s} = \frac{\partial h}{\partial \mathbf{X}}\frac{d\mathbf{X}}{dx^i_t} + \frac{\partial h}{\partial \mathbf{z}}\frac{d\mathbf{z}}{dx^i_t} + \frac{\partial h}{\partial \mathbf{C}}\frac{d\mathbf{C}}{dx^i_t} + \frac{\partial h}{\partial \mathbf{p}}\frac{d\mathbf{p}}{dx^i_t} + \frac{\partial h}{\partial \boldsymbol{\tau}}\frac{d\boldsymbol{\tau}}{dx^i_t} \tag{14}$$

and analogously for other state variables. All partial derivatives of h are given through its functional form while the derivatives of state variables, choices, and prices were previously computed through perturbation methods.

The computation of the coefficient is trivial once we make the observation that the first term is given by the derivative of h (which is given) and the second one has already been computed in the previous approximation. Thus, computing the distribution of any variable of interest within the economy is not more intricate than computing the distribution of capital.

[4] Judd (2002), Fernández-Villaverde and Rubio-Ramírez (2006), and Mertens (2009) explain nonlinear changes of variables in conjunction with perturbation methods.

3.6. Accuracy

The solution method gives rise to a natural way to check for its accuracy. The equilibrium conditions are satisfied when the functional G in Eq. (6) returns zero values for all of its components. Since we have asymptotic validity of the solution method, we specify a tolerance as a threshold for the error. Once the error is below the tolerance in some norm, we terminate the approximation process.

To get a meaningful measurement for the error, it makes sense to normalize the optimality conditions such that they are unit-free. For example, we rewrite the Euler Eq. (4) in the form

$$\beta E_t\left[(1-\delta+r^k_{t+1})\frac{u'_c(c^i_{t+1})}{u'_c(c^i_t)}\right]=0 \tag{15}$$

to avoid the error scaling with marginal utility. This measurement provides a way to check for accuracy after adding an order of approximation. Thus one can decide at each step whether the approximation suffices the criteria or not. As an additional benefit, there is no need to recompute previous orders after each step. The approximation method keeps previous coefficients unaltered when refining the solution.

3.7. Generality

The numerical method described previously is applicable to a variety of complete and incomplete market models. This section lays out the features that a model needs to possess to be handled by our technique.

First, the optimality conditions must be representable by a system of equations consistent with the operator in Eq. (6). Both competitive equilibria and dynamic programming problems can be of this type. The operator ensures that first-order conditions (or the Bellman equation for dynamic programming problems) along with the equations of motion, market clearing conditions, and budget constraints hold. To apply perturbation methods, we require three main assumptions for the models.

First, we require the model to have differentiable policy functions. We apply perturbation methods to the problem, which builds a Taylor series expansion of the optimal policies around the deterministic steady state. In many economic problems, optimal policies have smooth policy functions and the Taylor series converges within a radius. If the assumptions are met, the implicit function theorem guarantees existence of a differentiable policy function around the deterministic steady state (see e.g. the implicit function theorem for analytic functions in Judd and Guu, 2001). Second, we require the economy to have a well-defined ergodic distribution to avoid unit roots. And third, we require the existence of a deterministic steady state at which all households are identical.[5]

If the equilibrium conditions given by functionals $g^{\cdot}$ display the symmetry properties of Eq. (13), the Taylor series of only one household needs to be approximated. Furthermore, the number of distinct coefficients in the Taylor series is small and, for a given degree, independent of the number of households.

4. Results

We first discuss the choice of functional forms and parameter values for the model in Section 2. Then we show the accuracy of the solution method before discussing our findings.

4.1. Calibration

For easier comparison with the previous literature, we implement the calibration to quarterly time periods in Den Haan et al. (2010) that was given to a number of research teams that discussed the relative performance of different algorithms for models with a cross-section of heterogeneous agents. We implement the specification in Kim et al. (2010), who presented a numerical algorithm based on perturbation methods.

We therefore specify the utility function to have constant relative risk aversion with a coefficient of one, i.e., logarithmic utility. The time preference factor β is 0.99. The parameters governing the technology are standard in the literature with a capital share of output α of 0.36 and a depreciation rate of 2.5%.

For the shocks to aggregate productivity and shocks to individual labor productivity, Kim et al. (2010) use a continuous representation of the first-order Markov processes in Den Haan et al. (2010). Therefore, we transform the idiosyncratic labor income shock $\psi=\frac{\hat\psi}{\hat\Psi}$ and use the autoregressive process for $\hat\psi$ from Kim et al. (2010)

$$\hat\psi^i_{t+1}=\bar\phi_\psi+\phi_\psi\hat\psi^i_t+\phi_\theta(\hat\psi^i_t)\sigma\theta^i_{t+1}. \tag{16}$$

[5] Imposing identical households in the steady state does not mean that we cannot allow for heterogeneity. We can either allow for a finite number of different types. Alternatively, we can scale heterogeneity among agents, say in their risk aversion, by the perturbation parameter σ. That way, the steady-state remains unaffected but the higher-order expansion will introduce heterogeneity across agents.

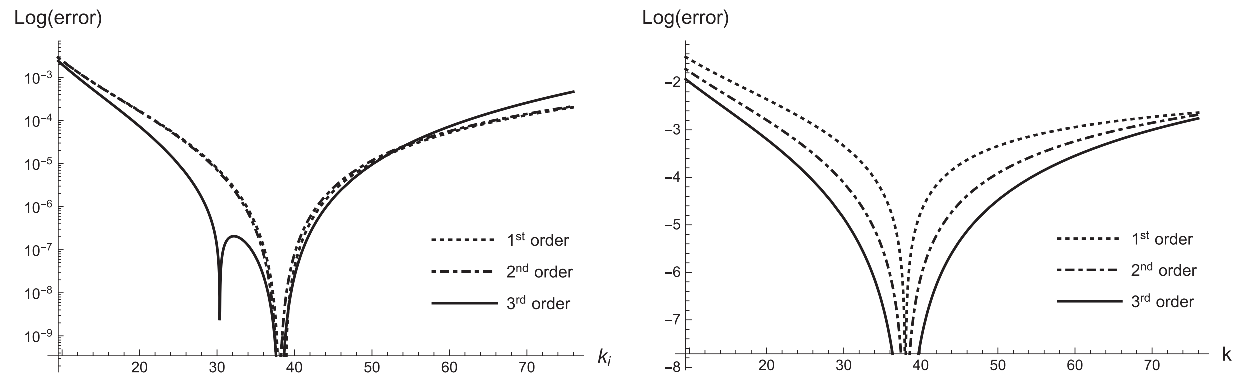

Fig. 2. Euler equation error in the deterministic economy for deviations of capital holdings for one agent (left graph) or all agents (right graph).

The parameter ϕ_ψ governs the degree of persistence in the evolution of the shock. $\bar{\phi}_\psi$ adjusts the long-run mean and $\phi_\theta(\psi_t^i)$ governs the standard deviation of the shock that we allow to be a function of $\hat{\psi}_t^i$. θ is white noise with unit variance.

Following the calibration from the literature, we set the process for the transformation of labor income shocks to

$$\hat{\psi}_{t+1}^i = 0.4 + 0.55555\hat{\psi}_t^i + (0.48989 - 0.28381\hat{\psi}_t^i)\theta_{t+1}^i.$$

We set upper and lower bounds for innovations to ± 0.1.

Since we work with a finite number of households, the average of all individual shocks $\hat{\Psi}_t = \frac{1}{I}\sum_{i=1}^{I} \hat{\psi}_t^i$ will be stochastic. To keep the average labor income shock constant over time, we normalize individual labor income shocks by $\hat{\Psi}_t$.

Total factor productivity follows an AR(1) process

$$z_{t+1} = \bar{\phi}_z + \phi_z z_t + \phi_\varepsilon \sigma \varepsilon_{t+1} \tag{17}$$

where the parameter ϕ_z determines the degree of mean reversion in total factor productivity. The calibration sets $\bar{\phi}_z$ to 0.25, ϕ_z to 0.75, and ϕ_ε to 0.00661.

We specify all shocks to z and ψ to have a mean of zero and a unit standard deviation. It is standard in the literature in macroeconomics to specify a normal distribution for these shocks and truncate them when solving the model numerically. For the application of perturbation methods, we use a finite support and therefore want to think of these shocks as following a truncated normal distribution.

The numerical parameter for the penalty function ν is set to $\frac{1}{1,000,000}$. Despite being a rather arbitrary choice, it ensures that the impact of ν on the coefficients in the Taylor series is negligible.

The last parameter governs the number of households in the economy. As demonstrated when describing the solution method, the computing power required is the same for any number of individuals. To generate the results of this section, we set this number to roughly the current number of citizens in the United States, 324,420,000 leaving us with 648,840,002 state variables.

4.2. Convergence

As mentioned in Section 3, the numerical method produces a polynomial solution for which we report the coefficients in Appendix C. To check the accuracy of this solution, we normalize the Euler equation by dividing by marginal utility on both sides as in Eq. (15). Fig. 2 plots the logarithm of the Euler equation error as a function of one household's capital stock. This check for accuracy corresponds to the deterministic version of the economy.

The deterministic steady state satisfies the deterministic optimality conditions. Thus, the Euler equation error is zero at this point and its logarithm at negative infinity.

Two observations stand out from this graph. First, we see convergence. The Euler equation error decreases for the interval. And second, the result approximates the solution not just locally but globally on a sizable interval.

4.3. Impact of heterogeneity

Heterogeneity with aggregate risk increases the steady-state level of capital, as known from previous literature. There are two reasons for it. First, idiosyncratic risk leads to precautionary savings on the part of households. Since households cannot trade claims contingent on their labor income, they try to partially insure against these shocks by building up a buffer stock of savings. Second, due to aggregate productivity shocks, holding capital is risky. There are two opposing forces. On the one hand, households are risk averse and demand a higher risk premium for holding risky capital. Each unit should thus return a higher dividend, which implies a higher marginal product of capital and thus a lower steady-state level of capital. On the other hand though, since returns to capital are risky, households respond by building up savings, which implies a higher steady-state level of capital. With our utility specification of constant relative risk aversion, the latter effect dominates. Thus, heterogeneity with aggregate risk increases the steady-state level of capital.

The use of perturbation methods lends itself to a novel representation of the stochastic discount factor in heterogeneous household modes. A nonlinear change of variables transforms the equilibrium policy functions into an expansion of the stochastic discount factor. This representation allows us to study the pricing of assets in our incomplete markets economy.

We see that an aggregate stochastic discount factor defined as the average of all individual discount factors is approximated by

$$\frac{\beta}{I}\sum_{i=1}^{I}\frac{u_c'(c_{t+1}^i)}{u_c'(c_t^i)} = c(\mathbf{X}_t, \mathbf{z}_t, \mathbf{X}_{t+1}) + \left(c_z^{(1)}e^{z_{t+1}} + c_z^{(2)}\mathrm{var}(e^{z_{t+1}}) + \ldots\right) + \left(c_\psi^{(2)}\mathrm{var}(\psi^i) + \ldots\right) \tag{18}$$

The derivation of this expression can be found in Appendix D.

The first term as well as all coefficients in the expansion are known at time t. The approximation thus tells us directly which assets demand a risk premium. Every security that comoves with total factor productivity, its variance (which we kept constant in our example), and higher moments will carry a risk premium.

But heterogeneity also enters the pricing of securities. The variance of individual labor income shocks is a risk factor. A time- or state-dependent variance will induce a risk premium for all assets that comove with this variance. Therefore, our economy displays risk pricing effects of idiosyncratic income risk similar to Constantinides and Duffie (1996). Other than in their work, the impact of heterogeneity varies with aggregate productivity and the amount of capital. This effect can be seen from the third-order terms in the Taylor expansion where the variance of idiosyncratic shocks is interacted with aggregate capital or total factor productivity.

5. Comparison between methods

The main alternative to our method is based on the idea of replacing the true law of motion by a function that depends on moments of state variables. We demonstrate the performance of a perturbation-based solution method in comparison with a general approach that underlies the standard methods in the literature, such as Krusell and Smith (1998), Kim et al. (2010), or Preston and Roca (2007).

To see the difference between the different approaches, we study a particularly challenging problem that we borrow from Panageas (2011). The model gives rise to an explicit solution that we can compare our approximate solutions to. Panageas (2011) provided this model as a challenge to methods that replace the law of motion by a low-order polynomial.

Panageas (2011) uses a simple representative agent economy with a Lucas tree where two taste shocks make up the state space. As opposed to the setup in Section 2, the multi-dimensionality of the state space thus stems from different shocks as opposed to individual state variables by different agents. This assumption allows us to control the persistence of different state variables, which is a crucial feature of this example. While this model does not feature a cross-section of agents, it shows a condition under which quasi-aggregation performs poorly.[6]

A representative household prices a stochastic stream of endowments C_t according to the following stochastic process

$$\log C_{t+1} = \log C_t + \mu + \varepsilon_{t+1}$$

where the innovation ε is distributed $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. The representative household's expected utility is given by the discounted stream of per-period utilities that have constant relative risk aversion with a coefficient γ and a preference shock ξ_t

$$\mathcal{U}_0 = E_0\left[\sum_{t=0}^{\infty}\beta^t e^{\xi_t}u_c(C_t)\right]. \tag{19}$$

The preference shock has two components

$$\xi_t = A_t + B_t$$

that evolve according to the stochastic processes

$$\log A_{t+1} = \rho_A \log A_t + \sigma_A \eta_{t+1}$$
$$\log B_{t+1} = \rho_B \log B_t + \sigma_B \eta_{t+1},$$

where the innovations are standard normally distributed.

To determine the value of the tree, we use the representative household's Euler equation

$$\frac{P_0}{C_0} = E_0\left[\beta\frac{A_1}{A_0}\frac{B_1}{B_0}\left(\frac{C_1}{C_0}\right)^{1-\gamma}\left(\frac{P_1}{C_1}+1\right)\right] \tag{20}$$

[6] Thanks to an anonymous referee for pointing this out.

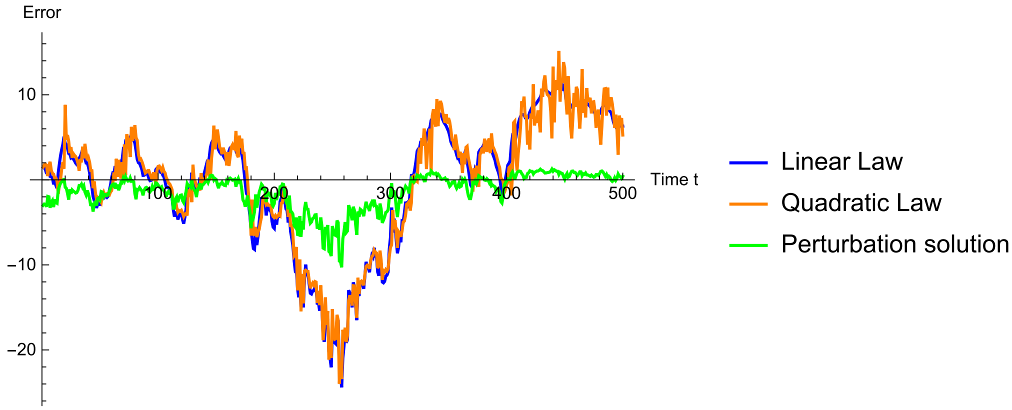

Fig. 3. This figure compares deviations of approximations based on a linear or quadratic law of motion versus perturbation from the explicit solution.

and iterate forward to get

$$\frac{P_0}{C_0} = \sum_{t=1}^{\infty} \beta^t e^{-\gamma\mu t + \frac{\gamma^2\sigma_\varepsilon^2}{2} t} A_0^{\rho_A^t - 1} B_0^{\rho_B^t - 1} e^{\frac{1}{2}\left[\sigma_A^2 \frac{1-\rho_A^{2t}}{1-\rho_A^2} + \sigma_B^2 \frac{1-\rho_B^{2t}}{1-\rho_B^2} + 2\sigma_A\sigma_B \frac{1-\rho_A^t\rho_B^t}{1-\rho_A\rho_B}\right]}. \tag{21}$$

Appendix E.1 contains a derivation.

We can evaluate the quasi-closed-form solution with arbitrary accuracy by forward iteration given by Eq. (21). This reference solution serves as a benchmark for two approximation methods. First, we can assume a linear law of motion. Second, we can use the approximation method described in this paper to solve for the pricing.

As the closed-form solution in Eq. (21) shows, the state space depends on the two preference shocks, A_t and B_t. The distribution of shocks that make up the preference shock is thus the state space. A common theme in approximation methods is now to approximate the law of motion by moments of this distribution. In most specifications, however, only the mean is being used.

To implement this law of motion, we replace the true law of motion with

$$\frac{P_{t+1}}{C_{t+1}} = \bar{\phi}_{PC} + \phi_{PC}\frac{P_t}{C_t} + \phi_\eta \eta_{t+1} \tag{22}$$

which is a linear law of motion for which we will choose the coefficients optimally. Using this law of motion, we arrive at an approximate closed-form expression for the price-dividend ratio

$$\frac{P_0}{C_0} = \frac{\beta e^{-\gamma\mu + \frac{\gamma^2\sigma_\varepsilon^2}{2}} A_0^{\rho_A - 1} B_0^{\rho_B - 1} e^{\frac{1}{2}(\sigma_A+\sigma_B)^2}\left((1+\bar{\phi}_{PC}) + \phi_\eta(\sigma_A+\sigma_B)\right)}{1 - \phi_{PC}\beta e^{-\gamma\mu + \frac{\gamma^2\sigma_\varepsilon^2}{2}} A_0^{\rho_A - 1} B_0^{\rho_B - 1} e^{\frac{1}{2}(\sigma_A+\sigma_B)^2}}. \tag{23}$$

The derivation of this formula and details on the quadratic law of motion are in Appendix E.3.

The coefficients in the linear law are chosen to maximize the fit with the dynamic evolution under this law of motion. Therefore, we fix the coefficients, solve, and simulate the economy. We run a linear regression of next period's price-consumption ratio on this period's ratio. Finally, we use the resulting coefficients as a new law of motion and iterate until a fixed point is found.

A related question is whether we could do better by adding the squared term of the price-consumption ratio in Eq. (22). To determine the coefficients, we project the next period's price-consumption ratio onto the linear and quadratic term.

We solve the economy for the identical parameter combination and realization of shocks as in Panageas (2011). For this parameter combination, the R^2 criterion for the linear law of motion provides values above 98%. Specifically, these parameters are a growth rate $\mu = 1.4\%$, risk aversion at $\gamma = 8$, and the time discount factor $\beta = 1.05$. The persistence of the two shocks is set to $\rho_A = 0.98$ and $\rho_B = 0.8$. The standard deviations are fixed at $\sigma_\varepsilon = 0.04$, $\sigma_A = 0.1$, and $\sigma_B = 0.04$.

Fig. 3 shows the quality of the best fit given a linear law of motion and the quadratic approximation as the deviation from the explicit solution. Several results emerge.

First, we observe that even the best fit with an approximated linear law of motion can substantially deviate from the closed-form solution. The approximation can deviate by more than 20 from the true solution where the price-dividend ratio varies between 30 and 90 (see Fig. 4). The discrepancy arises from the fact that the two components of the preference shock have different persistence and, as a result, the average of the two components is a poor statistic of the true law of motion.

Second, adding a quadratic term does not improve the overall fit much. In general, moments of the distribution cannot capture the entire dynamics of the state variables since each component has a differential impact on next period's state.

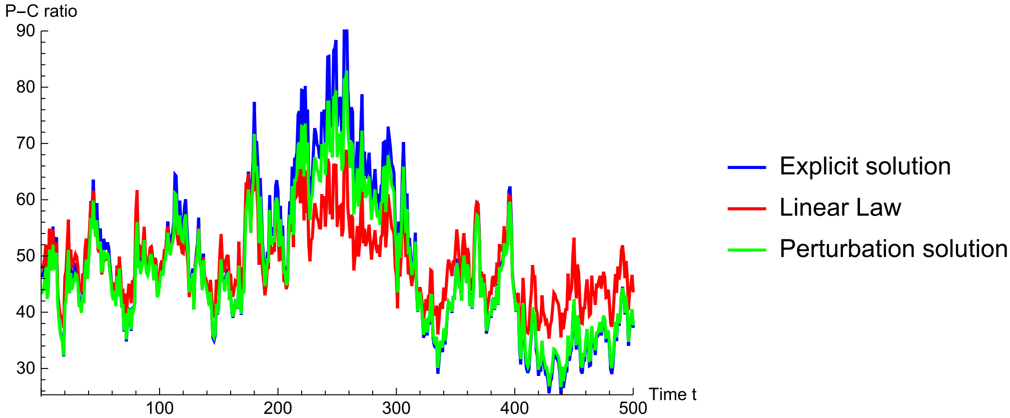

Fig. 4. Comparison between the true solution and approximation methods using a linear law of motion and the perturbation approach of this paper.

How does our solution method fare with regards to this asset pricing problem? A perturbation-based method of this paper is well suited to solve this economy. Heterogeneity in the two state variables arises from uncertainty. Both preference shocks would remain at their means in the deterministic economy.

We apply perturbation methods to the Euler equation in Eq. (20). The price-dividend ratio is a function of the two state variables A_t and B_t. We start with the deterministic steady state around which we approximate the price-dividend ratio. Then we proceed in the standard fashion by building a high-order perturbation in the two state variables. Finally, we take the derivatives (and cross-derivatives with the two state variables) with respect to the standard deviation of the shocks.

This is a particularly hard problem for the solution technique because the price-consumption ratios in the stochastic economy are in a different range from their deterministic counterparts, which lies at a price consumption ratio of roughly 15. If we set the standard deviations in Eq. (21) to zero, the deterministic price-consumption ration will be far smaller. In our later calibration, the stochastic price-consumption ratios will be more than ten times larger than the deterministic steady state.

Fig. 3 shows, however, that the fit for the perturbation-based approach is far better than for the approximations using a linear law of motion. For perturbation, we choose an approximation of order 5. Fig. 4 shows the time series in levels. As can be seen from there, the approximation based on perturbation tracks the true solution closely, while a linear law of motion deviates substantially although the R^2 diagnostic indicates a good fit (Panageas (2011).

Panageas (2011) created this example to demonstrate difficulties when assuming a linear or quadratic law of motion. A linear or quadratic law of motion might deliver a poor approximation when the model is either highly nonlinear or when it is comprised of several state variables with different persistence.

There is a key difference in our method compared to the use of a linear or quadratic law of motion. Our method approximates the system of equilibrium conditions. Therefore, we determine the magnitude of those errors on the domain over which we solve the problem. Those tests are demanding in that they demand that the Euler equation condition on each state be economically insignificant. Replacing the law of motion by linear or quadratic functions requires approximating both the system of equilibrium conditions for each individual as well as the law of motion of the state. Checking the error of the approximation for the law of motion is non-trivial, as, for example, pointed out by Panageas (2011).

An alternative to using perturbation methods would be projection methods. Projection methods have the advantage of being global methods. They also work for economies with hard constraints that produce non-differentiable policy functions. However, since solution functions are approximated on grids, the number of state variables is rather limited. Perturbation methods, on the other hand, can handle high-dimensional state spaces, particularly when there are symmetries with respect to the derivatives as shown in this paper. Perturbation methods are local by design but converge globally within in the radius of convergence. Typically, perturbation methods are fast since they can be implemented without iteration.

Which method is preferable depends on the specific economic model. The representative agent counterpart of our model in Section 2 is a standard real business cycle model. Aruoba et al. (2006) show for a similar model that perturbation methods perform well. Our Euler equation errors confirm this conclusion.

6. Extensions

The model of Section 2 is stylized and lacks some of the features of models at the frontier of research in macroeconomics and finance. This section shows how our perturbation-based approach can be used for more complex models.

6.1. State-of-the-art DSGE models

There are a number of extensions that have become standard in the business cycle literature. Here are a few examples: First, the utility function can be defined over consumption as well as leisure, which adds another optimality condition to the system. Second, adjustment costs based on investment growth such as in Christiano et al. (2005) deliver a better fit for business cycle moments. These specifications add lagged investment as an additional state variable. Third, the literature has moved towards matching asset pricing as well as macroeconomic data. Therefore, habit utility or Epstein-Zin preferences in conjunction with long-run risks in productivity have become the norm. Additionally, the inclusion of preference shocks as in Pavlova and Rigobon (2007) is feasible.

All of these extensions fall into the category of models that can be solved with the type of perturbation methods we propose in this paper. While additional state variables or additional choice variables add complexity to the solution, the extension of the algorithm is straightforward. In particular, perturbation methods have advantages for models with many state variables since they do not require the solution to be approximated on a grid that would be limited in size by the curse of dimensionality.

6.2. Borrowing constraints

The inclusion of borrowing constraints is another potentially interesting extension. Hard borrowing constraints typically induce non-differentiabilities at certain points in the policy function. Since perturbation methods work with derivatives and approximate optimal policies by their Taylor series expansion, the method cannot handle non-differentiabilities. There are two ways of dealing with the inclusion of borrowing constraints.

First, modeling borrowing constraints as penalties on interest rates that increase when the borrower is close to the credit limit renders the model tractable for perturbation methods. This smooth increase in the borrowing rate, until borrowing becomes prohibitively expensive, does not lead to a loss of derivatives in the relevant region and the perturbation methodology of this paper can be applied.

Second, we discuss the use of perturbation methods for models in which the borrowing limits take the form of inequality constraints on the amount of borrowing. As is standard in the optimization literature (see, for example, Nocedal and Wright, 2006) and implemented in many optimization algorithms (for example in sequential quadratic programming methods), the borrowing constraint can be approximated arbitrarily well by a barrier function that preserves differentiability. Barrier methods have been used successfully in conjunction with perturbation methods by Preston and Roca (2007) and Kim et al. (2010). Therefore, we would add a barrier function to the objective of maximizing utility, for example, of the form

$$u_k(k_t^i) = -\nu_1 \frac{1}{(k_t^i - \underline{k})^2} + 2\nu_2 k_t^i, \tag{24}$$

where $\underline{k}$ denotes the lower bound on capital and $\nu_\cdot > 0$ parameters for the barrier function. We impose the restriction $\nu_2 = \frac{\nu_1}{(\bar{k}-\underline{k})^3}$ that ensures a vanishing derivative at the deterministic steady state $\bar{k}$.

A hard constraint can be reformulated as a limiting sequence of barrier functions (see Appendix F for the theory). Under the barrier approach, borrowing is penalized with an increasing penalty close to the borrowing limit. Note that the barrier function only takes effect when going to higher-order expansions. It induces a singularity in the objective function and thus limits the radius of convergence to the relevant part of the state space. The singularity ensures that higher-order approximations deliver stronger penalties.

6.3. Portfolio choice

Lastly, perturbation methods of the form discussed here can deal with portfolio choice problems where agents choose between riskless and risky assets. The main complication is an indeterminacy of the portfolio for the auxiliary deterministic economy. This indeterminacy arises since returns to capital and bonds are identical (and both riskless) in a deterministic economy. There are two ways to deal with this indeterminacy. First, the Bifurcation Theorem, instead of the Implicit Function Theorem, justifies building a Taylor series approximation around the point to which the solution converges when risk is decreased to zero (see Judd and Guu, 2000). Second, we can introduce a penalty function for bond holdings different from a target level and multiply it by a penalty coefficient akin to the penalty introduced in Eq. (8). Any positive coefficient, no matter how small, will prevent households from deviating from identical portfolios. As a result, the perturbation methods of this paper apply to the penalty problem without modification.

7. Conclusion

In this paper, we presented a numerical method for solving an incomplete markets model with an arbitrarily large cross-section of households. Our algorithm builds on perturbation methods that use Taylor series expansions around a deterministic steady state. This solution method is particularly useful for models with many state and choice variables. Generally, this idea can be applied not only to competitive equilibria but also to dynamic programming problems. As a first example, we solved a dynamic stochastic general equilibrium model with idiosyncratic shocks to labor income. We demonstrated the convergence properties for this particular example. Furthermore, we showed that heterogeneity impacts macroeconomic quantities as well as the pricing of risk.

Appendix A. Structure of model setup

A1. Structure of variables

The structure of state variables is given by

$$\mathbf{X}_t = \begin{pmatrix} k_t^1 & \psi_t^1 \\ \vdots & \vdots \\ k_t^I & \psi_t^I \end{pmatrix} \qquad \mathbf{z}_t = (z_t).$$

Equilibrium choice, price, and transition functions are

$$\mathbf{C}(\mathbb{S}) = \begin{pmatrix} c^1(\mathbb{S}) & k_{+1}^1(\mathbb{S}) \\ \vdots & \vdots \\ c^I(\mathbb{S}) & k_{+1}^I(\mathbb{S}) \end{pmatrix} \qquad \mathbf{p}(\mathbb{S}) = \begin{pmatrix} r^k(\mathbb{S}) \\ w(\mathbb{S}) \\ K(\mathbb{S}) \\ K_{+1}(\mathbb{S}) \\ Y(\mathbb{S}) \\ C(\mathbb{S}) \end{pmatrix} \qquad \boldsymbol{\tau}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) = \begin{pmatrix} \psi_{t+1}^1(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) \\ \vdots \\ \psi_{t+1}^I(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) \\ \Psi(\mathbb{S}_t) \\ z_{t+1}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) \\ w(\mathbb{S}_t) \\ r^k(\mathbb{S}_t) \\ C(\mathbb{S}_t) \\ K_{+1}(\mathbb{S}_t) \\ K(\mathbb{S}_t) \\ Y(\mathbb{S}_t) \end{pmatrix}$$

where k_{+1}^i denotes household i's optimal policy of capital holdings for the following period and K_{+1} the corresponding aggregate capital stock.

A2. System of equations

This section describes the operator G which defines the equilibrium conditions. The first function g_1 describes the Euler equations

$$g_1(\mathbf{X}_t, \mathbf{z}_t, \mathbf{X}(\sigma), \mathbf{z}(\sigma), \mathbf{C}(\sigma), \mathbf{p}(\sigma)) = \begin{pmatrix} u_c'(c_t^1) - \beta\Big((1-\delta+r_{t+1}^k)u_c'(c_{t+1}^1)\Big) \\ \vdots \\ u_c'(c_t^I) - \beta\Big((1-\delta+r_{t+1}^k)u_c'(c_{t+1}^I)\Big) \end{pmatrix}.$$

The second part of the operator g_2 collects the budget constraints, laws of motion for random variables, market clearing conditions, and aggregates variables. We assume mean reverting stochastic processes for preference shocks, total factor productivity, and idiosyncratic labor income shocks as defined in Section 4.1.

$$g_2(\mathbf{X}_t, \mathbf{z}_t, \mathbf{C}, \mathbf{p}, \boldsymbol{\tau}, \mathbf{e}_{t+1}) = \begin{pmatrix} c^1(\mathbb{S}_t) + k_{+1}^1(\mathbb{S}_t) - (1-\delta+r^k(\mathbb{S}_t))k_t^1 - w(\mathbb{S}_t)\psi_t^1 \\ \vdots \\ c^I(\mathbb{S}_t) + k_{+1}^I(\mathbb{S}_t) - (1-\delta+r^k(\mathbb{S}_t))k_t^I - w(\mathbb{S}_t)\psi_t^I \\ \boldsymbol{\tau}_1(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - (\bar{\phi}_\psi + \phi_\psi\psi_t^1 + \phi_\theta(\psi_t^1)\sigma\theta_{t+1}^1) \\ \vdots \\ \boldsymbol{\tau}_I(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - (\bar{\phi}_\psi + \phi_\psi\psi_t^I + \phi_\theta(\psi_t^I)\sigma\theta_{t+1}^I) \\ \boldsymbol{\tau}_{I+1}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - \frac{1}{I}\sum_i \psi_t^i \\ \boldsymbol{\tau}_{I+2}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - (\bar{\phi}_z + \phi_z z_t + \phi_\varepsilon\sigma\varepsilon_{t+1}) \\ \boldsymbol{\tau}_{I+4}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - (1-\alpha)z_t K_t^\alpha L_t^{-\alpha} \\ \boldsymbol{\tau}_{I+5}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - \alpha z_t K_t^{\alpha-1} L_t^{1-\alpha} \\ \boldsymbol{\tau}_{I+6}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - \sum_i c^i(\mathbb{S}_t) \\ \boldsymbol{\tau}_{I+7}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - \sum_i k_{+1}^i(\mathbb{S}_t) \\ \boldsymbol{\tau}_{I+8}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - \sum_i k_t^i \\ \boldsymbol{\tau}_{I+9}(\mathbb{S}_t, \sigma\mathbf{e}_{t+1}) - z_t K_t^\alpha L_t^{1-\alpha} \end{pmatrix}$$

Appendix B. Approximation of the Law of Motion

B1. The linear law of motion

Next period's average capital as $\frac{1}{I}K_{t+1} = \frac{1}{I}\sum_{i=1}^{I} k_{t+1}^{i}(\mathbf{X}_t, \mathbf{z}_t, \sigma)$. Using Eq. (10), we get a first-order approximation of the form

$$\begin{aligned}\frac{1}{I}K_{t+1} \approx k_0 & \\ & + \sum_{i=1}^{I} k^i_{\mathbf{X}_{11}}(\mathbf{X}_{11} - \bar{\mathbf{X}}_{11}) + k^i_{\mathbf{X}_{21}}(\mathbf{X}_{21} - \bar{\mathbf{X}}_{21}) + k^i_{\mathbf{X}_{31}}(\mathbf{X}_{31} - \bar{\mathbf{X}}_{31}) + k^i_{\mathbf{X}_{41}}(\mathbf{X}_{41} - \bar{\mathbf{X}}_{41}) + \ldots \\ & + \sum_{i=1}^{I} k^i_{\mathbf{X}_{12}}(\mathbf{X}_{12} - \bar{\mathbf{X}}_{12}) + k^i_{\mathbf{X}_{22}}(\mathbf{X}_{22} - \bar{\mathbf{X}}_{22}) + k^i_{\mathbf{X}_{32}}(\mathbf{X}_{32} - \bar{\mathbf{X}}_{32}) + k^i_{\mathbf{X}_{42}}(\mathbf{X}_{42} - \bar{\mathbf{X}}_{42}) + \ldots \\ & + \ldots \\ & + \sum_{i=1}^{I} k^i_{\mathbf{z}_1}(\mathbf{z}_1 - \bar{\mathbf{z}}_1) + k^i_{\mathbf{z}_2}(\mathbf{z}_2 - \bar{\mathbf{z}}_2) + k^i_{\mathbf{z}_3}(\mathbf{z}_3 - \bar{\mathbf{z}}_3) + \ldots\end{aligned}$$

where $\bar{\mathbf{X}}_{\cdot 1}$ is the deterministic steady-state value.

With the assumptions on symmetry, coefficients on expansions as well as steady-state values are identical and summarize to

$$\begin{aligned}\frac{1}{I}K_{t+1} \approx k_0 & \\ & + (k^1_{\mathbf{X}_{11}} - k^1_{\mathbf{X}_{21}})(\mathbf{X}_1 - \bar{\mathbf{X}}_{11}) + Ik^i_{\mathbf{X}_{21}}(\mathbf{X}_1 - \bar{\mathbf{X}}_{11}) \\ & + (k^1_{\mathbf{X}_{12}} - k^1_{\mathbf{X}_{22}})(\mathbf{X}_2 - \bar{\mathbf{X}}_{12}) + Ik^i_{\mathbf{X}_{22}}(\mathbf{X}_2 - \bar{\mathbf{X}}_{12}) \\ & + \sum_{i=1}^{I} k^i_{\mathbf{z}_1}(\mathbf{z}_1 - \bar{\mathbf{z}}_1) + k^i_{\mathbf{z}_2}(\mathbf{z}_2 - \bar{\mathbf{z}}_2) + k^i_{\mathbf{z}_3}(\mathbf{z}_3 - \bar{\mathbf{z}}_3) + \ldots\end{aligned}$$

where $\mathbf{X}_j = \frac{1}{I}\sum_{i=1}^{I}\mathbf{X}_{ij}$.

B2. The quadratic law of motion

For the second-order terms, we again build an expansion for one policy and sum up over all households. Thereby, we invoke symmetry in the analogous fashion. Simply regrouping the terms from Eq. (10) yields

$$\begin{aligned}k^1(X,z) \approx k_0 & + \text{first-order terms} \\ & + (k^1_{\mathbf{X}_{11},\mathbf{X}_{11}} - k^1_{\mathbf{X}_{21},\mathbf{X}_{21}} - 2k^1_{\mathbf{X}_{11},\mathbf{X}_{21}} + k^1_{\mathbf{X}_{21},\mathbf{X}_{31}})(\mathbf{X}_{11} - \bar{\mathbf{X}}_1)^2 \\ & + (k^1_{\mathbf{X}_{21},\mathbf{X}_{21}} - k^1_{\mathbf{X}_{21},\mathbf{X}_{31}})\sum_{i=1}^{I}(\mathbf{X}_{i1} - \bar{\mathbf{X}}_1)^2 \\ & + (2k^1_{\mathbf{X}_{11},\mathbf{X}_{21}} - k^1_{\mathbf{X}_{21},\mathbf{X}_{31}})(\mathbf{X}_{11} - \bar{\mathbf{X}}_1)(\mathbf{X}_1 - \bar{\mathbf{X}}_1) \\ & + k^1_{\mathbf{X}_{21},\mathbf{X}_{31}}(\mathbf{X}_1 - \bar{\mathbf{X}}_1)^2.\end{aligned}$$

To get to average capital, we average across all households and invoke symmetry

$$\begin{aligned}\frac{1}{I}\sum_{i=1}^{I}k^i(X,z) \approx k_0 & + \text{first-order terms} \\ & + (k^1_{\mathbf{X}_{11},\mathbf{X}_{11}} - 2k^1_{\mathbf{X}_{11},\mathbf{X}_{21}})\sum_{i=1}^{I}(\mathbf{X}_{i1} - \bar{\mathbf{X}}_1)^2 + 2k^1_{\mathbf{X}_{11},\mathbf{X}_{21}}(\mathbf{X}_1 - \bar{\mathbf{X}}_1)^2.\end{aligned}$$

B3. Approximation of any symmetric variable

The previous logic goes through for every approximation of a variable $f(\mathbf{k})$ where

$$\frac{\partial f}{\partial k^i} = \frac{\partial f}{\partial k^j} \qquad \forall i, j$$

Table 1
Coefficients in the approximate solution of the policy function for the model in Section 2. The table reports cross-derivatives where 1 stands for constants such that the cross-derivative between 1 and 1 is the steady-state level, between 1 and k^1 is the first derivative with respect to k^1, etc.

	1	k^1	k^2	k^3	ψ^1	ψ^2	ψ^3	z	σ
1	37.9893	0.9977	-6.7×10^{-11}		2.3415	-1.6×10^{-10}		3.5793	0
k^1		1.6×10^{-6}	-4.1×10^{-12}		3.9×10^{-6}	2.8×10^{-10}		2.5293	0
k^2			-3.1×10^{-15}	1.8×10^{-20}	2.0×10^{-9}	-7.2×10^{-15}	-3.2×10^{-18}	-5.8×10^{-9}	0
ψ^1					9.1×10^{-6}	-6.6×10^{-10}		8.1962	0
ψ^2						-1.7×10^{-14}	-7.3×10^{-18}	-1.4×10^{-8}	0
z								2.1961	0
σ									0.4371

The first-order expansion delivers

$$f(\mathbf{X}) \approx f(\bar{\mathbf{X}}) + f_{\mathbf{X}_{11}}(\bar{\mathbf{X}})(\mathbf{X}_1 - \bar{\mathbf{X}}_1) + f_{\mathbf{X}_{21}}(\bar{\mathbf{X}})(\mathbf{X}_2 - \bar{\mathbf{X}}_2) + \dots$$

The symmetry conditions are simply

(i) $\frac{d^2 f}{dx_i^2} = \frac{d^2 f}{dx_1^2}$

(ii) $\frac{d^2 f}{dx_i dx_j} = \frac{d^2 f}{dx_1\ dx_2}$ for $i \neq j$

Using these conditions, we can simplify the expansion of f analogously to before. We get

$$f(\mathbf{X}) \approx f_0 + \text{first-order terms} + I\left(\frac{d^2 f}{dx_1^2} - \frac{d^2 f}{dx_1\ dx_2} + I^2 \frac{d^2 f}{dx_1\ dx_2}\right)(X_1 - \bar{X}_1)^2$$

Appendix C. Approximate solution of policy function

Table 1 reports the coefficients in a second-order Taylor expansion of the optimal policy function for capital holdings for agent 1. The cross-derivatives including “1” denote constants such that the number in the upper left denotes the steady-state level of capital holdings. The rest of the first row are the first-order terms whereas the rest of the table represents second-order derivatives between the variables in the first column and row.

Appendix D. Approximation of a Stochastic Discount Factor

One stochastic discount factor is given by the average individual stochastic discount factor. We use the technique of a nonlinear change of variables to approximate it with perturbation methods. Therefore, we recognize that the marginal utility of next period’s consumption is a function of $\mathbf{X}_{t+1}$ and $\mathbf{z}_{t+1}$ while the marginal utility of consumption today is a function of today’s state variables. Together, we build one Taylor expansion with respect to all these state variables. Applying the logic from Eq. (10), we arrive at

$$\frac{\beta}{I}\sum_{i=1}^{I}\frac{u'(c_{t+1}^i)}{u'(c_t^i} \approx \sum_{o=0}^{\infty}\sum_{\|\mathbf{I}\|+|\mathbf{j}|+k=o}\frac{1}{\mathbf{I}!\cdot\mathbf{j}!\cdot k!}\left(h_{t,\mathbf{I},\mathbf{j},k}(\mathbf{X}_t,\mathbf{z}_t) + h_{t+1,\mathbf{I},\mathbf{j},k}(\mathbf{X}_{t+1},\mathbf{z}_{t+1})\right) \tag{25}$$

where $h_{\cdot,\mathbf{I},\mathbf{j},k}(\mathbf{X}_t, \mathbf{z}_t) = \mathbf{U}^{(\mathbf{I},\mathbf{j},k)}(\mathbf{X}^0, \mathbf{z}^0, 0)(\mathbf{X}_\cdot - \mathbf{X}^0)^{\mathbf{I}}(\mathbf{z}_\cdot - \mathbf{z}^0)^{\mathbf{j}}\sigma^k$. The function U is represents marginal utility of consumption for period $t+1$ and the inverse thereof for period t. The derivatives at the deterministic steady-state are computed using the nonlinear change of variables.

Given this expansion, collect all monomial terms merely depending on state variables known in period t, i.e., $\mathbf{X}_t$, $\mathbf{z}_t$, and $\mathbf{X}_{t+1}$. The collection of those terms is denoted by $c(\mathbf{X}_t, \mathbf{z}_t, \mathbf{X}_{t+1})$ in Eq. (18). Next, collect all terms in which total factor productivity appears linearly. These terms are the first-order term in the expansion of productivity and all cross-terms with variables known at time t. Collect those in a term $c_z^{(1)}(\mathbf{X}_t, \mathbf{z}_t, \mathbf{X}_{t+1})$ where we drop the arguments in Eq. (18). We similarly collect the terms for second-order expansions with respect to total factor and individual productivity and arrive at Eq. (18).

Appendix E. Asset Pricing Example — Derivations

E1. Derivation of the closed form solution for the asset pricing economy

We start from the Euler equation of the tree (20) which we iterate to get

$$\frac{P_0}{C_0} = E_0\left[\sum_{t=1}^{\infty}\beta^t \frac{A_t}{A_0}\frac{B_t}{B_0}\left(\frac{C_t}{C_0}\right)^{1-\gamma}\right]. \tag{26}$$

Now we plug in the stochastic processes that we iterate to get

$$\frac{A_t}{A_0} = A_0^{\rho_A^t - 1} e^{\sigma_A \sum_{j=0}^{t-1} \rho_A^j \eta_{t-j}} \tag{27}$$

and

$$\frac{B_t}{B_0} = B_0^{\rho_B^t - 1} e^{\sigma_B \sum_{j=0}^{t-1} \rho_B^j \eta_{t-j}}. \tag{28}$$

As a result, the expectation over the product of these ratios reads

$$\begin{aligned} E_0\left[\frac{A_t}{A_0}\frac{B_t}{B_0}\right] &= A_0^{\rho_A^t - 1} B_0^{\rho_B^t - 1} e^{\sum_{j=0}^{t-1}(\sigma_A \rho_A^j + \sigma_B \rho_B^j)\eta_{t-j}} \\ &= A_0^{\rho_A^t - 1} B_0^{\rho_B^t - 1} e^{\frac{1}{2}\left[\sigma_A^2 \frac{1-\rho_A^{2t}}{1-\rho_A^2} + \sigma_B^2 \frac{1-\rho_B^{2t}}{1-\rho_B^2} + 2\sigma_A\sigma_B \frac{1-\rho_A^t\rho_B^t}{1-\rho_A\rho_B}\right]}. \end{aligned} \tag{29}$$

We plug this equation in our iterated Euler equation

$$\begin{aligned} \frac{P_0}{C_0} &= \sum_{t=1}^{\infty} E_0\left[\beta^t \frac{C_t^{1-\gamma}}{C_0^{1-\gamma}}\right] \cdot E_0\left[\frac{A_t}{A_0}\frac{B_t}{B_0}\right] \\ &= \sum_{t=1}^{\infty} \beta^t e^{-\gamma\mu t + \frac{\gamma^2\sigma_\varepsilon^2}{2} t} A_0^{\rho_A^t - 1} B_0^{\rho_B^t - 1} e^{\frac{1}{2}\left[\sigma_A^2 \frac{1-\rho_A^{2t}}{1-\rho_A^2} + \sigma_B^2 \frac{1-\rho_B^{2t}}{1-\rho_B^2} + 2\sigma_A\sigma_B \frac{1-\rho_A^t\rho_B^t}{1-\rho_A\rho_B}\right]} \end{aligned} \tag{30}$$

which yields our result in Eq. (21).

E2. Linear law of motion

From the linear law of motion (22) and the Euler Eq. (20), we receive the equation

$$\frac{P_0}{C_0} = E_0\left[\beta \frac{A_1}{A_0}\frac{B_1}{B_0}\left(\frac{C_1}{C_0}\right)^{1-\gamma}\left(1 + \bar{\phi}_{PC} + \phi_{PC}\frac{P_0}{C_0} + \phi_\eta \eta_1\right)\right] \tag{31}$$

that we need to solve. We rearrange it to

$$\frac{P_0}{C_0} = \frac{\beta e^{-\gamma\mu + \frac{\gamma^2\sigma_\varepsilon^2}{2}} E_0\left[\frac{A_1}{A_0}\frac{B_1}{B_0}(1 + \bar{\phi}_{PC} + \phi_\eta \eta_1)\right]}{1 - \phi_{PC}\beta e^{-\gamma\mu + \frac{\gamma^2\sigma_\varepsilon^2}{2}}} \tag{32}$$

and solve for the different parts. First note that

$$E_0\left[\frac{A_1}{A_0}\frac{B_1}{B_0}\right] = A_t^{\rho_A - 1} B_t^{\rho_B - 1} e^{\frac{1}{2}(\sigma_A + \sigma_B)^2} \tag{33}$$

which simplifies the denominator. For the numerator, we make use of the fact that consumption growth and growth of taste shocks are independent. Thus, we can treat the terms in the expectation separately. For the preference shocks, we get

$$E_1\left[\frac{A_1}{A_0}\frac{B_1}{B_0}(1 + \bar{\phi}_{PC} + \phi_\eta e^{\eta_1})\right] = A_1^{\rho_A - 1} B_1^{\rho_B - 1}\left(e^{\frac{1}{2}(\sigma_A+\sigma_B)^2}(1 + \bar{\phi}_{PC}) + \phi_\eta(\sigma_A + \sigma_B)e^{\frac{(\sigma_A+\sigma_B)^2}{2}}\right) \tag{34}$$

where the first part comes from a standard iteration as before. The second part follows from

$$E_1\left[\frac{A_1}{A_0}\frac{B_1}{B_0}\phi_\eta \eta_1\right] = A_0^{\rho_A - 1} B_0^{\rho_B - 1} \phi_\eta E_0\left[e^{(\sigma_A+\sigma_B)\eta_1}\eta_1\right]. \tag{35}$$

The last expectation can be computed by solving the integral

$$\begin{aligned} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} x e^{-\frac{1}{2}(x-(\sigma_A+\sigma_B))^2} e^{-\frac{x^2}{2}} dx &= e^{\frac{(\sigma_A+\sigma_B)^2}{2}} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} x e^{-\frac{1}{2}(x-(\sigma_A+\sigma_B))^2} dx \\ &= (\sigma_A + \sigma_B) e^{\frac{(\sigma_A+\sigma_B)^2}{2}}. \end{aligned} \tag{36}$$

E3. Quadratic law of motion

Households perceive the equation of motion to be

$$\frac{P_{t+1}}{C_{t+1}} = \bar{\phi}_{QPC} + \phi_{QPC1}\frac{P_t}{C_t} + \phi_{QPC2}\left(\frac{P_t}{C_t}\right)^2 + \phi_{Q\eta}\eta_{t+1}.$$

Plug the equation of motion into the Euler equation to get

$$\frac{P_t}{C_t} = E_t\left[\beta \frac{A_{t+1}}{A_t}\frac{B_{t+1}}{B_t}\left(\frac{C_{t+1}}{C_t}\right)^{1-\gamma}\left(\bar{\phi}_{\mathrm{QPC}} + \phi_{\mathrm{QPC1}}\frac{P_t}{C_t} + \phi_{\mathrm{QPC2}}\left(\frac{P_t}{C_t}\right)^2 + \phi_{\mathrm{Q}\eta}\eta_{t+1} + 1\right)\right]$$

or simplified as

$$\Lambda_t(\alpha+1) + (\rho_{p1}\Lambda_t - 1)\frac{P_t}{C_t} + \Lambda_t\rho_{p1}\left(\frac{P_t}{C_t}\right)^2 + E_t\left[\beta \frac{A_{t+1}}{A_t}\frac{B_{t+1}}{B_t}\left(\frac{C_{t+1}}{C_t}\right)^{1-\gamma}\delta\eta_{t+1}\right] = 0$$

where

$$\Lambda_t = E_t\left[\beta \frac{A_{t+1}}{A_t}\frac{B_{t+1}}{B_t}\left(\frac{C_{t+1}}{C_t}\right)^{1-\gamma}\right].$$

The formula for quadratic equations delivers two solutions, one of which is the desired one.

Appendix F. Barrier Functions

For an optimization problem of the form

$$\max_x f(x) \text{ subject to } c_i(x) \geq 0 \tag{37}$$

where i serves as an index for constraints c_i, we can write a smooth version as

$$\max_x f(x) - \mu \sum_i B(c_i(x)) \tag{38}$$

where $B(\cdot)$ is continuously differentiable and $\lim_{x\to 0} B(x) = \infty$. Take a sequence of barrier parameters $\{\mu_k\}$ which leads to a sequence of solutions $\{x_k^*\}$ to Eq. (38). Then every limit point x^* of the sequence of solutions $\{x_k^*\}$ is a global solution to problem (37). For details, e.g. see Nocedal and Wright (2006).

References

Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. Q. J. Econ. 109, 659–684.
Aruoba, S.B., Fernández-Villaverde, J., Rubio-Ramírez, J.F., 2006. Comparing solution methods for dynamic equilibrium economies. J. Econ. Dyn. Control 30 (12), 2477–2508. doi:10.1016/j.jedc.2005.07.008.
Bewley, T.F., 1977. The permanent income hypothesis: a theoretical foundation. J. Econ. Theory 16, 252–292.
Chiarella, C., Dieci, R., Gardini, L., 2006. Asset price and wealth dynamics in a financial market with heterogeneous agents. J. Econ. Dyn. Control 30 (9–10), 1755–1786. https://doi.org/10.1016/j.jedc.2005.10.011.
Chiarella, C., He, X.-Z., 2001. Asset price and wealth dynamics under heterogeneous expectations. Quant. Financ. 1, 509–526.
Chien, Y., Cole, H., Lustig, H., 2010. A multiplier approach to understanding the macro implications of household finance. Rev. Econ. Stud. 78, 199–234.
Chien, Y., Lustig, H., 2010. The market price of aggregate risk and the wealth distribution. Rev. Financ. Stud. 23(4), 1596–1650.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. J. Polit. Econ. 113 (1), 1–45.
Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. J. Polit. Econ. 104 (2), 219–240.
Den Haan, W.J., 2010. Comparison of solutions to the incomplete markets model with aggregate uncertainty. J. Econ. Dyn. Control 34, 4–27.
Den Haan, W.J., Judd, K.L., Juillard, M., 2010. Computational suite of models with heterogeneous agents: incomplete markets and aggregate uncertainty. J. Econ. Dyn. Control 34, 1–3.
Dumas, B., Lyasoff, A., 2012. Incomplete market equilibria solved recursively on an event tree. J. Financ. 67 (5), 1897–1941.
Favilukis, J., Ludvigson, S., van Nieuwerburgh, S., 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. J. Polit. Econ. 125 (1), 140–223.
Feng, Z., Miao, J., Peralta-Alva, A., Santos, M.S., 2014. Numerical simulation of nonoptimal dynamic equilibrium models. Int. Econ. Rev. (Philadelphia) 55 (1), 83–110.
Fernández-Villaverde, J., Rubio-Ramírez, J.F., 2006. Solving DSGE models with perturbation methods and a change of variables. J. Econ. Dyn. Control 30 (12), 2509–2531.
Garlappi, L., Skoulakis, G., 2010. Solving consumption and portfolio choice problems: the state variable decomposition method. Rev. Financ. Stud. 23(9), 3346–3400.
Gomes, F., Michaelides, A., 2008. Asset pricing with limited risk sharing and heterogeneous agents. Rev. Financ. Stud. 21 (1), 415–448.
Gomes, J.F., Schmid, L., 2010. Equilibrium Credit Spreads and the Macroeconomy. Mimeo, Fuqua School of Business, Duke University.
Hassan, T.A., Mertens, T.M., 2017. The social cost of near-rational investment. Am. Econ. Rev 107 (4), 1059–1103.
Heathcote, J., Storesletten, K., Violante, G.L., 2014. Consumption and labor supply with partial insurance: an analytical framework. Am. Econ. Rev. 104 (7), 2075–2126.
Jin, H., Judd, K.L., 2002. Perturbation Methods for General Dynamic Stochastic Models. Mimeo, Hoover Institution, Stanford University.
Judd, K.L., 2002. Perturbation Methods with Nonlinear Change of Variables. Mimeo, Hoover Institution, Stanford University.
Judd, K.L., Guu, S.-M., 2000. The Economic Effects of New Assets: An Asymptotic Approach. Mimeo, Hoover Institution.
Judd, K. L., Guu, S.-M., 2001. Asymptotic Methods for Asset Market Equilibrium. NBER Working Paper 8135.
Judd, K. L., Maliar, L., Maliar, S., 2009. Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models. NBER Working Paper 15296.
Judd, K. L., Maliar, L., Maliar, S., 2010. A Cluster-Grid Projection Method: Solving Problems with High Dimensionality. NBER Working Paper 15965.
Kim, S.H., Kollmann, R., Kim, J., 2010. Solving the incomplete market model with aggregate uncertainty using a perturbation method. J. Econ. Dyn. Control 34, 50–58.
Krusell, P., Smith, A.A., 1998. Income and wealth heterogeneity in the macroeconomy. J. Polit. Econ. 106 (5), 867–896.
Lustig, H., van Nieuwerburgh, S., 2010. How much does household collateral constrain regional risk sharing? Rev. Econ. Dyn. 13 (2), 265–294.
Mertens, T.M., 2009. Volatile Stock Markets: Equilibrium Computation and Policy Analysis. Mimeo, NYU Stern.
Moll, B., 2009. Creditor Rights, Inequality and Development in a Neoclassical Growth Model. Mimeo, University of Chicago.

Nocedal, J., Wright, S.J., 2006. Numerical Optimization, second ed. Springer.
Panageas, S., 2011. Discussion of 'the macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium'. Discussion given at the NBER Asset Pricing Spring Meeting, Chicago, IL.
Pavlova, A., Rigobon, R., 2007. Asset prices and exchange rates. Rev. Financ. Stud. 20 (4), 1139–1180.
Preston, B., Roca, M., 2007. Incomplete Markets, Heterogeneity and Macroeconomic Dynamics.
Reiter, M., 2009. Solving heterogeneous-agent models by projection and perturbation. J. Econ. Dyn. Control 33 (3), 649–665.
Storesletten, K., Telmer, C.I., Yaron, A., 2007. Asset pricing with idiosyncratic risk and overlapping generations. Rev. Econ. Dyn. 10, 519–548.
Swanson, E., Anderson, G., Levin, A., 2005. Higher-Order Perturbation Solutions to Dynamic, Discrete-Time Rational Expectations Models. Mimeo, Federal Reserve Bank of San Francisco.
Taylor, M.E., 2010. Partial Differential Equations I: Basic Theory, 2nd Springer.