Journal of Economic Dynamics & Control 000 (2018) 1–20

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 xxx

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

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