Exactly Solved Economies with Heterogeneity

Jules H. van Binsbergen Christian C. Opp*

January 17, 2019

Abstract

We propose a novel modeling approach for general equilibrium economies with persistent heterogeneity that yields exact solutions. This is an important advancement relative to previous approaches where approximations are necessary, either by summarizing the distributional state space with a fixed number of moments or, alternatively, by perturbing the problem around a given fixed point for which the solution is known. As our approach does not impose any restriction on the shape of the state variable distribution, it can also be used to evaluate the conditions under which previous solution methods are likely to succeed.

Keywords: Persistent heterogeneity, exact solutions, occasionally binding constraints, nonlinearities

*The authors are at the Wharton School of the University of Pennsylvania and can be reached at julesv@wharton.upenn.edu and opp@wharton.upenn.edu. We thank Urban Jermann, Vincent Glode and Tom Sargent for helpful comments and suggestions.

1. Introduction

The increasing availability of micro data on the behavior of individual agents and firms has spurred the need to incorporate important observed heterogeneity in macroeconomic models. However, due to the complexity that this heterogeneity induces, simplifying assumptions are usually imposed. The purpose of these simplifications is to reduce the infinite dimensionality of the state space of the model (i.e., an entire state space distribution) without sacrificing the key elements of the problem at hand. One often-imposed assumption is the completeness of financial markets, which for many applications in economics and finance allows agents to be aggregated into a single representative agent. While the complete markets assumption makes solving models typically tractable, it is not ex ante obvious whether interesting features of the economic environment are assumed away. In other cases, market incompleteness is maintained and approximate model solutions are introduced to keep the analysis tractable (e.g., in Krusell and Smith (1998)). This latter approach has led to a literature debating under what circumstances the accuracy of the approximated aggregate law of motion suffices and under what circumstances it does not.[1]

In this paper, we introduce a novel approach to modeling general equilibrium economies with heterogeneity. In the class of models we propose, exact solutions are available, without imposing ex ante restrictions on the degree of persistent heterogeneity that an economy can feature. Our approach can easily handle features that have traditionally been challenging for existing solution methods, such as, non-linearities, occasionally binding constraints, and lumpy capital adjustment. These are important innovations relative to previous approaches that provide only approximate solutions, either by summarizing the distributional state space with a fixed number of moments (Krusell and Smith (1998)) or alternatively, by perturbing the problem around a given fixed point for which the solution is known (Judd and Mertens (2019)). The

[1] See, e.g., Den Haan (2010a) for a discussion.

validity of these approximation methods has been questioned in the literature particularly in environment with the above-mentioned features.

To illustrate this modeling approach, we consider a classic Krusell and Smith (1998) type economy. Recasting this type of economy with our approach involves specifying agents and technology such that both capital and distributions of agents take values in discrete sets. Following the *Stochastic Lumpy Adjustments* (SLA) technology proposed in Binsbergen and Opp (Forthcoming), agents control hazard rates with which they upgrade or downgrade their capital stocks by lumpy increments. Moreover, agents have positive measure. As both the measure of each agent and the size of capital increments can theoretically be set to arbitrarily small values, this environment can capture any degree of granularity appropriate for an application. While this setting allows evaluating interesting theoretical limiting cases where granularity approaches zero, lumpiness is a desirable feature in many relevant economic problems: almost any type of economic entity in practice exhibits some degree of granularity (e.g., agents and firms). In fact, a growing literature that documents the empirical relevance of power laws highlights that in cross-sectional distributions of firms and agents, the largest and wealthiest entities account for a large fraction of economic activity.[2]

Solutions to the proposed class of models are fully characterized by continuous-time Markov generator matrices. These matrices are generically *sparse*, a result that emerges in our continuous-time environment where any state-variable can locally move to only a relatively small set of neighboring states. Sparsity, in turn, dramatically increases the numerical efficiency of inverting large matrices, which is the only operation required to obtain exact model solutions. Moreover, given the Markov generator matrix characterizing the economy, exact conditional and stationary distributions are immediately available. As a result, the proposed class of models can be evaluated and estimated without the need to use time-consuming simulations.

It may be tempting to view our approach as another way of approximating existing models

[2] See for example Gabaix (2011) and the references therein.

using a discretized state space, much like previous grid-based solution methods. However, our contribution is different. By incorporating the discreteness of the state space as an inherent part of the economic model, the model does not have to be approximated. This has the important advantage of not having to second guess whether an approximation is sufficiently accurate. Instead, the assumptions of the model are clearly stated upfront, and conditional on those assumptions, the model solution is exact. Further, as we have argued above, we view discreteness as an important empirical regularity, that is, agents, firms, and investment opportunities are in fact not continuous objects in practice.

While our class of models requires inverting increasingly large matrices when studying environments with a large number of possible state variable realizations and/or a large number of agents, recent and projected advances in computational capabilities (e.g., via parallel computing) ameliorate such concerns. Further, even given the currently available computational power, the fact that the model is exactly solvable provides a number of important benefits relative to existing work. First, with our approach, researchers can gradually increase the degree of heterogeneity while maintaining exact solutions, allowing to assess the economic relevance of heterogeneity. Second, our method is also ideally suited to (re-)evaluate how well previous numerical solution methods fare when used to approximate our exactly solved class of models. Our approach does not impose any restrictions on the shape of the state variable distribution. It can therefore be used to evaluate the conditions under which previous solution methods are likely to succeed. While our results from the Krusell and Smith (1998) type environment are very encouraging, it is straightforward to apply this method to economies with material nonlinearities, which are particularly challenging for existing solution approaches.

In our paper, we focus on a Krusell and Smith (1998) (KS) type economy that features idiosyncratic labor (employment) shocks against which agents cannot fully insure.[3] While insurance is imperfect, agents can buy and sell an asset (capital) subject to an exogenous lower

[3]See, e.g., Bewley (1977), Scheinkman and Weiss (1986).

bound on assets holdings (Aiyagari (1994). KS argue that in their environment, the utility costs from fluctuations in consumption are quite small and that this finding is consistent with a previous literature that suggests that self-insurance with only one asset is quite effective.[4]

Even though the findings in KS are similar to models where self-insurance is effective, the results are not in fact driven by a highly effective self-insurance channel. In their model, the unconditional standard deviation of individual consumption is about four times that of aggregate consumption, and the unconditional correlation of the consumption of any two agents is very close to zero. Instead, the reason for their findings is that in their stationary equilibria, agents are insured well *enough* that the marginal propensity to save out of current wealth is almost completely independent of the levels of wealth and labor income, except at the very lowest levels of wealth. While some very poor agents therefore do have substantially different savings rates, this heterogeneity is not important enough to materially affect the equilibrium due to the small wealth that these agents possess.

KS also suggest a range of extensions that could potentially increase the importance of wealth (capital) heterogeneity including (1) models with endogenous borrowing constraints (such as in Kiyotaki and Moore (1997)) (2) models where the return to savings is dependent on consumers' own production technology rather than the aggregate savings technology (where the return is equal for all agents), and (3) fixed costs in capital accumulation (Banerjee and Newman (1993)).

Since the KS paper, the literature on heterogeneity in finance and macroeconomics has substantially evolved. Examples of methods largely inspired by the KS method include Storesletten et al. (2007), Gomes and Michaelides (2008), and Favilukis et al. (2017). In addition, new approaches for modeling and/or solving environments with heterogeneous agents have been developed in Den Haan et al. (2010), Den Haan (2010b), Judd et al. (2010), and Judd and Mertens

[4]See, e.g., Lucas (1987), Cochrane (1989), and Krusell and Smith (1996) as well as the asset pricing literature that focuses on incomplete markets (Marcet and Singleton (1999), Telmer (1993), Lucas (1994), Heaton and Lucas (1996), den Haan (1996), and Krusell and Smith (1997)).

(2019). While the upside of the most recent solution methods is that they can handle a large(r) number of agents and/or state variables, they do still rely on approximations in one form or another. Notable exceptions that do generate closed-form solutions are Heathcote et al. (2014) and Han et al. (2018). The former maintains closed-form solutions by setting up the problem such that (i) individual wealth is a redundant state variable, and (ii) agents have access to perfect insurance against some shocks and no explicit insurance against others, the latter of which is particular to the island structure that they assume. To achieve this insurance dichotomy as an equilibrium outcome, the island-economy structure is important. The latter who recast the Aiyagari-Bewley-Huggett model of income and wealth distribution in continuous time and show for the special case of two income groups how to maintain closed-form solutions.

2. The Economy

The economy we consider is in continuous time. There is a measure one population of infinitely lived consumers. There is only one good, and preferences over streams of consumption of each agent are given by

$$\mathbb{E}_0 \int_0^\infty e^{-\beta\tau} U(C_\tau) d\tau, \tag{1}$$

where the flow utility is given by CRRA preferences

$$U(C) = \frac{C^{1-\gamma} - 1}{1-\gamma}. \tag{2}$$

Production of the good, Y, follows from a Cobb-Douglas function of capital, K, and labor input L

$$Y_t = Z_t K_t^\alpha L_t^{1-\alpha}, \tag{3}$$

with $\alpha \in [0, 1]$.

Groups. The measure one of agents consists of n_g groups of equal measure. Shocks to agents occur at the group-level (including investment and (un)employment shocks). As the number of groups is increased, this setup converges to a specification where agents are atomistic. That said, the specification also accommodates economic shocks that affect larger groups of agents (for example, when a large firm goes bankrupt). For notational simplicity, we drop subscripts indicating groups whenever doing so does not create ambiguity.

Investment in capital. Two essential features of the proposed modeling approach are (1) that capital investment decisions are made at the group level, and (2) that the investment technology is of the SLA type, following Binsbergen and Opp (Forthcoming). Under this technology, log-capital $k_t = \log[K_t]$ takes values in a discrete set Ω_k, the n_k elements of which constitute an equidistant grid with lower bound $\min\{\Omega_k\} \in \mathbb{R}$ and grid increments of size $\Delta_k > 0$.[5] By choosing Δ_k small enough, this specification can approximate a continuous support for capital arbitrarily well. Moreover, the discrete state space structure allows capturing the lumpy nature of investment observed in the data.

Let $N^+_{k,t}$ and $N^-_{k,t}$ denote Poisson processes that keep track of successful capital acquisitions and divestments, and let $N^\delta_{k,t}$ denote a Poisson process for capital depreciation shocks. The corresponding capital evolution equation is given by

$$dk_t = \Delta_k \cdot (dN^+_{k,t} - dN^-_{k,t} - dN^\delta_{k,t}). \tag{4}$$

Groups incur flow costs (payoffs) when aiming to upgrade (downgrade) their capital stocks. Specifically, they control the Poisson intensities of the processes $N^+_{k,t}$ and $N^-_{k,t}$. Each group

[5] This technology specification can also easily accommodate grids of log-capital that are non-equidistant.

chooses its expected investment rate

$$i_t^+ \equiv (e^{\Delta_k} - 1)\mathbb{E}_t \left[dN_{k,t}^+\right] \geq 0, \tag{5}$$

and stochastically succeeds in upgrading its capital to the next-higher level, that is, by a log change of size Δ_k, with Poisson intensity $i_t^+/(e^{\Delta_k} - 1)$. Throughout, $\mathbb{E}$ denotes the expectation operator.

In the process of attempting to upgrade capital, agents within a group incur a flow cost equal to $i_t^+ K_t$. Once an upgrade occurs ($dN_{k,t}^+ = 1$), the additional capital is useable immediately, at no additional costs. When a group's capital stock reaches the upper bound $e^{\max\{\Omega_k\}}$, further upgrades are infeasible. By choosing n_k high enough, this restriction will be immaterial, as optimal investment will be zero above some endogenous threshold for capital.

Similarly, groups choose their expected disinvestment rate

$$i_t^- \equiv (1 - e^{-\Delta_k})\mathbb{E}_t \left[dN_{k,t}^-\right] \geq 0, \tag{6}$$

and downgrade their capital by a log change of size Δ_k with Poisson intensity $i_t^-/(1 - e^{-\Delta_k})$. The divestment process yields flow proceeds equal to $i_t^- K_t$. Conditional on a Poisson arrival, there is no additional payment to the agents in the group. Divestments are infeasible when capital reaches the lower bound $e^{\min\{\Omega_k\}}$. Again, by specifying $\min\{\Omega_k\}$ low enough, we can ensure that a firm would never optimally attempt to divest at this lower bound, such that this restriction is also non-binding.

Capital depreciates stochastically to the next-lower level with a Poisson intensity $\delta/(1 - e^{-\Delta})$, except at the lower bound $\min\{\Omega_k\}$ where the Poisson intensity is zero. Thus, for $k_t > \min\{\Omega_k\}$, the expected depreciation rate is δ, and the expected growth rate of capital is given by

$$\frac{\mathbb{E}_t[dK_t]}{K_t} = (i_t^+ - i_t^- - \delta)dt. \tag{7}$$

We introduce the generator matrix Λ_k that collects the transition rates between all capital states Ω_k. This matrix depends on the endogenous investment controls $i = (i^+, i^-)$.

It is straightforward to introduce additional adjustment costs in this setting (as specified in Binsbergen and Opp (Forthcoming)). We choose not to introduce them at this point to stay closer to the original Krusell and Smith (1998) framework.

Budget constraint. The flow of output Y can be used for consumption C or for investment

$$Y_t = C_t + i_t^+ K_t - i_t^- K_t. \tag{8}$$

Aggregate productivity. There is also a stochastic shock to aggregate productivity, which is denoted by Z. Z follows an N_Z state continuous-time Markov chain with generator Matrix Λ_Z.

Labor supply. Each agent in a group is endowed with $\epsilon\tilde{L}$ units of labor input, where ϵ is stochastic and can take on the value zero or one at the group level. When $\epsilon = 1$, each agent in the group is employed and supplies $\tilde{L}$ units of labor. When $\epsilon = 0$, agents in the group are unemployed. The generator matrix governing employment shocks is given by $\Lambda_\epsilon(Z)$, which indicates that the transition rates can depend on the aggregate state Z.

Market Arrangement. As in Krusell and Smith (1998), the economy features incomplete markets. Capital is the only asset in which agents can invest. As a result, capital is used both as a store of value and as a means of insurance against employment shocks.

Equilibrium prices. Consumers collect income from working and from the services of their capital. Let the total amount of capital in the economy be denoted by $\bar{K}$ and the total amount of labor supplied by $\bar{L}$, then the constant returns-to-scale production function implies that the relevant prices are given by:

$$w(\bar{K}, \bar{L}, Z) = (1 - \alpha) \cdot Z \cdot \left(\frac{\bar{K}}{\bar{L}}\right)^{\alpha}, \tag{9}$$

$$r(\bar{K}, \bar{L}, Z) = \alpha \cdot Z \cdot \left(\frac{\bar{K}}{\bar{L}}\right)^{\alpha-1}. \tag{10}$$

State variables and transition dynamics. The aggregate state of the economy consists of (1) the current value of the state Z, and (2) the distribution of the n_g groups across $n_k \times 2$ states (recall that there are two possible employment statuses for a group). The number of possible distributions of n_g groups across $n_k \times 2$ states is:

$$n_h = \frac{(n_g + 2n_k - 1)!}{n_g!(2n_k - 1)!} \tag{11}$$

Thus, the total number of possible aggregate states is $n = n_h \times n_Z$.

Say $s_h \in \{1, ..., n_h\}$ denotes a particular histogram index, and let $\mathbf{h}(s_h, j)$ denote the j-th element of the vector of length $n_\epsilon \times n_k$ representing the histogram buckets. Finally, let $K(j)$ denote the capital level associated with the j-th histogram bucket. Aggregate capital is thus given by the sum of the products of capital levels in each histogram state times the number of groups in that state.

$$\bar{K}(s_h) = \sum_{j=1}^{2n_k} \mathbf{h}(s_h, j) \cdot K(j) \tag{12}$$

Similarly, the aggregate labor supply is given by

$$\bar{L}(s_h) = \sum_{j=1}^{2n_k} \mathbf{h}(s_h, j) \cdot L(j), \tag{13}$$

where $L(j)$ denotes the labor supply of a group that is in the j-th bucket.

For an individual group, the relevant state variables are its own log-capital k, its employment status ϵ, and the aggregate state index $s \in \{1, ..., n\}$. Let Λ denote the generator matrix collecting all exogenous and endogenous transition rates between the possible values of the state tuple (k, ϵ, s). We will use the notation $\Lambda(k, \epsilon, s)$ to refer to the row of the matrix Λ corresponding to a specific state (k, ϵ, s).

Optimization. The Hamilton-Jacobi-Bellman (HJB) equation faced by a group is given by:

$$0 = \max_{\substack{C(k,\epsilon,s)\geq 0 \\ i^+(k,\epsilon,s)\geq 0 \\ i^-(k,\epsilon,s)\geq 0}} \left\{ U(C(k, \epsilon, s)) - \beta V(k, \epsilon, s) + \Lambda(k, \epsilon, s)\mathbf{V} \right\}, \tag{14}$$

where the notation $\mathbf{V}$ denotes a vector that collects the values of $V(k, \epsilon, s)$ for all possible values of the state tuple (k, ϵ, s), and where the ordering of states follows the same convention as the one used for the matrix Λ. Note that the endogenous investment controls $(i^+(k, \epsilon, s), i^-(k, \epsilon, s))$ enter the vector $\Lambda(k, \epsilon, s)$. The budget constraint for an individual in a group is given by:

$$r(\bar{K}_t, \bar{L}_t, Z_t)K_t + w(\bar{K}_t, \bar{L}_t, Z_t)\tilde{L}_t\epsilon_t = C_t + i_t^+ K_t - i_t^- K_t. \tag{15}$$

It is convenient to define the amount of consumption absent investment, that is, the *base level* of consumption

$$C_b(k, \epsilon, s) = r(\bar{K}, \bar{L}, Z)K + w(\bar{K}, \bar{L}, Z)\tilde{L}\epsilon. \tag{16}$$

The first-order conditions with respect to i^+ and i^- yield the optimal investment controls:

$$i^+(k,\epsilon,s) = \frac{1}{K}\max\left[C_b(k,\epsilon,s) - \left(\frac{V(k+\Delta_k,\epsilon,s^{k_g+}) - V(k,\epsilon,s)}{K(e^{\Delta_k}-1)}\right)^{-\frac{1}{\gamma}}, 0\right], \tag{17}$$

$$i^-(k,\epsilon,s) = \frac{1}{K}\max\left[\left(\frac{V(k,\epsilon,s) - V(k-\Delta_k,\epsilon,s^{k_g-})}{K(1-e^{-\Delta_k})}\right)^{-\frac{1}{\gamma}} - C_b(k,\epsilon,s), 0\right], \tag{18}$$

where s^{k_g+} and s^{k_g-} denote the new aggregate state indices that obtain when the individual's group capital level is increased and decreased by one increment, respectively (which also changes the histogram of groups in the economy). Optimal consumption is given by:

$$C(k,\epsilon,s) = C_b(k,\epsilon,s) - i^+(k,\epsilon,s)K + i^-(k,\epsilon,s)K. \tag{19}$$

3. An Exactly Solved Economy

In this section we solve the model using four capital states and seven groups of agents. The particular grid values for capital are given by the set $\{0.84, 1.41, 2.38, 4.00\}$, which corresponds to a log capital increment $\Delta_k = 0.52$. The transition rates for the aggregate states (boom and bust) as well as the productivity factor (Z) are summarized in the top panel of Table 1. The bottom panel of the table summarizes the state invariant parameters, which we all set to standard values. The annual depreciation rate δ is set to 12%, the rate of time preference β is set to 0.1 (corresponding to a log discount rate of 0.9), and the capital share in the Cobb-Douglas production function is set to 0.6. The transition rates in and out of employment are set such that the average employment rate is 5%, and employed agents stay employed for 3.33 years on average.

Before we present the exact solutions to the model, it is important to relay the two types of cross-sectional heterogeneity that emerge in this setting. In this environment, the joint distri-

Table 1

Parameters. The table lists parameters of the economy. The economy features $N_g = 7$ groups and $N_k = 4$ capital states at the individual level taking the values in the set $\{0.84, 1.41, 2.38, 4.00\}$, which corresponds to $\Delta_k = 0.52$.

State-dependent Parameters			
Parameter	Variable	Recession	Boom
Transition rates for aggregate states	λ_Z	0.500	0.100
Productivity factor	Z	1.100	1.200

State-invariant Parameters		
Parameter	Variable	Values
Rate of time preference	β	0.100
Coefficient of relative risk aversion	γ	3.000
Capital share	α	0.600
Labor supply	$\bar{L}$	1.000
Expected rate of depreciation	δ	0.120
Transition rate of becoming unemployed	$\lambda_{\epsilon=0}$	0.300
Transition rate of becoming employed	$\lambda_{\epsilon=1}$	5.700

bution of capital and employment matters. The employment status does not directly pin down a group's capital, as capital upgrades and downgrades occur stochastically, following the SLA technology, and because agents optimally spread out efforts to adjust their capital over time (to smooth consumption). Thus, even if the impact of investment on capital were deterministic (a limiting case of our model where $\Delta_k \to 0$), the joint distribution of capital and employment would matter.

To illustrate the relevance of this joint distribution, compare the histograms in Figure I. The two histograms both represent an economy with the same aggregate capital stock $\bar{K}$, the same cross-sectional variation in the capital stock, and the same aggregate labor supply $\bar{L}$. They only differ in how capital and employment are related.

The second type of heterogeneity that the model features, is a standard one: two histograms of capital that differ in shape but represent the same aggregate amount of capital $\bar{K}$. Figure II below illustrates two such capital distributions, which each have aggregate capital approximately equal to 14.[6]

FIGURE I

Joint distribution of capital and employment. The histograms illustrate two identical capital distributions with identical aggregate employment in which the unemployed own differing amounts of capital.

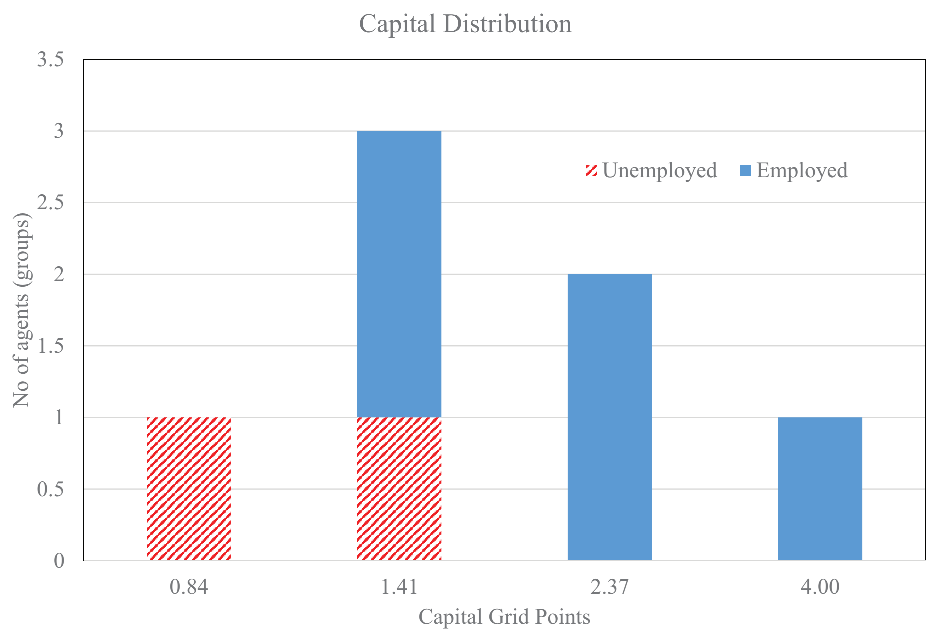

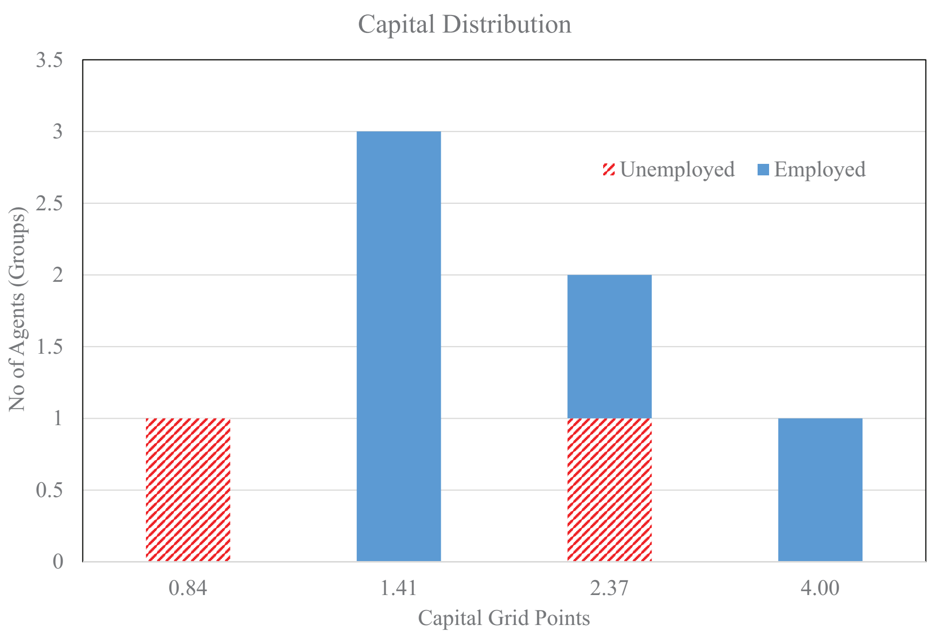

[6]Note that due to the recombination properties of our equidistant log-capital grid, it is not possible to get exactly the same amount of aggregate capital for any two distinct histograms. The amounts of aggregate capital are in fact 14.3 and 13.8, respectively.

FIGURE II

Shape of the capital distribution. The two histograms illustrate two capital distributions that imply approximately the same amount of aggregate capital $\bar{K}$.

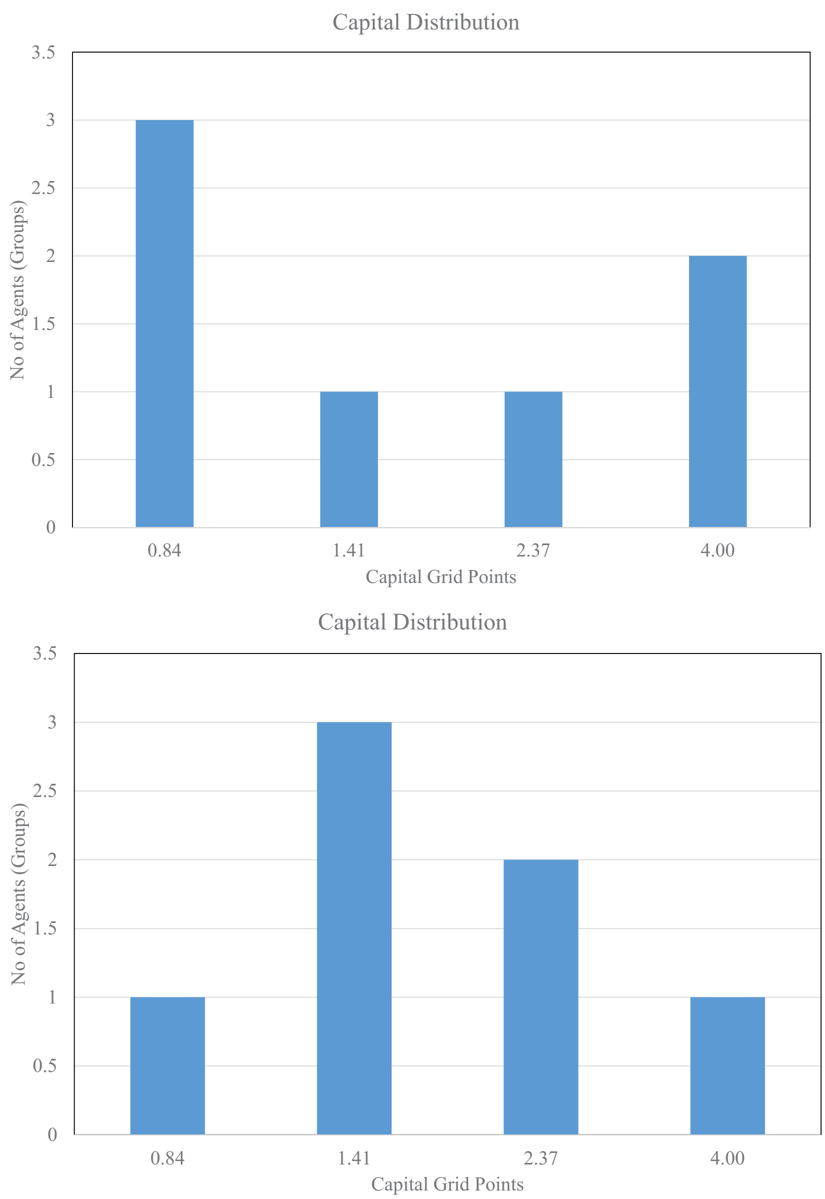

We now discuss the solution to the model. We start considering the stationary distribution of aggregate capital $\bar{K}$, which is plotted in panel (a) of Figure III. With our modeling approach, this distribution is available in closed-form, once the matrix Λ is determined. Despite the simple example, the shape of the stationary distribution is quite smooth; even a setting with just seven groups and four capital grid points leads to significant diversification effects at the aggregate level. Roughly 80% of the probability mass is concentrated between $\bar{K} = 1.8$ and $\bar{K} = 2.8$. Panel (b) of Figure III illustrates the stationary distribution of labor, which implies a unconditional unemployment rate of 5%.

FIGURE III

Stationary distribution of aggregate capital and labor. The graph plots the stationary distribution of aggregate capital $\bar{K}$ (panel (a)) and aggregate labor $\bar{L}$ (panel (b)).

(a) Aggregate capital **(b)** Aggregate labor

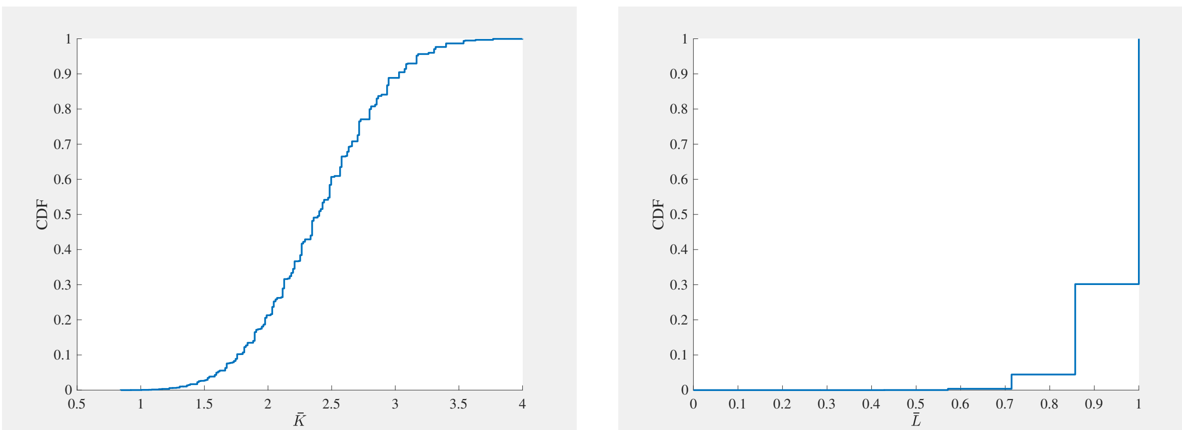

Next, we consider plots that illustrate a group's optimal consumption policy as a function of the aggregate amount of capital in the economy ($\bar{K}$). In each graph we additionally condition on a particular aggregate employment level $\bar{L}$ and productivity Z, and an individual group's own employment status and capital. For ease of exposition, we focus on states that have a large probability of occurring under the stationary distribution.[7] As a consequence, we will focus on cases where the aggregate labor supply $\bar{L}$ is 0.86 or 1.[8] Further, we focus on the case where the individual group has capital intermediate levels of capital K.

[7]Needless to say that our method also provides exact solutions for low probability state realizations, they are just less interesting to study when comparing solutions.

[8]For example, states with aggregate employment less than 0.86 have strictly positive probabilities, these probabilities are extremely small given the parameter values and thus hardly affect the solution. This is not surprising as those states correspond to aggregate unemployment rates of 29% or higher.

FIGURE IV

Optimal consumption as function of aggregate capital. The graph plots the optimal consumption policies (blue circles) as well as the best linear approximation (solid black line) for states in which the individual is employed ($\epsilon = 1$) and aggregate employment employment is $\bar{L} = 1$. The panels illustrate policies for different individual capital levels ($K = 1.41$ versus $K = 2.38$) and different aggregate states ($Z = G$ versus $Z = B$). We fit linear curves (black lines) using a weighted least squares minimization, where the weights are the conditional probabilities of a particular $\bar{K}$-consumption pair occurring under the stationary distribution. We report the corresponding R^2 values in the graphs.

(a) $K = 1.41$, $Z = G$

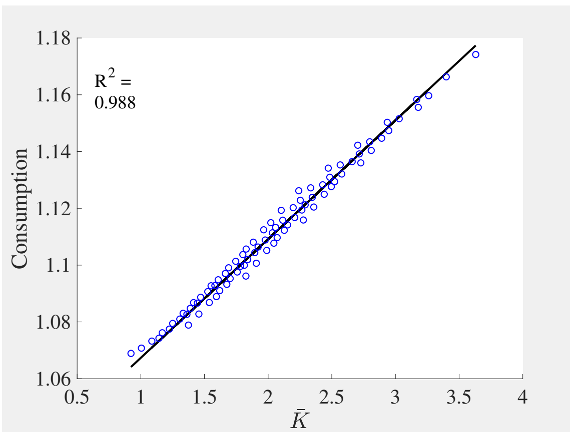

(b) $K = 1.41$, $Z = B$

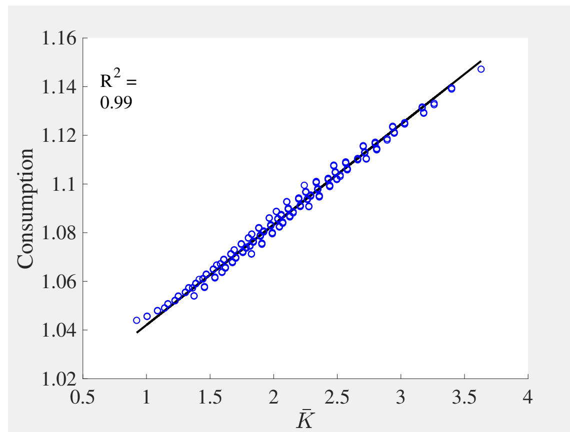

(c) $K = 2.38$, $Z = G$

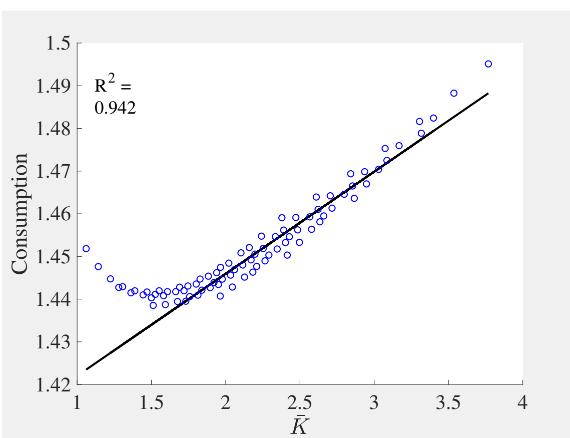

(d) $K = 2.38$, $Z = B$

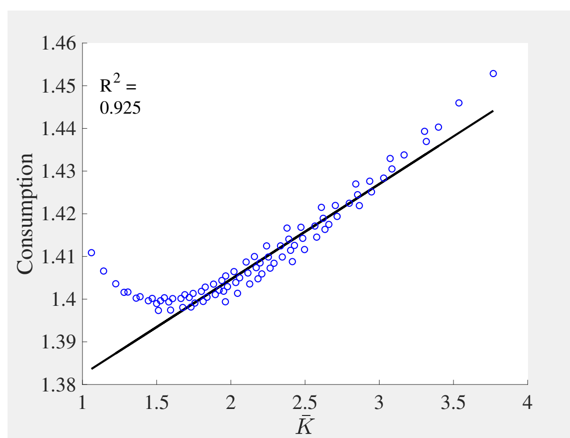

In all panels of Figure IV the individual group is employed. Each blue circle in a graph represents a particular joint histogram of capital and employment, keeping aggregate employment fixed (as illustrated in Figure I). The aggregate capital implied by such a distribution is plotted on the x-axis. As discussed in reference to Figures I, even conditional on identical histograms of capital and aggregate employment, the optimal consumption policy can still differ depending on how capital and employment are *jointly* distributed. As a result, two blue circles can have the same value of $\bar{K}$ but a different optimal consumption value (as the underlying states differ). Moreover, as highlighted in Figure II, a second source of variation in optimal consumption given a particular level of $\bar{K}$ stems from different possible shapes of the capital distribution. The full employment states illustrated in Figure IV isolates this latter type of variation.

The panels on the left ((a) and (c)) of Figure IV condition on the aggregate boom state ($Z = G$), whereas the panels on the right ((b) and (d)), condition on a recession ($Z = B$). Finally, the top row conditions on a group with the capital level $K = 1.41$, whereas the bottom row conditions on a group with a higher capital level of $K = 2.38$. At the lower capital level (panels (a) and (b)), the relation between the group's consumption and aggregate capital is almost perfectly linear; the R^2 of a linear fitted curve (black line), computed based on the stationary probability distribution, is close to 1.

In contrast, at the higher capital level (panels (c) and (d)), the group starts to consume more in states where the aggregate economy has little capital ($\bar{K} < 1.5$). When aggregate capital is lower, the return on capital r is higher, and the group with $K = 2.38$ is a positive outlier in the capital distribution. As a result, this group has a particularly high level of income (that is, the group's base consumption C_b increases as $\bar{K}$ as decreased). While the group invests a positive, increasing amount for low levels of $\bar{K}$, this choice does not fully undo the increases in total income for low levels of $\bar{K}$, causing the U-shaped patterns in consumption shown in panels (c) and (d) of Figure IV. A linear approximation to the conditional consumption policy then provides a significantly worse fit — the R^2 drops below 95%. Overall, the results suggest that if agents were to approximate the aggregate capital distribution using the first moment ($\bar{K}$) only, they would deviate substantially from the optimal consumption policies. Thus, in the present environment, standard approximation techniques are unlikely to provide accurate solutions.

In Figure V, we provide similar graphs, but consider for the case in which a group is unemployed. When the group's own capital level is $K = 1.41$ (panels (a) and (b)), the group divests

FIGURE V

Optimal consumption as function of aggregate capital. The graph plots the optimal consumption policies (blue circles) as well as the best linear approximation (solid black line) for states in which the individual is unemployed ($\epsilon = 0$) and aggregate employment employment is $\bar{L} = 0.86$. The panels illustrate policies for different individual capital levels ($K = 1.41$ versus $K = 2.38$) and different aggregate states ($Z = G$ versus $Z = B$). We fit linear curves (black lines) using a weighted least squares minimization, where the weights are the conditional probabilities of a particular $\bar{K}$-consumption pair occurring under the stationary distribution. We report the corresponding R^2 values in the graphs.

(a) $K = 1.41$, $Z = G$

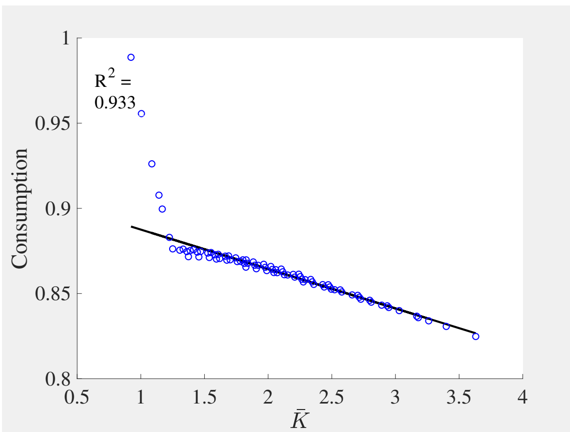

(b) $K = 1.41$, $Z = B$

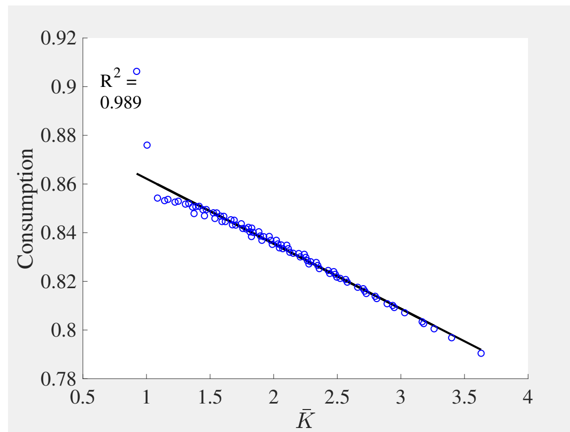

(c) $K = 2.38$, $Z = G$

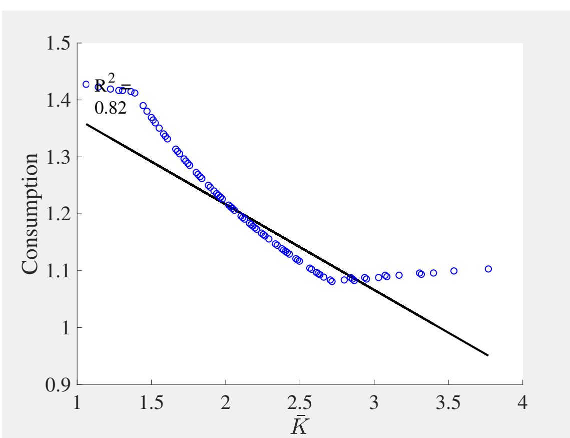

(d) $K = 2.38$, $Z = B$

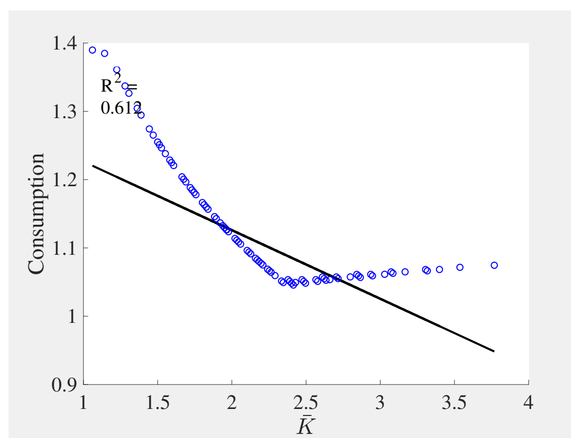

positive amounts (for $\bar{K} > 1.25$ in panel (a) and $\bar{K} > 1.1$ in panel (b)) in order to obtain a consumption level in excess of the income from capital. Once aggregate capital is very low and the return on capital very high (for $\bar{K} < 1.25$ in panel (a) and $\bar{K} < 1.1$ in panel (b)), the group stops divesting and switches to simply consuming its income, that is, $C = C_b$. This switch occurs at the kink that is visible in the consumption plots. This type inaction region is a natural consequence of the lumpy capital adjustment decisions also observed in the data.[9] At a higher level of capital ($K = 2.38$, see panels (c) and (d)), the individual group has a higher level of income C_b. With more disposable income, the group then has greater incentives to invest some of its income, relative to the case where the group has less capital (that is, $K = 1.41$, see panels (a) and (b)). As a result, when the return on capital is particularly high (for $\bar{K} < 1.5$ in panel (c) and $\bar{K} < 1.1$ in panel (d)), the group invests positive amounts. For intermediate levels of aggregate capital $\bar{K}$ the group simply consumes its income, and finally, for high enough values of $\bar{K}$, it starts to divest (starting at the second kink).

4. Additional Computational Efficiency Gains

Further efficiency gains in computational speed can be obtained by employing the characteristics of the particular economic problem. While the number of theoretically possible distributional shapes of the state variables can be large (through the combinatorics in equation 11), the effective number of histograms that are relevant is typically much smaller, for two reasons. First, historical data on the cross-sectional distribution of variables such as employment and wealth provide evidence on the potential shapes that such distributions can plausibly assume. For example, as unemployment rates above 50% have never been observed in U.S. data, it appears desirable to specify an aggregate employment distribution that satisfies this property. This can be achieved by augmenting the Markov processes for each group's employment status with this aggregate restriction.

Second, apart from these data-based restrictions, many histograms are theoretically very unlikely to occur in any given model economy. To illustrate this, consider the CDF of aggregate labor in Figure IIIb. The figure shows that an employment rate above 50% has close to zero

[9] In our environment in which agents control hazard rates of lumpy adjustments, inaction regions can be eliminated by specifying adjustment costs that are quadratic (as opposed to linear).

probability of occurring. To be precise, full employment happens with a 69.8% probability, one out of seven groups being unemployed happens with 25.7% probability, two out of seven happens with 4.1% probability, and three out of seven with a 0.4% probability. The joint probability of an even higher unemployment rate is less than 0.0002%. It therefore seems unlikely that the capital/employment histograms (as the ones illustrated in Figure I) that have more than three groups unemployed will materially affect the solution of the specified economy.

Based on these insights, we propose the following approach to gradually increase the number of groups and/or capital grid points:

Step 1: Solve the model with a computationally feasible number of grid points and groups.

Step 2: Evaluate the stationary distributions of the key aggregate variables (such as capital and labor), and evaluate which states have a small probability of occurring.

Step 3: Re-solve the model setting the probabilities of those states to zero (i.e., restricting the number of histograms), thereby lowering the number of states.

Step 4: Compare the solutions obtained from Step 3 to those from Step 1.

Step 5: If the solutions are numerically close, increase the number of groups and/or grid points while restricting the number of states as considered in Step 3.

Step 6: Repeat steps 3 through 5.

5. Conclusion

In this paper we have introduced a novel tractable method to introduce heterogeneity in macroeconomic models. The only computational operation needed to solve this class of models is the inversion of large sparse matrices. The main benefit relative to previous modeling environments is that no approximation is required even for environments with more than two agents (or groups of agents). Even though the class of models we propose has many applications in finance and economics, we illustrate its appeal by focusing on a generalized version of the well-known Krusell and Smith (1998) type setting. While our results from this analysis are very encouraging, in future work, we intend to apply this method to economies with material non-linearities, which are particularly challenging for existing solution methods.

References

Aiyagari, S. Rao, 1994, Uninsured Idiosyncratic Risk and Aggregate Saving, *Quarterly Journal of Economics* 109, 659–684.

Banerjee, Abhijit V., and Andrew F. Newman, 1993, Occupational Choice and the Process of Development, *Journal of Political Economy* 101, 27498.

Bewley, Truman, 1977, The Permanent Income Hypothesis: A Theoretical Formulation, *Journal of Econic Theory* 16, 252–292.

Binsbergen, Jules H. van, and Christian C. Opp, Forthcoming, Real Anomalies, *Journal of Finance* .

Cochrane, John, 1989, The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives, *American Economic Review* 79, 319–337.

den Haan, Wouter, 1996, Heterogeneity, aggregate uncertainty, and the short-term interest rate, *Journal of Business and Econ. Statistics* 14, 399411.

Den Haan, Wouter, 2010a, Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents, *Journal of Economic Dynamics and Control* 34, 79–99.

Den Haan, Wouter J., 2010b, Comparison of solutions to the incomplete markets model with aggregate uncertainty, *Journal of Economic Dynamics and Control* 34, 4–27.

Den Haan, Wouter J., Kenneth L. Judd, and Michel Juillard, 2010, Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty, *Journal of Economic Dynamics and Control* 34, 1–3.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, 2017, The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium, *Journal of Political Economy* 125, 140–223.

Gabaix, Xavier, 2011, The granular origins of aggregate fluctations, *Econometrica* 79, 733772.

Gomes, Francisco, and Alexander Michaelides, 2008, Asset Pricing with Limited Risk Sharing and Heterogeneous Agents, *Review of Financial Studies* 21, 415–448.

Han, Jiequn, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, 2018, Income and wealth distribution in macroeconomics: A continuous-time approach Working Paper, Princeton University.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, 2014, Consumption and Labor Supply with Partial Insurance: An Analytical Framework, *American Economic Review* 104, 2075–2126.

Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–87.

Judd, K., and T. Mertens, 2019, Solving an incomplete markets model with a large cross-section of agents, *Journal of Economic Dynamics and Control* forthcoming.

Judd, Kenneth L., Lilia Maliar, and Serguei Maliar, 2010, A Cluster-Grid Projection Method: Solving Problems with High Dimensionality, NBER Working Papers 15965, National Bureau of Economic Research, Inc.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit Cycles, *Journal of Political Economy* 105, 211–48.

Krusell, Per, and Anthony A. Smith, 1996, Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis, *Journal of Econonmic Dynamics and Control* 20, 527558.

Krusell, Per, and Anthony A. Smith, 1997, Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns, *Macroeconomic Dynamics* 1, 387422.

Krusell, Per, and Anthony A. Smith, 1998, Income and Wealth Heterogeneity in the Macroeconomy, *Journal of Political Economy* 106, 867–896.

Lucas, Deborah J., 1994, Asset pricing with undiversifiable income risk and short sales constraints: Deepening the equity premium puzzle, *Journal of Monetary Economics* 1994, 325–341.

Lucas, Robert E., 1987, *Models of business cycles* (New York: Basil Blackwell).

Marcet, Albert, and Kenneth J. Singleton, 1999, Equilibrium asset prices and savings of heterogeneous agents in the presence of incomplete markets and portfolio constraints, *Macroeconomic Dynamics* 3, 243277.

Scheinkman, Jose, and Laurence Weiss, 1986, Borrowing Constraints and Aggregate Economic Activity, *Econometrica* 54, 23–45.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, 2007, Asset Pricing with Idiosyncratic Risk and Overlapping Generations, *Review of Economic Dynamics* 10, 519–548.

Telmer, Chris I., 1993, Asset-pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.