

Merging Simulation and Projection Approaches to Solve High-Dimensional Problems with an Application to a New Keynesian model

Lilia Maliar and Serguei Maliar

March 9, 2014

Abstract

We introduce a numerical algorithm for solving dynamic economic models that merges stochastic simulation and projection approaches: we use simulation to approximate the ergodic measure of the solution, we cover the support of the constructed ergodic measure with a fixed grid, and we use projection techniques to accurately solve the model on that grid. The construction of the grid is the key novel piece of our analysis: we replace a large cloud of simulated points with a small set of "representative" points. We present three alternative techniques for constructing representative points: a clustering method, an epsilon-distinguishable set method, and a locally-adaptive variant of the epsilon-distinguishable set method. As an illustration, we solve one- and multi-agent neoclassical growth models and a large-scale new Keynesian model with a zero lower bound on nominal interest rates. The proposed solution algorithm is tractable in problems with high dimensionality (hundreds of state variables) on a desktop computer.

JEL classification : C61, C63, C68, E31, E52

Key Words : ergodic set; ε -distinguishable set; clusters; adaptive grid; discrepancy; large-scale model; new Keynesian model; ZLB; stochastic simulation

Ken Judd was also an author on this paper. See page 43

1 Introduction

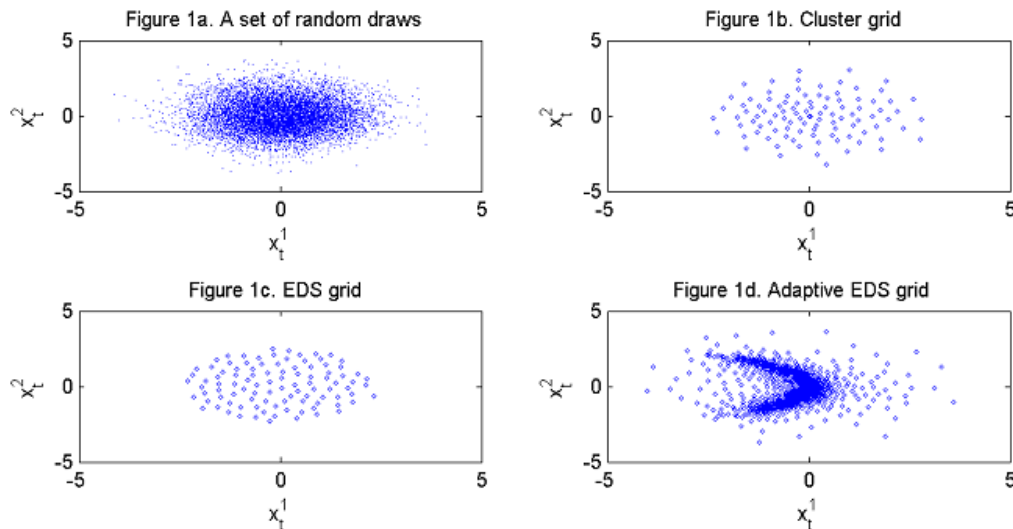
We introduce a numerical algorithm for solving dynamic economic models that merges stochastic simulation and projection approaches: we use simulation to approximate the ergodic measure of the solution, we cover the support of the constructed ergodic measure with a fixed grid, and we use projection techniques to accurately solve the model on that grid. The construction of the grid is the key novel piece of our analysis: we replace a large cloud of simulated points with a small set of "representative" points. We present three alternative techniques for constructing representative points: a clustering method, an epsilon-distinguishable set method, and a locally-adaptive variant of the epsilon-distinguishable set method. As an illustration, we solve one- and multi-agent neoclassical growth models and a large-scale new Keynesian model with a zero lower bound on nominal interest rates. The proposed solution algorithm is tractable in problems with high dimensionality (hundreds of state variables) on a desktop computer.

One broad class of numerical methods for solving dynamic economic models builds on stochastic simulation. First, this class includes iterative methods for solving rational expectations models; see, e.g., Marcet (1988), Smith (1993), Maliar and Maliar (2005) and Judd et al. (2011a). Second, it includes learning-based analysis; see, e.g., Marcet and Sargent (1989), Bertsekas and Tsitsiklis (1996), Pakes and McGuire (2001) and Powell (2011). Finally, it includes methods that use simulation to reduce information sets of decision makers; see, e.g., Krusell and Smith (1998), and Benkard et al. (2008). The key advantage of stochastic simulation methods is that the geometry of the set on which the solution is computed is adaptive. Namely, these methods solve dynamic economic models on a set of points produced by stochastic simulation, avoiding thus the cost of finding solutions in areas of the state space that are effectively never visited in equilibrium; see Judd et al. (2011a) for a discussion. However, a set of simulated points itself is not an efficient choice either as a grid for approximating a solution (it contains many closely-located and hence, redundant points) or as a set of nodes for approximating expectation functions (the accuracy of Monte Carlo integration is low).

Another broad class of numerical methods for solving dynamic economic models relies on projection techniques; see, e.g., Wright and Williams (1984), Judd (1992), Christiano and Fisher (2000), Krueger and Kubler (2004), Aruoba et al. (2006), Anderson et al. (2010), Malin et al. (2011), Pichler (2011) and Judd et al. (2013). Projection methods use efficient discretizations of the state space and effective deterministic integration methods, and they deliver very accurate solutions. However, a conventional projection method is limited to a fixed geometry such as a multidimensional hypercube. In order to capture all points that are visited in equilibrium, a hypercube must typically include large areas of the state space that have a low probability of happening in equilibrium. Moreover, the fraction of the irrelevant areas of a hypercube domain grows rapidly with the dimensionality of the problem.

The solution method introduced in this paper combines best features of stochastic simulation and projection methods, namely, it combines the adaptive geometry of stochastic

simulation methods with efficient discretization techniques of projection methods. As an example, in Figure 1a, we show a set of points that is obtained by simulating two state variables x_t^1 and x_t^2 of a typical dynamic economic model; this set of points identifies a high-probability area of the state space. In Figures 1b, 1c and 1d, we distinguish three different subsets of the simulated points – we call them "grids".



The grid shown in Figure 1b is constructed using methods from clustering analysis: we partition the simulated data into clusters, and we compute the centers of the clusters. The resulting *cluster grid* mimics the density function by placing more points in regions where the cloud of simulated points is more dense and fewer points where it is less dense. The grid shown in Figure 1c is produced by an ε -*distinguishable set* (EDS) technique: we select a subset of points that are situated at the distance at least ε from one another, where $\varepsilon > 0$ is a parameter. The EDS grid is roughly uniform. Finally, in Figure 1d, we show an example of a locally-adaptive EDS grid: instead of using a constant ε , we allow it to vary across the domain, i.e., for each point (x_t^1, x_t^2) , we have a different ε , i.e., it is a function $\varepsilon = \varepsilon(x_t^1, x_t^2)$ (in this specific example, we decrease the value of ε as we approach a hyperbola $x_t^1 = [x_t^2]^2$). This kind of grid construction enables us to control the density of grid points (and hence, the quality of approximation) over the solution domain.

An important question is: Which of these techniques delivers the best grid of points to be used within a projection method? Crude simulation shown in Figure 1a is not an efficient choice: having many closely-located grid points does not increase accuracy but cost. Cluster grids tend to produce a good fit in a high-probability area of the state space, but may result in larger errors in low-probability areas of the state space. EDS grids with a constant ε tend to deliver more uniform accuracy. Finally, locally-adaptive EDS grids allow us to automate the control of accuracy over the solution domain using the following two-step procedure: (i) compute a solution using an EDS grid with a constant ε and evaluate the quality of approximation; (ii) define ε to be a decreasing function of the size

of approximation errors, construct an EDS grid with a space-dependent ε and recompute the solution; iterate on steps (i) and (ii) if necessary. Thus, in those areas in which errors are large, we use a smaller ε ; this increases the density of grid points and hence, augments the accuracy.

An important role in our analysis plays the choice of an interpolation method, i.e. the way in which we approximate functions off the EDS grid. We consider two kinds of interpolant. One is a global polynomial function that approximates a given decision function on the whole solution domain. The other is a combination of local polynomial bases each of which approximates a decision function only in a small neighborhood of a given EDS grid point; a global approximation is then obtained by tying up local approximations together. There are many ways to construct local approximations and to tie them up into a global approximation. Our baseline technique is as follows: For each grid point x in the EDS grid, we construct a hypercube centered at that specific point, we populate this hypercube with low-discrepancy points (namely, Sobol points), and we solve the model using those points as a grid; see Niederreiter (1992) for a review of low-discrepancy methods. Here, we compute a solution to the model as many times as the number of points in the EDS grid (since we construct a Sobol grid around each EDS grid point). Finally, to simulate a solution, we use a nearest-neighbor approach.

We next incorporate the EDS and cluster-grid techniques into a numerical method for solving dynamic economic models. Our solution method requires some initial guess about the true solution to the model at the initialization stage, such as a log-linearized solution. In particular, we need an initial guess to produce simulated points which we can use for constructing a grid. We therefore proceed iteratively: guess a solution, simulate the model, construct a grid, solve the model on that grid using a projection method, and perform few iterations on these steps until the grid converges. We complement the efficient grid construction with other computational techniques suitable for high-dimensional problems, namely, low-cost monomial integration rules and a fixed-point iteration method for finding parameters of equilibrium rules.¹ Taken together, these techniques make our solution algorithm tractable in problems with high dimensionality – hundreds of state variables!

We first apply the EDS method to the standard neoclassical growth models with one and multiple agents (countries). The EDS method delivers accuracy levels comparable to the best accuracy attained in the related literature. In particular, we are able to compute global quadratic solutions for equilibrium problems with up to 80 state variables on a desktop computer using a serial MATLAB software (the running time ranges from 30 seconds to 24 hours). The maximum unit-free approximation error on a stochastic simulation is always smaller than 0.01%.

Our second and a more novel application is a new Keynesian model which includes a Taylor rule with a zero lower bound (ZLB) on nominal interest rates. This model has eight state variables and is characterized by a kink in equilibrium rules due to the ZLB.

¹In the present paper, we focus on equilibrium problems in which solutions are characterized by Euler equations, however, in a working paper version of the paper, Judd et al. (2012), we also show applications of the EDS technique to dynamic programming.

We focus on equilibrium in which target inflation coincides with actual inflation in the steady state. We parameterize the model using the estimates of Smets and Wouters (2003, 2007), and Del Negro et al. (2007). The EDS method is tractable for global polynomial approximations of degrees 2 and 3 (at least): the running time is less than 25 minutes in all the cases considered. For comparison, we also assess the performance of perturbation solutions of orders 1 and 2. We find that if the volatility of shocks is low and if we allow for negative nominal interest rates, both the EDS and perturbation methods deliver sufficiently accurate solutions. However, if either the ZLB is imposed or the volatility of shocks increases, the perturbation method is significantly less accurate than the EDS method. In particular, under some empirically relevant parameterizations, the perturbation methods of orders 1 and 2 produce errors that are as large as 25% and 38% on a stochastic simulation, while the corresponding errors for the EDS method are less than 5%. The difference between the EDS and perturbation solutions is economically significant. Namely, when the ZLB is active, the perturbation method considerably understates the duration of the ZLB episodes and the magnitude of the crises. We also solve the new Keynesian model with an active ZLB using a locally-adaptive EDS method, and we find that consecutive refinements of the EDS grid can considerably increase the quality of approximation.

The mainstream of the literature on new Keynesian models relies on local perturbation solution methods.² However, recent developments in numerical analysis triggered a quickly growing body of literature that computes non-linear solutions to medium- and large-scale new Keynesian models; see Judd et al. (2011b), Braun et al. (2012), Coibion et al. (2012), Fernández-Villaverde et al. (2012), Gust et al. (2012), Schmitt-Grohé and Uribe (2012), Aruoba and Schorfheide (2013), Gavion et al. (2013), Mertens and Ravn (2013), Richter and Throckmorton (2013). As is argued in Judd et al. (2011b), Fernández-Villaverde et al. (2012) and Braun et al. (2012), perturbation methods, which were traditionally used in this literature, do not provide accurate approximation in the context of new Keynesian models with the ZLB. Moreover, recent papers of Schmitt-Grohé and Uribe (2012), Mertens and Ravn (2013) and Aruoba and Schorfheide (2013) argue that new Keynesian economies may have multiple equilibria in the presence of ZLB. In particular, the last paper accurately computes a deflation and sunspot equilibria with a full set of stochastic shocks using a modified variant of a *cluster-grid algorithm (CGA)* introduced in Judd et al. (2010, 2011b). Namely, to increase the accuracy of solutions in the ZLB area, first, they add grid points near the ZLB area using the actual data on the U.S. economy; and second, they apply two piecewise bases to separately approximate the solution in the areas with active and non-active ZLB.

Our locally-adaptive EDS methods are related to several other methods in the literature. First, the EDS technique with local bases has similarity to finite element approxima-

²Most papers use linear approximations, however, there are also papers that compute quadratic approximations (e.g., Kollmann, 2002, and Schmitt-Grohé and Uribe, 2007) and cubic approximations (e.g., Rudebusch and Swanson, 2008). Earlier applications of nonlinear solution methods either focus on low-dimensional problems or employ simplifying assumptions; see Adam and Billi (2006), Anderson et al. (2010), and Adjemian and Juillard (2011).

tion methods that construct a global approximation using a combination of disjoint local approximations; see Hughes (1987) for a mathematic review of finite element methods, and see McGrattan (1996) for their applications to economics. Second, a locally-adaptive EDS technique with space-dependent ε resembles locally-adaptive sparse-grid techniques which refines an approximation by introducing new grid points and bases functions in those areas in which the quality of approximation is low; see Ma and Zabararas (2009) for a review of such methods, and see Brumm and Scheidegger (2013) for their applications to economic problems. Finally, a locally-adaptive EDS technique is also related to the analysis of Aruoba and Schorfheide (2013) who show the benefits of adaptive grid points and bases functions in the context of a new Keynesian model with the ZLB.

The CGA and EDS projection methods can be used to accurately solve small-scale models that were previously studied using other global solution methods.³ However, a comparative advantage of these algorithms is their ability to solve large-scale problems that other methods find intractable or expensive. The speed of the CGA and EDS algorithms also makes them potentially useful in estimation methods that solve economic models at many parameters vectors; see Fernández-Villaverde and Rubio-Ramírez (2007), and Winschel and Krätzig (2010). Finally, cluster grids and EDS grids can be used in other applications that require us to produce a discrete approximation to the ergodic distribution of a stochastic process with a continuous density function, in line with Tauchen and Hussey (1991).

The rest of the paper is as follows: In Section 2, we describe the construction of the simulation-based grids using EDS, locally-adaptive EDS and clustering techniques. In Section 3, we integrate the EDS grid into a projection method for solving dynamic economic models. In Section 4, we apply the EDS algorithm to solve one- and multi-agent neoclassical growth models. In Section 5, we compute a solution to a new Keynesian model with the ZLB. In Section 6, we conclude.

2 Discrete approximations to the ergodic set

In this section, we introduce techniques that produce a discrete approximation to the ergodic set of a stochastic process with a continuous density function. Later, we will use the resulting discrete approximation as a grid for finding a solution in the context of a projection-style numerical method for solving dynamic economic models.

³For reviews of methods for solving dynamic economic models, see Taylor and Uhlig (1990), Gaspar and Judd (1997), Judd (1998), Marimon and Scott (1999), Santos (1999), Christiano and Fisher (2000), Adda and Cooper (2003), Aruoba et al. (2006), Den Haan (2010), Kollmann et al. (2011), and Maliar and Maliar (2013).

2.1 A class of stochastic processes

We focus on a class of discrete-time stochastic processes that can be represented in the form

$$x_{t+1} = \varphi(x_t, \epsilon_{t+1}), \quad t = 0, 1, \dots, \quad (1)$$

where $\epsilon \in E \subseteq \mathbb{R}^p$ is a vector of p independent and identically distributed shocks, and $x \in X \subseteq \mathbb{R}^d$ is a vector of d (exogenous and endogenous) state variables. The distribution of shocks is given by a probability measure Q defined on a measurable space (E, \mathbb{E}) , and x is endowed with its relative Borel σ -algebra denoted by \mathbb{X} .

Many dynamic economic models have equilibrium laws of motion for state variables that can be represented by a stochastic system in the form (1). For example, the standard neoclassical growth model, described in Section 4, has the laws of motion for capital and productivity that are given by $k_{t+1} = K(k_t, a_t)$ and $a_{t+1} = a_t^\rho \exp(\epsilon_{t+1})$, respectively, where $\epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2)$, $\sigma > 0$ and $\rho \in (-1, 1)$; by setting $x_t \equiv (k_t, a_t)$, we arrive at (1).

To characterize the dynamics of (1), we use the following definitions.

Def 1. A transition probability is a function $\mathcal{P} : X \times X \rightarrow [0, 1]$ that has two properties: (i) for each measurable set $\mathcal{A} \in \mathbb{X}$, $\mathcal{P}(\cdot, \mathcal{A})$ is \mathbb{X} -measurable function; and (ii) for each point $x \in X$, $\mathcal{P}(x, \cdot)$ is a probability measure on (X, \mathbb{X}) .

Def 2. An (adjoint) Markov operator is a mapping $\mathcal{M}^* : X \rightarrow X$ such that $\mu_{t+1}(\mathcal{A}) = (\mathcal{M}^* \mu_t)(\mathcal{A}) \equiv \int \mathcal{P}(x, \mathcal{A}) \mu_t(dx)$.

Def 3. An invariant probability measure μ is a fixed point of the Markov operator \mathcal{M}^* satisfying $\mu = \mathcal{M}^* \mu$.

Def 4. A set \mathcal{A} is called invariant if $\mathcal{P}(x, \mathcal{A}) = 1$ for all $x \in \mathcal{A}$. An invariant set \mathcal{A}^* is called ergodic if it has no proper invariant subset $\mathcal{A} \subset \mathcal{A}^*$.

Def 5. An invariant measure μ is called ergodic if either $\mu(\mathcal{A}) = 0$ or $\mu(\mathcal{A}) = 1$ for every invariant set \mathcal{A} .

These definitions are standard to the literature on dynamic economic models; see Stokey, Lucas and Prescott (1989), and Stachursky (2009). $\mathcal{P}(x, \mathcal{A})$ is the probability that stochastic system (1) whose today's state is $x_t = x$ will move tomorrow to a state $x_{t+1} \in \mathcal{A}$. The Markov operator \mathcal{M}^* maps today's probability into tomorrow's probability, namely, if $\mu_t(\mathcal{A})$ is the probability that the system (1) is in \mathcal{A} at t , then $(\mathcal{M}^* \mu_t)(\mathcal{A})$ is the probability that the system will remain in the same set at $t+1$. Applying the operator \mathcal{M}^* iteratively, we can describe the evolution of the probability starting from a given $\mu_0 \in \mathbb{X}$. An invariant probability measure μ is a steady state solution of the stochastic system (1). An invariant set \mathcal{A} is the one that keeps the system (1) forever in \mathcal{A} , and an ergodic set \mathcal{A}^* is an invariant set of the smallest possible size. Finally, an invariant probability measure is ergodic if all the probability is concentrated in just one of the invariant sets.

The dynamics of (1) produced by economic models can be very complex. In particular, the Markov process (1) may have no invariant measure or may have multiple invariant measures. These cases represent challenges to numerical methods that approximate solutions to dynamic economic models. However, there is another challenge that numerical methods face – the curse of dimensionality. The most regular problem with a unique,

smooth and well-behaved solution can become intractable when the dimensionality of the state space gets large. The challenge of high dimensionality is the focus of our analysis. We employ the simplest possible set of assumptions that allows us to describe and to test computational techniques that are tractable in high-dimensional applications.

Assumption 1. *There exists a unique ergodic set \mathcal{A}^* and the associated ergodic measure μ .*

Assumption 2. *The ergodic measure μ admits a representation in the form of a density function $g : X \rightarrow \mathbb{R}^+$ such that $\int_{\mathcal{A}} g(x) dx = \mu(\mathcal{A})$ for every $\mathcal{A} \subseteq \mathbb{X}$.*

2.2 An EDS technique for approximating the ergodic set

We propose a two-step procedure for forming a discrete approximation to the ergodic set. First, we identify an area of the state space that contains nearly all the probability mass. Second, we cover this area with a finite set of points that are roughly evenly spaced.

2.2.1 An essentially ergodic set

We define a high-probability area of the state space using the level set of the density function.

Def 6. *A set $\mathcal{A}^\eta \subseteq \mathcal{A}^*$ is called a η -level ergodic set if $\eta > 0$ and*

$$\mathcal{A}^\eta \equiv \{x \in X : g(x) \geq \eta\}.$$

The mass of \mathcal{A}^η under the density $g(x)$ is equal to $p(\eta) \equiv \int_{g(x) \geq \eta} g(x) dx$. If $p(\eta) \approx 1$, then \mathcal{A}^η contains all X except for points where the density is lowest, in which case \mathcal{A}^η is called an *essentially ergodic set*.

By construction, the correspondence $\mathcal{A}^\eta : \mathbb{R}^+ \rightrightarrows \mathbb{R}^d$ maps η to a compact set. The correspondence \mathcal{A}^η is upper semi-continuous but may be not lower semi-continuous (e.g., if x is drawn from a uniform distribution $[0, 1]$). Furthermore, if g is multimodal, then for some values of η , \mathcal{A}^η may be disconnected (composed of disjoint areas). Finally, for $\eta > \max_x \{g(x)\}$, the set \mathcal{A}^η is empty.

Our approximation to the essentially ergodic set builds on stochastic simulation. Formally, let P be a set of n independent random draws $x_1, \dots, x_n \subseteq \mathbb{R}^d$ generated with the distribution function $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^+$. For a given subset $J \subseteq \mathbb{R}^d$, we define $C(P; J)$ as a characteristic function that counts the number of points from P in J . Let \mathcal{J} be a family generated by the intersection of all subintervals of \mathbb{R}^d of the form $\prod_{i=1}^d [-\infty, v_i)$, where $v_i > 0$.

Assumption 3. *The empirical distribution function $\hat{\mu}(J) \equiv \frac{C(P; J)}{n}$ converges to the true distribution function $\mu(J)$ for every $J \in \mathcal{J}$ when $n \rightarrow \infty$.*

If random draws are independent, the asymptotic rate of convergence of $\hat{\mu}$ to μ is given by the so-called law of iterated logarithm of Kiefer's (1961), namely, it is $(\log \log n)^{1/2} (2n)^{-1/2}$. For serially correlated processes like (1), the convergence rate depends on specific assumptions; see Zhao and Woodroffe (2008) for the results on general stationary processes.

We use the following algorithm to select a subset of simulated points that belongs to an essentially ergodic set \mathcal{A}^η .

(Algorithm \mathcal{A}^η): Selection of points within an essentially ergodic set.

Step 1. Simulate (1) for T periods.

Step 2. Select each κ th point to get a set P of n points $x_1, \dots, x_n \in X \subseteq \mathbb{R}^d$.

Step 3. Estimate the density function $\hat{g}(x_i) \approx g(x_i)$ for all $x_i \in P$.

Step 4. Remove all points for which the density is below η .

In Step 2, we include in the set P only each κ th observation to make random draws (approximately) independent. As far as Step 3 is concerned, there are various methods in statistics that can be used to estimate the density function from a given set of data; see Scott and Sain (2005) for a review. We use one of such methods, namely, a multivariate kernel algorithm with a normal kernel which estimates the density function in a point x as

$$\hat{g}(x) = \frac{1}{n (2\pi)^{d/2} \bar{h}^d} \sum_{i=1}^n \exp \left[-\frac{D(x, x_i)}{2\bar{h}^2} \right], \quad (2)$$

where \bar{h} is the bandwidth parameter, and $D(x, x_i)$ is the distance between x and x_i . The complexity of Algorithm \mathcal{A}^η is $O(n^2)$ because it requires to compute pairwise distances between all the sample points. Finally, in Step 3, we do not choose the density cutoff η but a fraction of the sample to be removed, δ , which is related to η by $p(\eta) = \int_{g(x) \geq \eta} g(x) dx = 1 - \delta$. For example, $\delta = 0.05$ means that we remove 5% of the sample which has the lowest density.

2.2.2 An ε -distinguishable set (EDS)

Our next objective is to construct a uniformly-spaced set of points that covers the essentially ergodic set (to have a uniformly-spaced grid for a projection method). We proceed by selecting an ε -distinguishable subset of simulated points in which all points are situated at least on the distance ε from one another. Simulated points are not uniformly-spaced but the EDS subset will be roughly uniform, as we will show in Appendix A3.

Def 7. Let (X, D) be a bounded metric space. A set P^ε consisting of points $x_1^\varepsilon, \dots, x_M^\varepsilon \in X \subseteq \mathbb{R}^d$ is called ε -distinguishable if $D(x_i^\varepsilon, x_j^\varepsilon) > \varepsilon$ for all $1 \leq i, j \leq M : i \neq j$, where $\varepsilon > 0$ is a parameter.

EDSs are used in mathematical literature that studies the entropy; see Temlyakov (2011) for a review. This literature focuses on a problem of constructing an EDS that covers a given subset of \mathbb{R}^d (such as a multidimensional hypercube). We study a different problem, namely, we construct an EDS for a given discrete set of points. To this purpose, we introduce the following algorithm.

(Algorithm P^ε): Construction of an EDS.

Let P be a set of n points $x_1, \dots, x_n \in X \subseteq \mathbb{R}^d$.

Let P^ε begin as an empty set, $P^\varepsilon = \{\emptyset\}$.

Step 1. Select $x_i \in P$. Compute $D(x_i, x_j)$ to all x_j in P .

Step 2. Eliminate from P all x_j for which $D(x_i, x_j) < \varepsilon$.

Step 3. Add x_i to P^ε and eliminate it from P .

Iterate on Steps 1-3 until all points are eliminated from P .

The complexity of Algorithm P^ε is of order $O(nM)$, where M is the number of points into the set P^ε . Indeed, consider the worst-case scenario such that ε is smaller than all inter-point distances for the first M points. Then, the algorithm will go through $n - M$ iterations without eliminating any point, and it will eliminate $n - M$ points at the end. Under this scenario, the complexity is $(n - 1) + (n - 2) \dots + (n - M) = \sum_{i=1}^M (n - i) = nM - \frac{M(M+1)}{2} \leq nM$. When no points are eliminated from P , i.e., $M = n$, the complexity is quadratic, $O(n^2)$. However, the number of points M in an EDS is bounded from above if X is bounded; see Proposition 2 in Appendix A2. This means that asymptotically, when $n \rightarrow \infty$, the complexity of Algorithm P^ε is linear, $O(n)$.

2.2.3 Distance between points

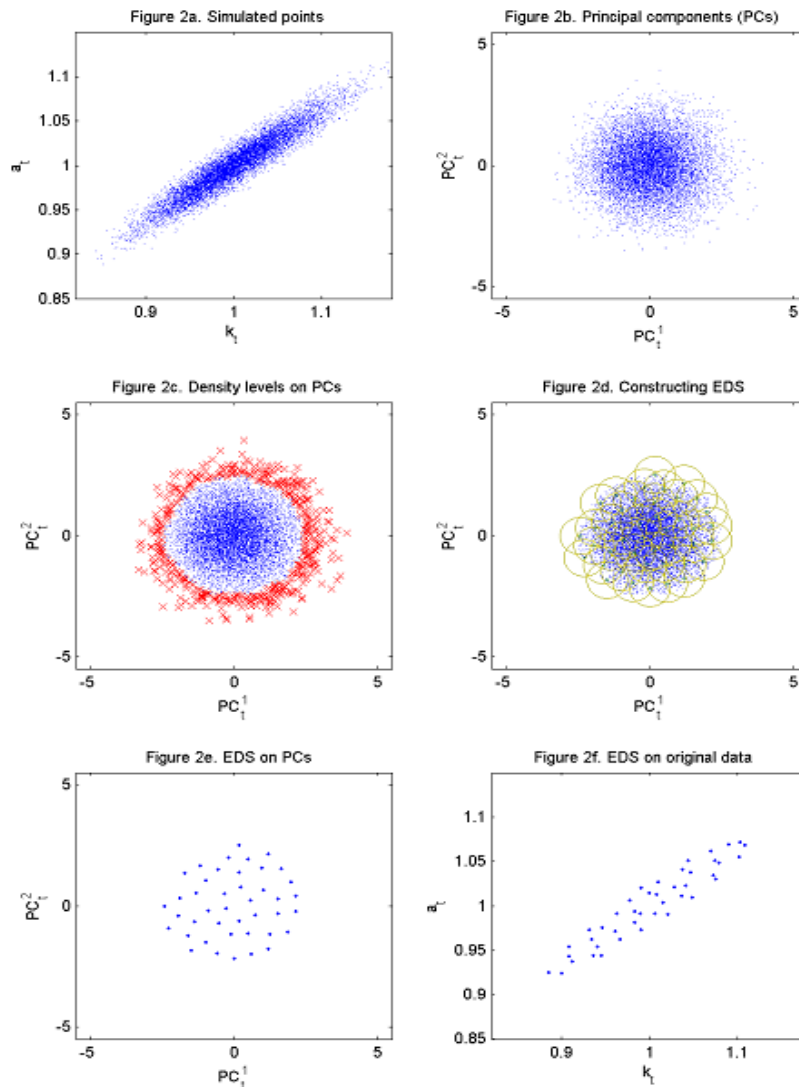
Both estimating the density function and constructing an EDS requires us to measure the distance between simulated points. Generally, variables in economic models have different measurement units and are correlated. This affects the distances between the simulated points and hence, affects the resulting EDS. Therefore, prior to using Algorithm \mathcal{A}^η and Algorithm P^ε , we normalize and orthogonalize the simulated data.

To be specific, let $X \in \mathbb{R}^{n \times d}$ be a set of simulated data normalized to zero mean and unit variance. Let $x_i \equiv (x_i^1, \dots, x_i^d)$ be an observation $i = 1$ (there are n observations), and let $x^\ell \equiv (x_1^\ell, \dots, x_n^\ell)^\top$ be a variable ℓ (there are d variables), i.e., $X = (x^1, \dots, x^d) = (x_1, \dots, x_n)^\top$. We first compute the singular value decomposition of X , i.e., $X = UQV^\top$, where $U \in \mathbb{R}^{n \times d}$ and $V \in \mathbb{R}^{d \times d}$ are orthogonal matrices, and $Q \in \mathbb{R}^{d \times d}$ is a diagonal matrix. We then perform a linear transformation of X using $PC \equiv XV$. The variables $PC = (PC^1, \dots, PC^d) \in \mathbb{R}^{n \times d}$ are called *principal components* (PCs) of X , and are orthogonal (uncorrelated), i.e., $(PC^{\ell'})^\top PC^\ell = 0$ for any $\ell' \neq \ell$. As a measure of distance between two observations x_i and x_j , we use the Euclidean distance between their PCs, namely, $D(x_i, x_j) = \left[\sum_{\ell=1}^d (PC_i^\ell - PC_j^\ell)^2 \right]^{1/2}$, where all principal components PC^1, \dots, PC^d are normalized to unit variance.

2.2.4 An illustration of the EDS technique

In this section, we will illustrate the EDS technique described above by way of example. We consider the standard neoclassical stochastic growth model with a closed-form solution (see Section 4 for a description of this model). We simulate time series for capital

and productivity level of length 1,000,000 periods, and we select a sample of 10,000 observations by taking each 100th point (to make the draws independent); see Figure 2a.



We orthogonalize the data using the principal component (PC) transformation, and we normalize the PCs to unit variance; see Figure 2b. We estimate the density function using the multivariate kernel algorithm with the standard bandwidth of $\bar{h} = n^{-1/(d+4)}$, and we remove from the sample 5% of simulated points in which the density is lowest; see Figure 2c. We construct an EDS; see Figure 2d. We plot such a set in the PC and original coordinates in Figure 2e and Figure 2f, respectively. As we see, the EDS technique delivers a set of points that covers the same area as does the set of simulated points but that is spaced roughly uniformly.⁴

⁴Our two-step procedure produces an approximation not only to the ergodic set but to the whole ergodic distribution (because in the first step, we estimate the density function in all simulated points

2.2.5 Dispersion and discrepancy of EDS grids

In our examples, the EDS grids constructed on simulated series appear to be uniform. However, an important question is whether our construction guarantees the uniformity of grid points in general. We address this question in Appendices A1 and A3, specifically, we provide formal results about the degrees of dispersion and discrepancy of EDS grids from a uniform distribution.

Our analysis is related to recent mathematical literature on covering-number problems (see, Temlyakov, 2011) and random sequential packing problems (see, Baryshnikov et al., 2008). A well-known example from this literature is a car-parking model of Rényi (1958). Cars of a length ε park at random locations along the roadside of a length one subject to a non-overlap with the previously parked cars. It is known that when cars arrive at uniform random positions, they are also distributed uniformly in the limit $\varepsilon \rightarrow 0$.⁵

Our analysis differs from Rényi's (1958) analysis in that cars can arrive at random positions with an arbitrary density function (normally, we do not know density functions of stochastic processes arising in an economic model that we try to solve). In terms of Rényi's (1958) problem, our results are as follows: We show that EDS grids are low-dispersion sequences for any density function, namely, any two points (cars) in the EDS grid are situated on the distance between ε and 2ε from each other, and this distance converges to 0 as $\varepsilon \rightarrow 0$ (see Proposition 1 in Appendix A1). However, we find that this fact alone is not sufficient to guarantee the asymptotic uniformity (low discrepancy) of the EDS grids (see Proposition 3 in Appendix A3). To see the intuition, consider Rényi's (1958) setup such that on the interval $[0, \lambda^*]$, evil drivers park their cars on the distance 2ε to leave as little parking space to other drivers as possible, and on the interval $[\lambda^*, 1]$, a police officer directs the cars to park on the distance ε in a socially optimal way. Under this construction, there are twice as many points in the second subinterval as in the first one for any ε (and this non-uniformity is not reduced when $\varepsilon \rightarrow 0$). Finally, we establish that even though EDS grids do not possess the property of low-discrepancy sequences in general, their discrepancy from the uniform distribution is bounded from above for any density function; see Proposition 3 in Appendix A3.

2.2.6 Number of points in EDS grids

Under our baseline Algorithm P^ε , the cardinality of an EDS grid (i.e., the number of points in it) depends on the value of $\varepsilon > 0$: the smaller is ε , the more points are included in the EDS grid. We derive bounds on the number of points in the EDS grids in Proposition 2 of Appendix A2, however, the exact relation is hard to characterize analytically, in

including those that form an ε -distinguishable set). The density weights show what fraction of the sample each "representative" point represents, and can be used to construct weighted-average approximations. If our purpose is to construct a set of evenly-spaced points, we do not need to use the density weights and should treat all points equally.

⁵Rényi (1958) shows that they occupy about 75% of the roadside at jamming, namely, $\lim_{\varepsilon \rightarrow 0} E[M] \varepsilon \approx 0.748$.

particular, because the cardinality of the EDS grid depends on the order in which points are processed.

In applications, it may be necessary to control the number of grid points, for example, in a projection method, we need to construct a grid with a given number of grid points \overline{M} . To construct the relation between ε and $\overline{M} = \overline{M}(\varepsilon)$, we can use a simple numerical bisection method.

(Algorithm \overline{M}): Construction of an EDS with a target number of points \overline{M} .

For iteration $i = 1$, fix $\varepsilon_{\min}^{(1)}$ and $\varepsilon_{\max}^{(1)}$ such that $M(\varepsilon_{\max}^{(1)}) \leq \overline{M} \leq M(\varepsilon_{\min}^{(1)})$.

Step 1. On iteration i , take $\varepsilon = \frac{\varepsilon_{\min}^{(i)} + \varepsilon_{\max}^{(i)}}{2}$, construct an EDS and compute $M(\varepsilon)$.

Step 2. If $M(\varepsilon) > \overline{M}$, then set $\varepsilon_{\min}^{(i+1)} = \varepsilon$; and otherwise, set $\varepsilon_{\max}^{(i+1)} = \varepsilon$.

Iterate on Steps 1 and 2 until $M(\varepsilon)$ converges.

To find initial values of ε_{\max} and ε_{\min} , we use the bounds established in Appendix A2 (see Proposition 2), namely, we set $\varepsilon_{\max}^{(1)} = 0.5r_{\max}\overline{M}^{-1/d}$ and $\varepsilon_{\min}^{(1)} = r_{\min}\left(\overline{M}^{1/d} - 1\right)^{-1}$, where r_{\max} and r_{\min} are, respectively, the largest and smallest PCs of the simulated points. Since the essentially ergodic set is not necessarily a hypersphere (as is assumed in Proposition 2), we take r_{\max} and r_{\min} to be the radii of the limiting hyperspheres that contain none and all PCs of the simulated points, respectively.

2.3 Reducing the cost of constructing an EDS on the essentially ergodic set

The two-step procedure described in Section 2.2 has a complexity of $O(n^2)$. This is because the estimation of the density function in Step 3 of Algorithm \mathcal{A}^η has a complexity of order $O(n^2)$, and the construction of an EDS set in Step 1 of Algorithm P^ε has a complexity of order $O(nM)$. (The latter is significantly lower than the former if $M \ll n$). The complexity of order $O(n^2)$ does not imply a substantial cost for the size of applications we study in the present paper, however, it might be expensive for larger applications.

We now describe an alternative implementation of the two-step procedure that has a lower complexity, namely, $O(nM)$. The idea is to invert the steps in the two-step procedure described in Section 2.2, namely, we first construct an EDS with M points using all simulated points, and we then remove from the EDS a subset of points with the lowest density. Since we need to estimate the density function only in M simulated points, the complexity is reduced to $O(nM)$.

(Algorithm P^ε – cheap): Construction of an EDS.

Step 1. Simulate (1) for T periods.

Step 2. Select each κ th point to get a set P of n points $x_1, \dots, x_n \in X \subseteq \mathbb{R}^d$.

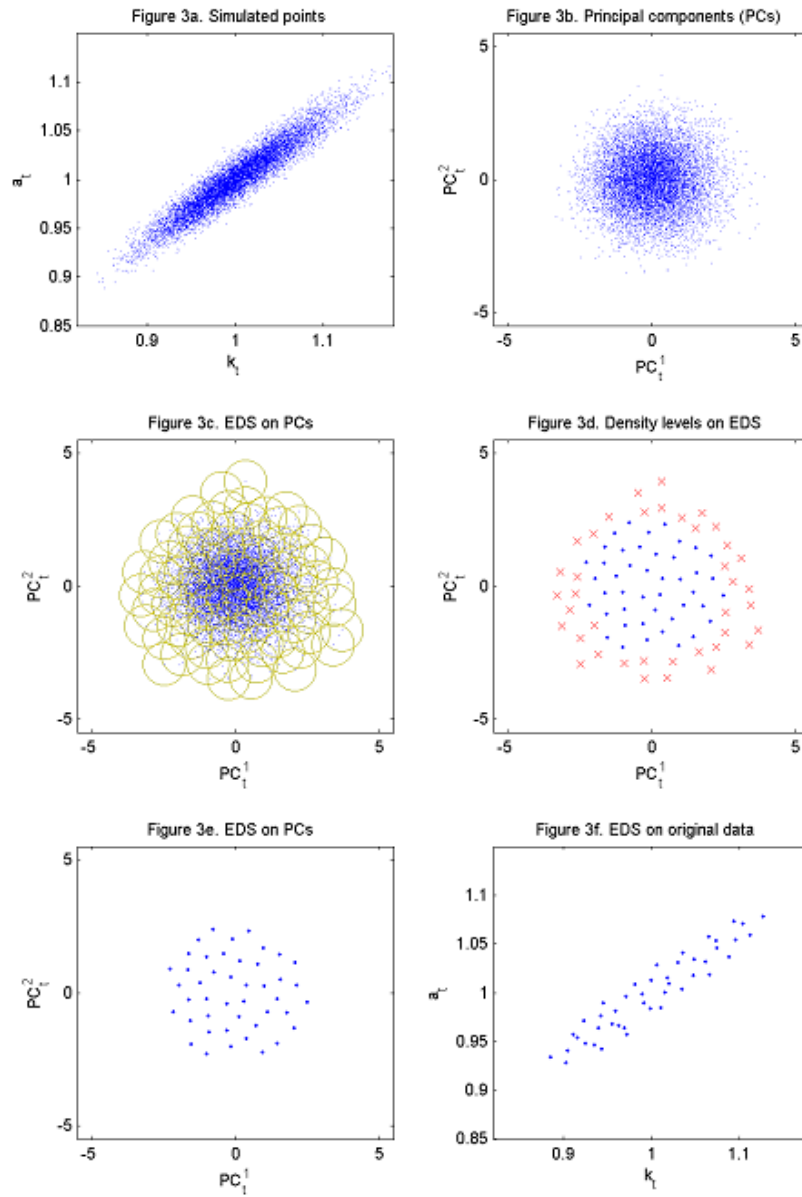
Step 3. Select an EDS P^ε of M points, $x_1^\varepsilon, \dots, x_M^\varepsilon$ using Algorithm P^ε .

Step 4. Estimate the density function $\hat{g}(x_i^\varepsilon) \approx g(x_i^\varepsilon)$ for all $x_i^\varepsilon \in P^\varepsilon$ using (2).

Step 5. Remove a fraction of points δ of P^ε which has the lowest density.

To control the fraction of the sample removed, we use the estimated density function \hat{g} . Note that when eliminating a point $x_j^\varepsilon \in P^\varepsilon$, we remove $\frac{\hat{g}(x_j^\varepsilon)}{\sum_{i=1}^M \hat{g}(x_i^\varepsilon)}$ of the original sample. We therefore can proceed with eliminations of points from the EDS one by one until their cumulative mass reaches the target value of δ .

To illustrate the application of the above procedure, we again use the example of the neoclassical stochastic growth model with a closed-form solution studied in Section 2.2.4; see Figures 3a-3f.



We first compute the normalized PCs of the original sample; see Figure 3b (this step in

the same as in Figure 2b). We compute an EDS P^ε on the normalized PCs; see Figure 3c. We then estimate the density function in all points of P^ε using the kernel density algorithm. We next remove from P^ε a set of points that has the lowest density function and that represents 5% of the sample. The removed points are represented with crosses in Figure 3d. The resulting EDS is shown in Figure 3e. Finally, we plot the EDS grid in the original coordinates in Figure 3f.

2.4 Cluster-grid technique

We have described one specific EDS procedure for forming a discrete approximation to the essentially ergodic set of the stochastic process (1). There are other procedures that can be used for this purpose. In particular, we can use methods from cluster analysis to select a set of representative points from a given set of simulated points; see Everitt et al. (2011) for a review of clustering techniques. Namely, we partition the simulated data into clusters (groups of closely-located points), and we replace each cluster with one representative point. In this paper, we study two clustering methods that can be used in the context of our analysis, an agglomerative hierarchical and K-means ones. The steps of an agglomerative hierarchical method are shown below, and a K-means method is described in Appendix B2.⁶

(Algorithm P^c): Agglomerative hierarchical clustering algorithm.

Initialization. Choose M , the number of clusters to be created.

In a zero-order partition $\mathcal{P}^{(0)}$, each simulated point represents a cluster.

Step 1. Compute all pairwise distances between the clusters in a partition $\mathcal{P}^{(i)}$.

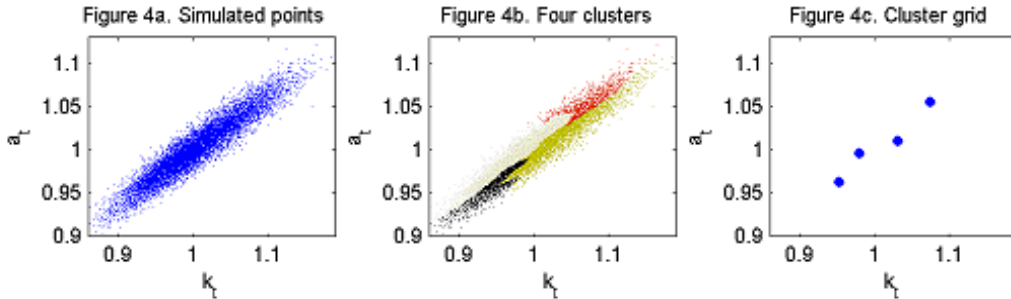
Step 2. Merge a pair of clusters with the smallest distance to obtain a partition $\mathcal{P}^{(i+1)}$.

Iterate on Steps 1 and 2. Stop when the number of clusters in the partition is M .

Represent each cluster with a simulated point which is closest to the cluster's center.

As a measure of distance between two clusters, we use Ward's measure of distance; see Appendix B1. In Figures 4a-4c, we show an example in which we partition a set of simulated points into 4 clusters and construct 4 representative points. A representative point is the closest point to the cluster center (computed as the average of all observations in the given cluster).

⁶The clustering methods were used to produce all the numerical results in the earlier versions of the paper, Judd et al. (2010, 2011b). In our examples, projection methods operating on cluster grids and those operating on EDSs deliver comparable accuracy of solutions.



The advantage of the clustering methods is that we can control the number of grid points directly (while the number of points in an EDS is controlled via ε). The drawbacks are that their complexity is higher, namely, it is of order $O(n^3)$ and $O(n^{dM+1} \log n)$ for the agglomerative hierarchical and K-means algorithms, respectively. Also, the properties of grids produced by clustering methods are hard to characterize analytically.

As is in the case of the EDS technique, two versions of the cluster grid technique can be constructed: we can first remove the low-density points and then construct representative points using clustering methods (this is parallel to the basic two-step EDS procedure of Section 2.2), or we can first construct clusters and then eliminate representative points in which the density function is the lowest (this is parallel to the cheap version of the two-step procedure described in Section 2.3). Prior to applying the clustering methods, we preprocess the data by constructing the normalized PCs, as we do when constructing an EDS grid in Section 2.2.3.

2.5 Locally-adaptive EDS grids

The locally-adaptive EDS grid technique makes it possible to control the quality of approximation over the state space. Namely, we place more grid points in those areas in which the accuracy is low.

In simple cases, we may know a priori that an approximation is less accurate in some area $X_1 \subseteq \mathbb{R}^d$ than in other areas. Consequently, we can use a small ε_1 in the area X_1 , and we can use a large ε_2 everywhere else. This produces more dense grid points in X_1 than in the rest of the domain; an example of this construction is shown in Figure 1d. However, in the typical case, it is not a priori known where the solution is accurate, and we proceed as follows:

(Algorithm P^ε – locally adaptive): Construction of a locally adaptive EDS.

Step 1. Define $\varepsilon_1, \dots, \varepsilon_n$ for a given set $x_1, \dots, x_n \in X \subseteq \mathbb{R}^d$ (initially, $\varepsilon_i = \varepsilon$ for all i).

Step 2. Construct an EDS P^ε by using ε_i for each $x_i \in X$ and approximate $\hat{f} \approx f$.

Step 3. Evaluate approximation errors $\mathcal{R}(x_i) = \left\| \hat{f}(x_i) - f(x_i) \right\|$ for all $x_i \in X$.

Step 4. Define $\mathcal{E}(\mathcal{R}(x_i))$ to be a decreasing function of approximation errors.

Step 5. Compute $\varepsilon_i = \mathcal{E}(\mathcal{R}(x_i))$ for all $x_i \in X$ and go to Step 2.

Under the above algorithm, the larger is the approximation error in a given data point x_i , the smaller is the corresponding value of $\varepsilon_i = \mathcal{E}(\mathcal{R}(x_i))$ and hence, the higher is the density of grid points. In certain sense, this construction is similar to locally-adaptive techniques in the sparse-grid literature in that it refines an approximation by introducing new grid points and bases functions in those areas in which the quality of approximation is low; see Ma and Zabaras (2009) for a review of this literature; and see Brumm and Scheidegger (2013) for examples of economic applications. The locally-adaptive EDS grid technique is especially useful in applications with kinks and strong non-linearities. In Section 5, we will study an example of such an application – a new Keynesian model with a ZLB on the nominal interest rate.

2.6 Approximating a function off the EDS grid

There is a variety of numerical techniques in mathematical literature that can be used to approximate functions off the grid. They typically require us to assume a flexible functional form $\hat{f}(x; b)$ characterized by a parameters vector b , and to find a parameters vector b that minimizes the approximation errors, $\epsilon(x_i^\varepsilon; b) \equiv \hat{f}(x_i^\varepsilon; b) - f(x_i^\varepsilon)$, on the constructed EDS grid $x_1^\varepsilon, \dots, x_m^\varepsilon \in P^\varepsilon$ according to some norm $\|\cdot\|$. If the constructed $\hat{f}(\cdot; b)$ coincides with f in all grid points, then we say that $\hat{f}(\cdot; b)$ interpolates f off the EDS grid (this requires that the number of grid points in the EDS grid is the same as the number of the parameters in b). Otherwise, we say that $\hat{f}(\cdot; b)$ approximates f on the EDS grid (this is similar to a regression analysis in econometrics when the number of data points is larger than the number of the regression coefficients).

Global polynomial basis functions. A convenient choice for an approximating function is a high-degree ordinary polynomial function. Such a function is easy to construct, and it can be fitted to the data using simple and reliable linear approximation methods; see Judd et al. (2011a). Orthogonal polynomial families are another useful choice even though the property of orthogonality is not satisfied for the simulation-based grid points; see Judd et al. (2011a) for a discussion. However, global polynomial approximations may be not sufficiently flexible to accurately approximate decision functions with strong non-linearities and kinks.

Piecewise local basis functions. Piecewise local bases are more flexible than global ones because each local polynomial basis function approximates a decision function just

in an a small neighborhood of a given EDS grid point. A global approximation is obtained by combing local approximations together. There are many ways to construct local approximations and to tie them up into a global approximation. Our baseline technique is as follows: For each grid point $x_i^{\varepsilon_i}$ in the EDS grid, we construct a hypercube centered at that specific point, cover the hypercube with a uniformly distributed set of points, and solve the model on this set of points. As a set of points that covers the hypercube uniformly, we use low-discrepancy sequences, namely, a Sobol sequence; an example of such a sequence is shown in Figure 5b; see Niederreiter (1992) for a review of low-discrepancy methods. Thus, we re-compute a solution to the model as many times as the number of points in the EDS grid. Under piecewise local polynomial approximations, we use low-degree polynomial bases, which helps us to keep the cost reasonably low. Finally, to simulate the solution, we rely on a nearest neighbor approach. Our construction of local bases has similarity to finite-element methods; see Hughes (1987) for a mathematic review of such methods; and see McGrattan (1996) for their applications to economics.

Piecewise local basis functions with locally-adaptive EDS grids. Piecewise local basis functions can be naturally combined with the locally-adaptive EDS grid technique. This combination enables us to refine the solution only in those areas in which the accuracy is not sufficient and to hold fixed the solution in the remaining points. That is, when we add new points to the EDS grid, we need to compute the solutions in the neighborhood of these new grid points but we need not re-compute it in the existing grid points. This useful feature is specific to approximations with local bases functions; for global approximations, we need to re-compute the solution entirely when changing either grid points or an approximating function.

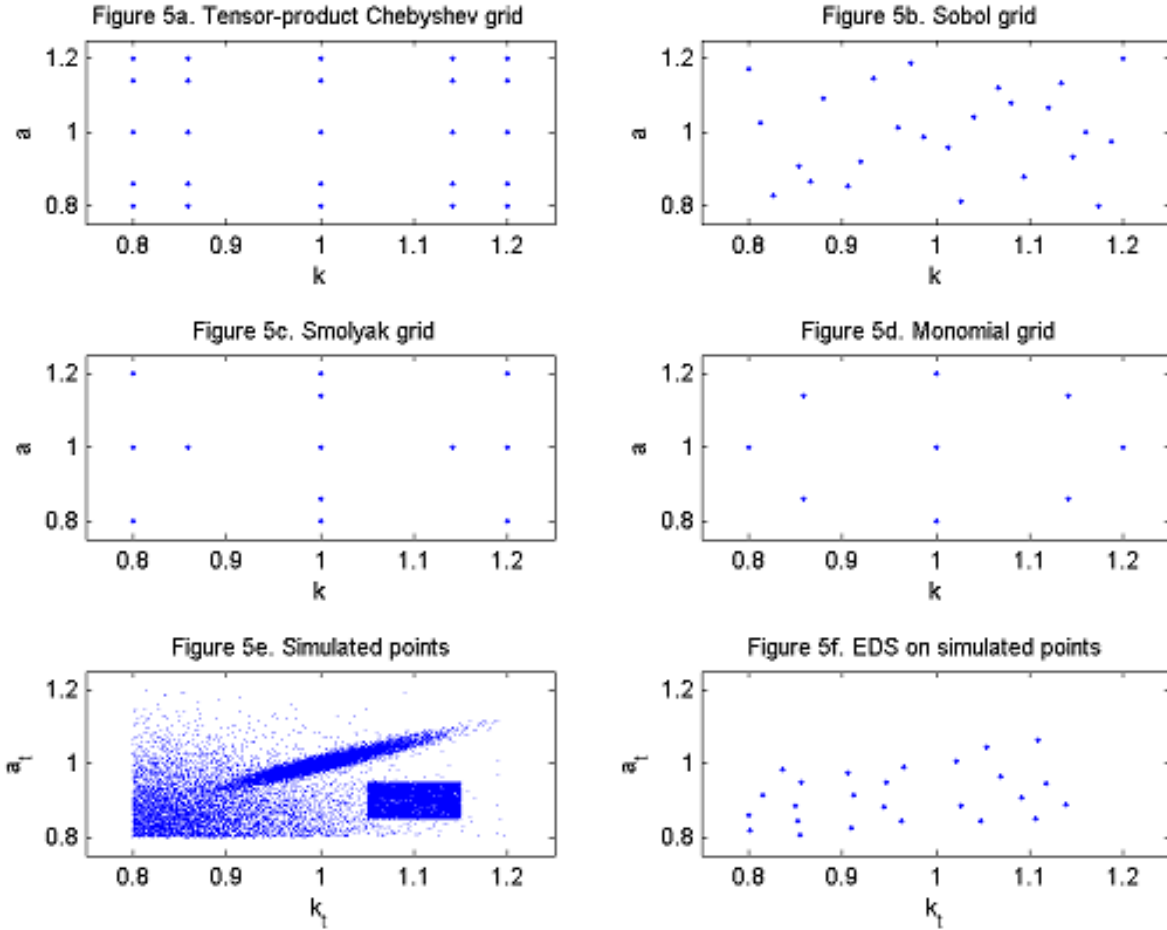
3 Incorporating the EDS grid into projection methods

In this section, we incorporate the EDS grid into projection methods for solving dynamic economic models, namely, we use the EDS grid as a set of points on which the solution is approximated.

3.1 Comparison of the EDS grid with other grids used in the context of numerical solution methods

Let us first compare the EDS grid to other grids used in the literature for solving dynamic economic models. We must make a distinction between a geometry of the set on which the solution is computed and a specific discretization of this set. A commonly-used geometry in the context of projection solution methods is a fixed multidimensional hypercube. Figures 5a-5d plot 4 different discretizations of the hypercube: a tensor-product Chebyshev

grid, a low-discrepancy Sobol grid, a sparse Smolyak grid and a monomial grid, respectively (in particular, these type of grids were used in Judd (1992), Rust (1998), Krueger and Kubler (2004), and Pichler (2011), respectively).⁷



In turn, stochastic simulation methods use the adaptive geometry; see Marcet (1988), Smith (1993), Maliar and Maliar (2005), Judd et al. (2011a) for examples of methods that compute solutions on simulated series.⁸ Focusing on the right geometry can be critical for the cost, as the following example shows.

⁷Also, Tauchen and Hussey (1991) propose a related discretization technique that delivers an approximation to a continuous density function of a given stochastic process. Their key idea is to approximate a Markov process with a finite-state Markov chain. This discretization technique requires to specify the distribution function of the Markov process explicitly and is primarily useful for forming discrete approximations of density functions of exogenous variables. In contrast, the EDS discretization technique builds on stochastic simulation and does not require to know the distribution function. It can be applied to both exogenous and endogenous variables.

⁸For a detailed description of Marcet's (1988) method, see Den Haan and Marcet (1990), and Marcet and Lorenzoni (1999).

Example. Consider a vector of uncorrelated random variables $x \in \mathbb{R}^d$ drawn from a multivariate Normal distribution $x \sim \mathcal{N}(0, I_d)$, where I_d is an identity matrix. An essentially ergodic set \mathcal{A}^n has the shape of a hypersphere. Let us surround such a hypersphere with a hypercube of a minimum size. For dimensions 2, 3, 4, 5, 10, 30 and 100, the ratio of the volume of a hypersphere to the volume of the hypercube is equal to 0.79, 0.52, 0.31, 0.16, $3 \cdot 10^{-3}$, $2 \cdot 10^{-14}$ and $2 \cdot 10^{-70}$, respectively. These numbers suggest that an enormous savings in cost are possible by focusing on an essentially ergodic set instead of the standard multidimensional hypercube.

However, a stochastic simulation is not an efficient discretization of a high-probability set: a grid of simulated points is unevenly spaced, has many closely-located, redundant points and contains some points in low-density regions.

The EDS grid is designed to combine the best features of the existing grids. It combines the adaptive geometry (similar to the one used by stochastic simulation methods) with an efficient discretization (similar to that produced by low-discrepancy methods on a hypercube). In Figure 5e, we show an example of a cloud of simulated points of irregular shape, and in Figure 5f, we plot the EDS grid delivered by the two-step procedure of Section 2.2. As we can see, the EDS grid appears to cover the high-probability set uniformly.

There are cases in which the EDS grid may be not a good choice. First, focusing on a high-probability set may not have advantages relatively to a hypercube; for example, if a vector $x \in \mathbb{R}^d$ is drawn from a multivariate uniform distribution, $x \sim [0, 1]^d$, then an essentially ergodic set coincides with the hypercube $[0, 1]^d$, and no saving in cost is possible. Second, in some applications, one may need to have a sufficiently accurate solution outside the high-probability set, for example, when analyzing a transition path of a developing economy with low initial endowment. In those cases, one may augment the grid on a high-probability set to include some "important" points situated outside this set. An example of this approach is shown by Aruoba and Schorfheide (2013) in the context of a new Keynesian model. They construct a grid by combining selected draws from the ergodic distribution of the model with a set of values for state variables filtered from the actual data. In this way, they augment the grid to include points from the 2008-2009 Great recession which do not naturally belong to a high-probability set of the studied model. Finally, our worst case analysis in Appendix A3 shows that there are scenarios in which EDSs constructed on simulated data are highly non-uniform. However, these scenarios require extreme assumptions about the density of simulated points, e.g., a set of highly uneven Dirac point masses. We did not observe the worst-case outcomes in our experiments. If we know that we are in those cases, we may opt for grids on a multidimensional hypercube.

3.2 General description of the EDS algorithm

In this section, we develop a projection method that uses the EDS grid. We focus on equilibrium problems, however, the EDS method can be also used to solve dynamic programming problems; see Judd et al. (2012) for examples.

3.2.1 An equilibrium problem

We study an equilibrium problem in which a solution is characterized by the set of equilibrium conditions for $t = 0, 1, \dots, \infty$,

$$E_t [G(s_t, z_t, y_t, s_{t+1}, z_{t+1}, y_{t+1})] = 0, \quad (3)$$

$$z_{t+1} = Z(z_t, \epsilon_{t+1}), \quad (4)$$

where (s_0, z_0) is given; E_t denotes the expectations operator conditional on information available at t ; $s_t \in \mathbb{R}^{d_s}$ is a vector of endogenous state variables at t ; $z_t \in \mathbb{R}^{d_z}$ is a vector of exogenous (random) state variables at t ; $y_t \in \mathbb{R}^{d_y}$ is a vector of non-state variables – prices, consumption, labor supply, etc. – also called non-predetermined variables; G is a continuously differentiable vector function; $\epsilon_{t+1} \in \mathbb{R}^p$ is a vector of shocks. A solution is given by a set of equilibrium functions $s_{t+1} = S(s_t, z_t)$, and $y_t = Y(s_t, z_t)$ that satisfy (3), (4) in the relevant area of the state space. In terms of notations of Section 2.1, we have $\varphi = (S, Y)$, $x_t = (s_t, z_t)$ and $d = d_s + d_z$. The solution (S, Y) is assumed to satisfy the assumptions of Section 2.1.

3.2.2 A projection algorithm based on the EDS grid

Our construction of the EDS grid in Section 2.2 is based on the assumption that the stochastic process (1) for the state variables is known. However, the law of motion for endogenous state variables is unknown before the model is solved: it is precisely our goal to approximate this law of motion numerically. We therefore proceed iteratively: guess a solution, simulate the model, construct an EDS grid, solve the model on that grid using a projection method, and iterate on these steps until the grid converges. Below, we elaborate a description of this procedure for the equilibrium problem (3), (4).

(EDS algorithm): A projection algorithm for equilibrium problems.

Step 0. Initialization.

- a. Choose (s_0, z_0) and simulation length, T .
 - b. Draw $\{\epsilon_{t+1}\}_{t=0, \dots, T-1}$. Compute and fix $\{z_{t+1}\}_{t=0, \dots, T-1}$ using (4).
 - c. Choose approximating functions $S \approx \widehat{S}(\cdot; b^s)$ and $Y \approx \widehat{Y}(\cdot; b^y)$.
 - d. Make an initial guess on b^s and b^y .
 - e. Choose integration nodes, ϵ_j , and weights, ω_j , $j = 1, \dots, J$.
-

Step 1. Construction of an EDS grid.

- a. Use $\widehat{S}(\cdot; b^s)$ to simulate $\{s_{t+1}\}_{t=0, \dots, T-1}$.
 - b. Construct an EDS grid, $\Gamma \equiv \{s_m, z_m\}_{m=1, \dots, M}$.
-

Step 2. Computation of a solution on EDS grid using a projection method.

- a. For $m = 1, \dots, M$, construct residuals

$$-\mathcal{R}(s_m, z_m) = \sum_{j=1}^J \omega_j \cdot G\left(s_m, z_m, y_m, s'_m, z'_{m,j}, y'_{m,j}\right),$$

where $y_m \equiv \widehat{Y}(s_m, z_m; b^y)$, $s'_m \equiv \widehat{S}(s_m, z_m; b^s)$, $z'_{m,j} \equiv Z(z_m, \epsilon_j)$;
 $-y'_{m,j} \equiv \widehat{Y}(s'_m, z'_{m,j}; b^y)$.

- b. Find b^s and b^y that minimize residuals according to some norm.
-

Iterate on Steps 1, 2 until convergence of the EDS grid.

3.2.3 Discussion of the computational choices

We construct the EDS grid as described in Section 2.2. We guess the equilibrium rule \widehat{S} , simulate the solution for T periods, construct a sample of n points by selecting each κ th observation, estimate the density function, remove a fraction δ of the sample with the lowest density, and construct an EDS grid with a target number of points \overline{M} using a bisection method. Below, we discuss some of the choices related to the construction of the EDS grids.

Initial guess on b^s . To insure that the EDS grid covers the right area of the state space, we need a sufficiently accurate initial guess about the equilibrium rules. Furthermore, the equilibrium rules used must lead to non-explosive simulated series. For many problems in economics, linear solutions can be used as an initial guess; they are sufficiently accurate, numerically stable and readily available from automated perturbation software (we use Dynare solutions; see Adjemian et al., 2011). Finding a sufficiently good initial guess can be a non-trivial issue in some applications, and techniques from learning literature can be useful in this context; see Bertsekas and Tsitsiklis (1996) for a discussion.

Choices of n and T . Our construction of an EDS relies on the assumption that simulated points are sufficiently dense on the essentially ergodic set. Technically, in Appendix A1, we require that each ball $B(x; \varepsilon)$ inside \mathcal{A}^η contains at least one simulated point. The probability $\Pr(0)$ of having no points in a ball $B(x; \varepsilon)$ inside \mathcal{A}^η after n draws satisfies

$\Pr(0) \leq (1 - p_\varepsilon)^n$ where $p_\varepsilon \equiv \int_{B(x;\varepsilon)} \eta dx \approx \lambda_d \varepsilon^d \eta$ and λ_d is the volume of a d -dimensional unit ball. (Note that on the boundary of \mathcal{A}^η where $g = \eta$, we have $\Pr(0) = (1 - p_\varepsilon)^n$). Thus, given ε and η , we must choose n and $T = n\kappa$, so that $\Pr(0)$ is sufficiently small. We use $T = 100,000$ and $\kappa = 10$, so that our sample has $n = 10,000$ points, and we choose η to remove 1% of points with the lowest density.

Choices of ε and \overline{M} . We need to have at least as many grid points in the EDS as the parameters b^s and b^y in \widehat{S} and \widehat{Y} (to identify these parameters). Conventional projection methods rely on collocation, when the number of grid points is the same as the number of parameters to identify. Collocation is a useful technique in the context of orthogonal polynomial constructions but is not convenient in our case (because our bisection method does not guarantee that the number of grid points is exactly equal to the target number \overline{M}). Hence, we target a slightly larger number of points than parameters, which also helps us to increase both accuracy and numerical stability.

Reconstructing the EDS grid iteratively. Under Assumptions 1 and 2, the convergence of the equilibrium rules implies the convergence of the time-series solution; see Peralta-Alva and Santos (2005). Therefore, we are left to check that the EDS grid constructed on the simulated series also converges. Let $\Gamma' \equiv \{x'_i\}_{i=1,\dots,M'}$ and $\Gamma'' \equiv \{x''_j\}_{j=1,\dots,M''}$ be the EDS grids constructed on two different sets of simulated points. Our criteria of convergence is $\sup_{x''_j \in \Gamma''} \inf_{x'_i \in \Gamma'} D(x'_i, x''_j) < 2\varepsilon$. That is, each grid point of Γ'' has a grid point of Γ' at the distance smaller than 2ε (this is the maximum distance between the grid points on the essentially ergodic set; see Proposition 3 in Appendix A3).

How often do we need to reconstruct the EDS grid? Constructing EDS grids may be costly, especially, in problems with high dimensionality because we need to produce a long simulation, to estimate the density function and to construct EDS grids several times until a bisection procedure locates a grid with the target number of grid points. The cost of constructing EDS grids can be especially high in those applications in which researchers must solve their models repeatedly using different parameters vectors, for example, in estimation or calibration studies.

Hence, an important question is: "How often do we need to reconstruct the EDS grid in a given application?" We found that typically, the properties of solutions are not sensitive to small changes in the EDS grid. For example, the EDS grid constructed on a log-linear solution would normally lead to as accurate non-linear solutions as the one constructed using highly accurate non-linear solutions. Furthermore, we found that small changes in the model's parameters do not require us to re-compute the grid. In the presence of kinks, such as the ZLB in new Keynesian model, the solution is more sensitive to a specific construction of the EDS grid, however, using more accurate solutions for constructing the grid does not necessarily lead to smaller approximation errors. Thus,

our experiments suggest that in many applications, we can construct an EDS grid just once using a relatively rough initial guess, and we can keep this grid when iterating on decision functions until convergence (without a visible accuracy loss).

Integration. Unlike simulation- and learning-based methods, we rely on deterministic integration methods such as the Gauss-Hermite quadrature and monomial integration methods. Deterministic methods dominate in accuracy the Monte Carlo method by orders of magnitude in the context of the studied class of models; see Judd et al. (2011a) for comparison results.⁹ The cost of Gaussian product rules is prohibitive in high-dimensional problems but monomial formulas are tractable even in models with hundreds of state variables; see Judd et al. (2011b) for the description of these formulas.

Solving systems of non-linear equations: the convergence issue. In Step 2 of the EDS algorithm, we must find the parameters vector $b \equiv (b^s, b^y)$ in the decision functions $\widehat{S}(\cdot; b^s)$ and $\widehat{Y}(\cdot; b^y)$ that satisfy the model's equations. Here, we have a system of $n = M \times H$ non-linear equations, where H is the number of equations in the vector function G with n' unknown parameters in b . By construction, $n' \leq n$. If $n' = n$, i.e., we have the same number of unknowns (grid points) as equations, we may have a unique solution that satisfies all equations exactly (this case is referred to as collocation). However, if $n' < n$, we construct a solution that satisfies the model's equations by minimizing a weighed sum of residuals in the model's equations (this case is similar to regression in econometrics).

There is a variety of numerical methods in the literature that can be used to solve a system of non-linear equations in Step 2 of the EDS algorithm, see, e.g., Judd (1998, pp 93-128)) for a review of such methods. In the paper, we restrict attention to a simple derivative-free fixed-point iteration method; see Wright and Williams (1984), Marcet (1988), Den Haan (1990), Gaspar and Judd (1997) for early applications of fixed-point iteration to economic problems. In terms of our problem, fixed-point iteration can be written as follows:

(FPI): Fixed-point iteration with damping.

Initialization. Write a system of equations in the form $\widehat{b} = \Psi(b)$.

Fix initial guess $b^{(0)}$, a norm $\|\cdot\|$ and a convergence criterion ϖ .

Step 1. On iteration i , compute $\widehat{b} = \Psi(b^{(i)})$.

Step 2. If $\|\widehat{b} - b^{(i)}\| < \varpi$, then stop.

Otherwise, set $b^{(i+1)} = \xi \widehat{b} + (1 - \xi) b^{(i)}$, where $\xi \in (0, 1]$ and go to *Step 1*.

That is, for iteration i , we guess some $b^{(i)}$, compute new \widehat{b} and use it to update our guess

⁹For example, assume that a Monte Carlo method is used to approximate an expectation of $y \sim \mathcal{N}(0, \sigma_y)$ with n random draws. The distribution of $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ is $\bar{y} \sim \mathcal{N}\left(0, \frac{\sigma_y}{\sqrt{n}}\right)$. If $\sigma_y = 1\%$ and $n = 10,000$, we have approximation errors of order $\frac{\sigma_y}{\sqrt{n}} = 10^{-4}$. To bring the error to the level of 10^{-8} , which we attain using quadrature methods, we need to have $n = 10^{12}$. That is, such a slow, \sqrt{n} -rate of convergence makes it very expensive to obtain highly accurate solutions using stochastic simulation.

for iteration $i + 1$, where ξ is the damping parameter that controls the speed of updating. The advantage of fixed-point iteration is that it can iterate in this simple manner on objects of any dimensionality, for example, on a vector of the polynomial coefficients. The cost of this procedure does not grow rapidly with dimensionality of the problem, unlike does the cost of Newton-style methods. As other non-linear solvers, fixed-point iteration may fail to converge. The following example, borrowed from Judd (1998, p 159), illustrates the possibility of non-convergence.

Example. Let us find a solution to $x^3 - x - 1 = 0$ using a fixed point iteration. We can rewrite it as $x = (x + 1)^{1/3}$ and construct a sequence $x^{(i+1)} = (x^{(i)} + 1)^{1/3}$ starting from $x^{(0)} = 1$. This yields a sequence $x^{(1)} = 1.26$, $x^{(2)} = 1.31$, $x^{(3)} = 1.32, \dots$ which converges to a solution. However, we can also rewrite this equation as $x = x^3 - 1$ and construct a sequence $x^{(i+1)} = (x^{(i)})^3 - 1$ starting from $x^{(0)} = 1$ which diverges to $-\infty$.

This example shows that whether fixed-point iteration succeeds or not in finding a solution may depend on a specific way in which it is implemented. Judd (1998, pp 557-558) also shows that fixed-point iteration may fail to converge in growth models like the ones studied in the present paper under some parameterizations. Damping helps us to increase the likelihood of convergence; see Judd (1998, pp 78-84). Newton-style methods may have better convergence properties but they may also fail if an initial guess is not sufficiently accurate. In sum, CGA and EDS solution methods are effective numerical methods for solving high-dimensional applications, but they share limitations that are common for all projection methods, namely, they may fail to converge. In our examples, the studied EDS and CGA solution methods were highly accurate and reliable, however, the reader must be aware of the existence of the above potential problems and must be ready to detect and to address such problems if they arise in applications.

3.2.4 Evaluating the accuracy of solutions

Provided that the EDS algorithm succeeds in producing a candidate solution, we subject such a solution to a tight accuracy check. We specifically generate a set of points within the domain on which we want the solution to be accurate, and we compute residuals in all equilibrium conditions.

(Evaluation of accuracy): Residuals in equilibrium conditions.

- a. Choose a set of points $\{s_\tau, z_\tau\}_{\tau=1, \dots, \mathcal{T}^{\text{test}}}$ for evaluating the accuracy.
- b. For $\tau = 1, \dots, \mathcal{T}^{\text{test}}$, compute the size of the residuals:

$$\mathcal{R}(s_\tau, z_\tau) \equiv \sum_{j=1}^{J^{\text{test}}} \omega_j^{\text{test}} \cdot \left[G \left(s_\tau, z_\tau, y_\tau, s'_\tau, z'_{\tau,j}, y'_{\tau,j} \right) \right],$$

where $y_\tau = \widehat{Y}(s_\tau, z_\tau; b^y)$, $s'_\tau = \widehat{S}(s_\tau, z_\tau; b^s)$,

$$z'_{\tau,j} = Z \left(z_\tau, \epsilon_j^{\text{test}} \right), y'_{\tau,j} = \widehat{Y} \left(y_\tau, z'_{\tau,j}; b^y \right),$$

ϵ_j^{test} and ω_j^{test} are the integration nodes and weights.

- c. Find a mean and/or maximum of the residuals $\mathcal{R}(s_\tau, z_\tau)$.
-
-

If the quality of a candidate solution is economically unacceptable, we modify the choices made in the EDS algorithm (i.e., simulation length, number of grid points, approximating functions, integration method) and recompute the solution. In the paper, we evaluate the accuracy on a set of simulated points. This new set of points which is different from that used in the solution procedure: it is constructed under a different sequence of shocks (i.e., we test accuracy out of sample). Other possible accuracy checks include evaluating the residuals in the model's equations on a given set of points in the state space (Judd, 1992), and testing the orthogonality of residuals in the optimality conditions (Den Haan and Marcet, 1994); see Santos (2000) for a discussion.

4 Neoclassical stochastic growth model

In this section, we use the EDS approach to solve the standard neoclassical stochastic growth model. We discuss some relevant computational choices and assess the performance of the algorithm in one- and multi-agent setups.

4.1 The set up

The representative agent solves

$$\max_{\{k_{t+1}, c_t\}_{t=0, \dots, \infty}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \quad (5)$$

$$\text{s.t. } c_t + k_{t+1} = (1 - \delta) k_t + a_t A f(k_t), \quad (6)$$

$$\ln a_{t+1} = \rho \ln a_t + \epsilon_{t+1}, \quad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma^2), \quad (7)$$

where (k_0, a_0) is given; E_t is the expectation operator conditional on information at time t ; c_t , k_t and a_t are consumption, capital and productivity level, respectively; $\beta \in (0, 1)$; $\delta \in (0, 1]$; $A > 0$; $\rho \in (-1, 1)$; $\sigma \geq 0$; u and f are the utility and production functions, respectively, both of which are strictly increasing, continuously differentiable and concave. Under our assumptions, this model has a unique solution; see, e.g., Stokey and Lucas with Prescott (1989, p 392). For numerical experiments, we use $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$ with $\gamma \in \{\frac{1}{5}, 1, 5\}$ and $f(k) = k^\alpha$ with $\alpha = 0.36$, and we set $\beta = 0.99$, $\delta = 0.025$, $\rho = 0.95$ and $\sigma = 0.01$. A version of the model under $u(c) = \ln(c)$, $\delta = 1$ and $f(k) = k^\alpha$ admits a closed-form solution $k_{t+1} = \alpha\beta a_t A k_t^\alpha$.

4.2 An EDS algorithm iterating on the Euler equation

We describe an example of the EDS method that iterates on the Euler equation. For the model (5)–(7), the Euler equation is

$$1 = \beta E \left[\frac{u'(c')}{u'(c)} (1 - \delta + a' A f'(k')) \right], \quad (8)$$

where primes on the variables denote next-period values, and u' and f' denote the derivatives of u and f , respectively. We must solve for equilibrium rules $c = C(k, a)$ and $k' = K(k, a)$ that satisfy (6)–(8). To implement fixed-point iteration, we represent (8) in the form $k' = \Psi(k')$ by multiplying both sides with k' which yields $\widehat{k}' = \beta k' E \left[\frac{u'(c')}{u'(c)} (1 - \delta + a' A f'(k')) \right]$. In our iterative procedure, we substitute $(k')^{(i)}$ obtained in iteration i in the right side of this equation, compute \widehat{k}' in the left side and use the solution to improve our guess $(k')^{(i+1)}$ for iteration $i + 1$; see Appendix C for a detailed description of the EDS solution method.

In Table 1, we provide the results for the Euler equation EDS algorithm under the target number of grid points $\overline{M} = 25$ points.

Table 1: Accuracy and speed of the Euler equation EDS algorithm in the one-agent model.^a

Polynomial degree	$\gamma = 1/5$			$\gamma = 1$			$\gamma = 5$		
	$M(\varepsilon) = 21$			$M(\varepsilon) = 27$			$M(\varepsilon) = 25$		
	L_1	L_∞	CPU	L_1	L_∞	CPU	L_1	L_∞	CPU
1st	-4.74	-3.81	25.5	-4.29	-3.31	24.7	-3.29	-2.35	23.6
2nd	-6.35	-5.26	1.8	-5.94	-4.87	0.8	-4.77	-3.60	0.4
3rd	-7.93	-6.50	1.9	-7.26	-6.04	0.9	-5.97	-4.47	0.4
4th	-9.37	-7.60	2.0	-8.65	-7.32	0.9	-7.05	-5.26	0.4
5th	-9.82	-8.60	14.25	-9.47	-8.24	5.5	-7.89	-6.46	2.8

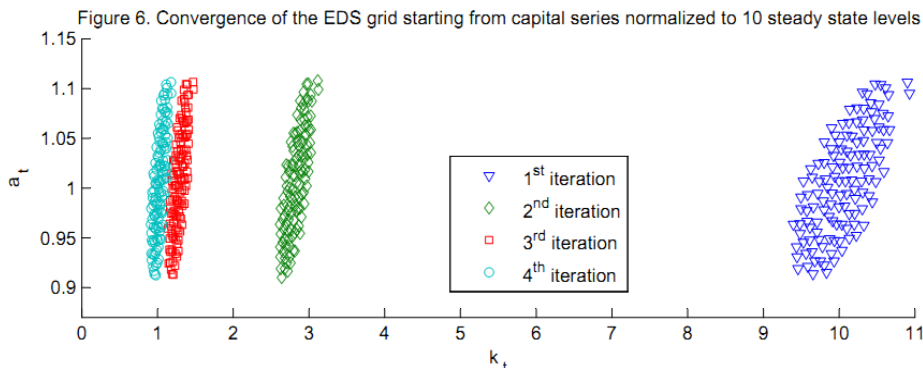
^a Notes: L_1 and L_∞ are, respectively, the average and maximum of absolute residuals across optimality condition and test points (in log10 units) on a stochastic simulation of 10,000 observations; CPU is the time necessary for computing a solution (in seconds); γ is the coefficient of risk aversion; $M(\varepsilon)$ is the realized number of points in the EDS grid (the target number of grid points is $\overline{M}=25$).

The accuracy of solutions delivered by the EDS algorithm is comparable to the highest accuracy attained in the related literature. The residuals in the optimality conditions decrease with each polynomial degree by one or more orders of magnitude. For the fifth-degree polynomials, the largest unit-free residual corresponding to our least accurate solution is still less than 10^{-6} (see the experiment with a high degree of risk aversion $\gamma = 5$). Computing high-degree polynomial solutions is relatively fast (a few seconds). Most of the cost of the EDS algorithm comes from the construction of the EDS grid (here, the EDS grid is constructed just once under the polynomial solution of degree 1, and the time for the grid construction is included in the total time for computing that solution). The cost increases gradually with the degree of approximating polynomial because we have a larger number of terms in the approximating functions. This results in a "U"-shaped pattern for the CPU time in the table.

We perform sensitivity experiments in which we vary the target number of grid points and find that the results are robust to the modifications considered. We also vary the number of nodes in the Gauss-Hermite quadrature rule, and we find that even the 2-node rule leads to essentially the same accuracy levels as the 10-node rule (except the fourth

and fifth-degree polynomials under which the accuracy is somewhat lower). This result is in line with the finding of Judd (1992) that in the context of the standard optimal growth model, even few quadrature nodes lead to very accurate solutions.

Autocorrection of the EDS grid. If an initial guess about the solution is poor, the simulated series will not cover the ergodic set. Will the EDS grid be able to autocorrect itself in the context of the EDS algorithm? A general answer to this question is unknown. However, we observe autocorrection of the EDS grid in numerical experiments. In one of such experiments, we scale up the time-series solution for capital by a factor of 10, and use the resulting series for constructing the first EDS grid (thus, the capital values in this grid are spread around 10 instead of 1). We solve the model on this grid and use the solution to construct the second EDS grid. We repeat this procedure two more times. Figure 6 shows that the EDS grid converges rapidly.



We tried out various initial guesses away from the essentially ergodic set, and we observed autocorrection of the EDS grid in all the experiments performed. Furthermore, the EDS grid approach had an autocorrection property in our challenging applications such as a multi-agent neoclassical growth model and a new Keynesian model with a zero lower bound on nominal interest rates.

EDS grid versus Smolyak grid. Krueger and Kubler (2004) and Malin et al. (2011) develop a projection method that relies on a Smolyak sparse grid. To isolate the role of the grid construction in the algorithm’s performance, we implement the Smolyak method in the same way as the EDS method, namely, we use an ordinary polynomial family for approximating decision functions, and we use fixed-point iteration for finding the polynomial coefficients. This implementation of the Smolyak method is in line with the one studied in Judd et al. (2013) and differs from the one in Krueger and Kubler (2004) and Malin et al. (2011) that builds on Smolyak polynomial function and time iteration (in particular, time iteration is more expensive than fixed-point iteration, and Smolyak polynomials have 4 times more basis functions and thus, are more flexible than ordinary

Table 2: Accuracy and speed in the one-agent model: Smolyak grid versus EDS grid.^a

Polynomial degree	Smolyak grid					Accuracy on a hypercube				
	Simulation		Rectangular		CPU	Simulation		Rectangular		CPU
	L ₁	L _∞	L ₁	L _∞		L ₁	L _∞	L ₁	L _∞	
1st	-3.31	-2.94	-3.25	-2.54	0.39	-4.23	-3.31	-3.26	-2.38	9.89
2nd	-4.74	-4.17	-4.32	-3.80	0.20	-5.89	-4.87	-4.41	-3.25	0.19
3rd	-5.27	-5.13	-5.39	-4.78	0.22	-7.19	-6.16	-5.44	-4.11	0.17

^a Notes: L₁ and L_∞ are, respectively, the average and maximum of absolute residuals across optimality condition and test points (in log10 units) on a stochastic simulation of 10,000 observations and on a fixed grid of 100×100 points covering the rectangular with the intervals [0.8, 1.2] and $\left[\exp\left(-\frac{0.8}{1-\rho}\right), \exp\left(\frac{0.8}{1-\rho}\right)\right]$ for capital and productivity, respectively; CPU is the time necessary for computing a solution (in seconds).

polynomials). Thus, under our implementation, the EDS and Smolyak methods differ only in the choice of grid points.

As in Malin et al. (2011), we use intervals [0.8, 1.2] and $\left[\exp\left(-\frac{0.8}{1-\rho}\right), \exp\left(\frac{0.8}{1-\rho}\right)\right]$ for capital and productivity level, respectively. The Smolyak grid has 13 points (see Figure 5c), so we use an EDS grid with the same number of points. With 13 grid points, we can identify the coefficients in ordinary polynomials up to degree 3. In this experiment, we evaluate the accuracy of solutions not only on a stochastic simulation but also on a set of 100 × 100 points which are uniformly spaced on the same domain as the one used by the Smolyak method for finding a solution. The accuracy results are shown in Table 2. The running time is similar for the Smolyak and EDS methods except that the EDS method needs additional time for constructing the grid.

In the test on a stochastic simulation, the EDS grid leads to considerably more accurate solutions than the Smolyak grid. This is because under the EDS grid, we fit a polynomial directly in the essentially ergodic set, while under the Smolyak grid, we fit a polynomial in a larger rectangular domain and face a trade-off between the fit inside and outside the ergodic set. However, in the test on the rectangular domain, the Smolyak grid produces significantly smaller maximum residuals than the EDS grid. This is because the EDS algorithm is designed to be accurate in the essentially ergodic set and its accuracy decreases more rapidly away from the essentially ergodic set than the accuracy of methods operating on larger hypercube domains. We repeated this experiment by varying the intervals for capital and productivity in the Smolyak grid, and we have the same regularities. These regularities are also observed in high-dimensional applications.¹⁰

¹⁰Kollmann et al. (2011b) compare the accuracy of solutions produced by several solution methods, including the cluster grid algorithm (CGA) introduced in the earlier version of the present paper and the Smolyak algorithm of Krueger and Kubler (2004) (see Maliar et al., 2011, and Malin et al., 2011, for implementation details of the respective methods in the context of those models). Their comparison is performed using a collection of 30 real-business cycle models with up to 10 heterogeneous agents. Their findings are the same as ours: on a stochastic simulation and near the steady state, the CGA solutions are more accurate than the Smolyak solutions whereas the situation reverses for large deviations from

4.3 EDS algorithm in problems with high dimensionality

We now explore the tractability of the EDS algorithm in problems with high dimensionality. We extend the one-agent model (5)–(7) to include multiple agents. This is a simple way to expand and to control the size of the problem.

The set up. There are N agents, interpreted as countries, which differ in initial capital endowment and productivity levels. The countries' productivity levels are affected by both country-specific and worldwide shocks. We study the social planner's problem. A social planner maximizes a weighted sum of expected lifetime utilities of N agents (countries),

$$\max_{\{c_t^h, k_{t+1}^h\}_{t=0, \dots, \infty}^{h=1, \dots, N}} E_0 \sum_{h=1}^N \lambda^h \left[\sum_{t=0}^{\infty} \beta^t u^h(c_t^h) \right] \quad (9)$$

subject to the aggregate resource constraint,

$$\sum_{h=1}^N c_t^h + \sum_{h=1}^N k_{t+1}^h = \sum_{h=1}^N k_t^h (1 - \delta) + \sum_{h=1}^N a_t^h A f^h(k_t^h), \quad (10)$$

where $\{k_0^h, a_0^h\}^{h=1, \dots, N}$ is given; E_t is the operator of conditional expectation; c_t^h , k_t^h , a_t^h and λ^h are, respectively, consumption, capital, productivity level and welfare weight of a country $h \in \{1, \dots, N\}$; $\beta \in (0, 1)$ is the discount factor; $\delta \in (0, 1]$ is the depreciation rate; A is the normalizing constant in the production function. The utility and production functions, u^h and f^h , respectively, are increasing, concave and continuously differentiable. The process for the productivity level of country h is given by

$$\ln a_{t+1}^h = \rho \ln a_t^h + \epsilon_{t+1}^h, \quad (11)$$

where ρ is the autocorrelation coefficient; $\epsilon_{t+1}^h \equiv \varsigma_{t+1}^h + \varsigma_{t+1}$ where $\varsigma_{t+1}^h \sim \mathcal{N}(0, \sigma^2)$ is specific to each country and $\varsigma_{t+1} \sim \mathcal{N}(0, \sigma^2)$ is identical for all countries.

We restrict our attention to the case in which the countries are characterized by identical preferences, $u^h = u$, and identical production technologies, $f^h = f$, for all h . The former implies that the planner assigns identical weights, $\lambda^h = 1$, and consequently, identical consumption $c_t^h = c_t$ to all agents. If an interior solution exists, it satisfies N Euler equations,

$$u'(c_t) = \beta E_t \left\{ u'(c_{t+1}) \left[1 - \delta + a_{t+1}^h A f'(k_{t+1}^h) \right] \right\}, \quad (12)$$

where u' and f' denote the derivatives of u and f , respectively. Thus, the planner's solution is determined by the process for productivity (11), the resource constraint (10), and the set of Euler equations (12). We use the same values of the parameters for the multicountry model as in the one-agent model; in particular, we assume $\gamma = 1$.

the steady state.

Solution procedure. Our objective is to approximate the planner’s solution in the form of N capital equilibrium rules, each of which depends on $2N$ state variables (N capital stocks and N productivity levels), i.e., $k_{t+1}^h = K^h \left(\{k_t^h, a_t^h\}^{h=1, \dots, N} \right)$, $h = 1, \dots, N$. Since the countries are identical in their fundamentals (preferences and technology), the planner chooses the same level of consumption for all countries. We could have used this symmetry to simplify the solution procedure, however, we do not do so. Instead, we compute a decision rule of each country separately, treating them as completely heterogeneous. In this manner, we can assess the cost of finding solutions in general multidimensional setups in which countries have heterogeneous preferences and technology. For each country, we essentially implement the same computational procedure as the one used in the representative-agent case; see Appendix D for details of the computational procedure.

The choice of an integration method used plays an important role in the accuracy and speed of our solution algorithm. The Monte Carlo method produces large sampling errors which dominate the overall accuracy of solutions. Quadrature product rules are accurate but their cost is prohibitive if the number of shocks is large. However, non-product monomial integration methods both produce very accurate solutions and are tractable in problems with high dimensionality. Moreover, in the studied class of models, an extremely simple and cheap integration method – a one-node quadrature rule – happens to produce accurate solutions; see Judd et al. (2011a) for a detailed description of this and other integration methods; and see Judd et al. (2011a, 2012) and Maliar et al. (2011) for accuracy comparisons of different integration methods in large-scale applications.

Determinants of cost in problems with high dimensionality. The cost of finding numerical solutions increases with the dimensionality of the problem for various reasons. There are more equations to solve and more decision functions to approximate. The number of terms in an approximating polynomial function goes up, and we need to increase the number of grid points to identify the polynomial coefficients. The number of nodes in integration formulas also increases. Finally, operating with large data sets can lead to a memory congestion. If a solution method relies on product rules in constructing a grid, integration nodes, optimization, its cost increases exponentially (curse of dimensionality). However, our design of the EDS method does not rely on product rules and its cost grows with dimensionality of the problem at a relatively moderate rate.

Accuracy and cost of solutions. We solve the model with N ranging from 2 to 200. The results about the accuracy and cost of solutions are provided in Table 3.

The accuracy of solutions here is similar to that we had for the one-agent model. For the polynomial approximations of degrees 1 and 2, the residuals are typically smaller than 0.1% and 0.01%, respectively. These regularities are robust to variations in the model’s parameters such as the volatility and persistence of shocks and the degrees of risk aversion; for sensitivity results, see Table 8 in an earlier version of the present paper, Judd et al. (2010).

Table 3: Accuracy and speed of the EDS algorithm in the multi-agent model.^a

Polyn. degree	$N = 2,$			$N = 20$			$N = 40$			$N = 200$		
	$M = 300, 2N$ nodes			$M = 1000, 2N$ nodes			$M = 4000, 1\text{-node}$			$M = 1000, 1\text{-node}$		
	L_1	L_∞	CPU	L_1	L_∞	CPU	L_1	L_∞	CPU	L_1	L_∞	CPU
1st	-4.70	-3.17	0.7	-4.76	-3.05	21	-4.79	-3.09	89	-4.66	-2.90	105
2nd	-6.01	-4.06	1.9	-5.88	-4.14	282	-5.48	-4.13	1463			

^a Notes: L_1 and L_∞ are, respectively, the average and maximum of absolute residuals across optimality condition and test points (in \log_{10} units) on a stochastic simulation of 10,000 observations; CPU is the time necessary for computing a solution (in minutes); \overline{M} is the target number of points in the EDS grid, respectively; " $2N$ " and "1-node" are the monomial rule with two nodes and one-node Gauss-Hermite integration rule, respectively.

The running time ranges from 36 seconds to 24 hours depending on the number of countries, polynomial degree and the integration technique used; see Judd et al. (2012) for sensitivity experiments. In particular, the EDS algorithm is able to compute quadratic solutions to the models with up to 40 countries and linear solutions to the models with up to 200 countries when using inexpensive one-node quadrature integration rule. Thus, the EDS algorithm is tractable in much larger problems than those studied in related literature. A proper coordination between the choices of approximating function and integration technique is critical in problems with high dimensionality. An example of such a coordination is a combination of a flexible second-degree polynomial with a cheap one-node Gauss-Hermite quadrature rule (in the given application, this cheap combination produces a very accurate approximation).

5 New Keynesian model with the ZLB

In this section, we use the EDS algorithm to solve a stylized new Keynesian model with Calvo-type price frictions and a Taylor (1993) rule.¹¹ Our setup builds on the models considered in Christiano et al. (2005), Smets and Wouters (2003, 2007), Del Negro et al. (2007). This literature estimates new Keynesian models using the data on actual economies, while we use their parameters estimates and compute solutions numerically. We solve two versions of the model, one in which we allow for negative nominal interest rates and the other in which we impose a zero lower bound (ZLB) on nominal interest rates.¹² The studied model has eight state variables and is large-scale in the sense that it is expensive or even intractable under conventional global solution methods that rely on

¹¹In an earlier version of the present paper, Judd et al. (2011b) use a cluster grid algorithm (CGA) to solve a new Keynesian model which is identical to the one studied here using an EDS algorithm except of parameterization.

¹²For the neoclassical growth model studied in Section 4, it would be also interesting to explore the case with occasionally binding borrowing constraints. Christiano and Fisher (2000) show how projection methods could be used to solve such a model.

product rules.

The literature that finds numerical solutions to new Keynesian models typically relies on local perturbation solution methods or applies expensive global solution methods to low-dimensional problems. As for perturbation methods, most papers compute linear approximations, however, there are papers that compute quadratic approximations (e.g., Kollmann, 2002, and Schmitt-Grohé and Uribe, 2007) and cubic approximations (e.g., Rudebusch and Swanson, 2008). The earlier literature that used global solution methods includes Adam and Billi (2006), Anderson et al. (2010), and Adjemian and Juillard (2011). The above studies have either few state variables or employ simplifying assumptions.¹³ However, recent literature, equipped with novel solution methods, started an exploration of medium- and large-scale new Keynesian models. In particular, Judd et al. (2011b), Fernández-Villaverde et al. (2012) and Braun et al. (2012) show that conventional perturbation methods do not provide accurate approximations in the context of new Keynesian models with the ZLB. Moreover, recent papers by Schmitt-Grohé and Uribe (2012), Mertens and Ravn (2013) and Aruoba and Schorfheide (2013) find that new Keynesian models may have multiple equilibria in the presence of ZLB, and in particular, the last paper accurately computes such multiple equilibria with a full set of stochastic shocks. In the present paper, we focus on the conventional equilibrium in which target inflation coincides with actual inflation in the steady state.

5.1 The set up

The economy is populated by households, final-good firms, intermediate-good firms, monetary authority and government; see Galí (2008, Chapter 3) for a detailed description of the baseline new Keynesian model.

Households. The representative household solves

$$\max_{\{C_t, L_t, B_t\}_{t=0, \dots, \infty}} E_0 \sum_{t=0}^{\infty} \beta^t \exp(\eta_{u,t}) \left[\frac{C_t^{1-\gamma} - 1}{1-\gamma} - \exp(\eta_{L,t}) \frac{L_t^{1+\vartheta} - 1}{1+\vartheta} \right] \quad (13)$$

$$\text{s.t. } P_t C_t + \frac{B_t}{\exp(\eta_{B,t}) R_t} + T_t = B_{t-1} + W_t L_t + \Pi_t, \quad (14)$$

where $(B_0, \eta_{u,0}, \eta_{L,0}, \eta_{B,0})$ is given; C_t , L_t , and B_t are consumption, labor and nominal bond holdings, respectively; P_t , W_t and R_t are the commodity price, nominal wage and (gross) nominal interest rate, respectively; $\eta_{u,t}$ and $\eta_{L,t}$ are exogenous preference shocks to the overall momentary utility and disutility of labor, respectively; $\eta_{B,t}$ is an exogenous premium in the return to bonds; T_t is lump-sum taxes; Π_t is the profit of intermediate-good firms; $\beta \in (0, 1)$ is the discount factor; $\gamma > 0$ and $\vartheta > 0$ are the utility-function

¹³In particular, Adam and Billi (2006) linearize all the first-order conditions except for the non-negativity constraint for nominal interest rates, and Adjemian and Juillard (2011) assume perfect foresight to implement an extended path method of Fair and Taylor (1984).

parameters. The processes for shocks are

$$\eta_{u,t+1} = \rho_u \eta_{u,t} + \epsilon_{u,t+1}, \quad \epsilon_{u,t+1} \sim \mathcal{N}(0, \sigma_u^2), \quad (15)$$

$$\eta_{L,t+1} = \rho_L \eta_{L,t} + \epsilon_{L,t+1}, \quad \epsilon_{L,t+1} \sim \mathcal{N}(0, \sigma_L^2), \quad (16)$$

$$\eta_{B,t+1} = \rho_B \eta_{B,t} + \epsilon_{B,t+1}, \quad \epsilon_{B,t+1} \sim \mathcal{N}(0, \sigma_B^2), \quad (17)$$

where $\rho_u, \rho_L, \rho_B \in (-1, 1)$, and $\sigma_u, \sigma_L, \sigma_B \geq 0$.

Final-good firms. Perfectly competitive final-good firms produce final goods using intermediate goods. A final-good firm buys $Y_t(i)$ of an intermediate good $i \in [0, 1]$ at price $P_t(i)$ and sells Y_t of the final good at price P_t in a perfectly competitive market. The profit-maximization problem is

$$\max_{Y_t(i)} P_t Y_t - \int_0^1 P_t(i) Y_t(i) di \quad (18)$$

$$\text{s.t. } Y_t = \left(\int_0^1 Y_t(i)^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad (19)$$

where (19) is a Dixit-Stiglitz aggregator function with $\varepsilon \geq 1$.

Intermediate-good firms. Monopolistic intermediate-good firms produce intermediate goods using labor and are subject to sticky prices. The firm i produces the intermediate good i . To choose labor in each period t , the firm i minimizes the nominal total cost, TC (net of government subsidy v),

$$\min_{L_t(i)} \text{TC}(Y_t(i)) = (1-v) W_t L_t(i) \quad (20)$$

$$\text{s.t. } Y_t(i) = \exp(\eta_{a,t}) L_t(i), \quad (21)$$

$$\eta_{a,t+1} = \rho_a \eta_{a,t} + \epsilon_{a,t+1}, \quad \epsilon_{a,t+1} \sim \mathcal{N}(0, \sigma_a^2), \quad (22)$$

where $L_t(i)$ is the labor input; $\exp(\eta_{a,t})$ is the productivity level; $\rho_a \in (-1, 1)$, and $\sigma_a \geq 0$. The firms are subject to Calvo-type price setting: a fraction $1 - \theta$ of the firms sets prices optimally, $P_t(i) = \tilde{P}_t$, for $i \in [0, 1]$, and the fraction θ is not allowed to change the price and maintains the same price as in the previous period, $P_t(i) = P_{t-1}(i)$, for $i \in [0, 1]$. A reoptimizing firm $i \in [0, 1]$ maximizes the current value of profit over the time when \tilde{P}_t remains effective,

$$\max_{\tilde{P}_t} \sum_{j=0}^{\infty} \beta^j \theta^j E_t \left\{ \Lambda_{t+j} \left[\tilde{P}_t Y_{t+j}(i) - P_{t+j} \text{mc}_{t+j} Y_{t+j}(i) \right] \right\} \quad (23)$$

$$\text{s.t. } Y_t(i) = Y_t \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon}, \quad (24)$$

where (24) is the demand for an intermediate good i (follows from the first-order condition of (18), (19)); Λ_{t+j} is the Lagrange multiplier on the household's budget constraint (14); mc_{t+j} is the real marginal cost of output at time $t+j$ (which is identical across the firms).

Government. Government finances a stochastic stream of public consumption by levying lump-sum taxes and by issuing nominal debt. The government budget constraint is

$$T_t + \frac{B_t}{\exp(\eta_{B,t}) R_t} = P_t \frac{\bar{G} Y_t}{\exp(\eta_{G,t})} + B_{t-1} + v W_t L_t, \quad (25)$$

where \bar{G} is the steady-state share of government spending in output; $v W_t L_t$ is the subsidy to the intermediate-good firms; $\eta_{G,t}$ is a government-spending shock,

$$\eta_{G,t+1} = \rho_G \eta_{G,t} + \epsilon_{G,t+1}, \quad \epsilon_{G,t+1} \sim \mathcal{N}(0, \sigma_G^2), \quad (26)$$

where $\rho_G \in (-1, 1)$ and $\sigma_G \geq 0$.

Monetary authority. The monetary authority follows a Taylor rule. When the ZLB is imposed on the net interest rate, this rule is $R_t = \max\{1, \Phi_t\}$ with Φ_t being defined as

$$\Phi_t \equiv R_* \left(\frac{R_{t-1}}{R_*} \right)^\mu \left[\left(\frac{\pi_t}{\pi_*} \right)^{\phi_\pi} \left(\frac{Y_t}{Y_{N,t}} \right)^{\phi_y} \right]^{1-\mu} \exp(\eta_{R,t}), \quad (27)$$

where R_t and R_* are the gross nominal interest rate at t and its long-run value, respectively; π_* is the target inflation; $Y_{N,t}$ is the natural level of output; and $\eta_{R,t}$ is a monetary shock,

$$\eta_{R,t+1} = \rho_R \eta_{R,t} + \epsilon_{R,t+1}, \quad \epsilon_{R,t+1} \sim \mathcal{N}(0, \sigma_R^2), \quad (28)$$

where $\rho_R \in (-1, 1)$ and $\sigma_R \geq 0$. . When the ZLB is not imposed, the Taylor rule is $R_t = \Phi_t$.

Natural level of output. The natural level of output $Y_{N,t}$ is the level of output in an otherwise identical economy but without distortions. It is a solution to the following planner's problem

$$\max_{\{C_t, L_t\}_{t=0, \dots, \infty}} E_0 \sum_{t=0}^{\infty} \beta^t \exp(\eta_{u,t}) \left[\frac{C_t^{1-\gamma} - 1}{1-\gamma} - \exp(\eta_{L,t}) \frac{L_t^{1+\vartheta} - 1}{1+\vartheta} \right] \quad (29)$$

$$\text{s.t. } C_t = \exp(\eta_{a,t}) L_t - G_t, \quad (30)$$

where $G_t \equiv \frac{\bar{G} Y_t}{\exp(\eta_{G,t})}$ is given, and $\eta_{u,t+1}$, $\eta_{L,t+1}$, $\eta_{a,t+1}$, and $\eta_{G,t}$ follow the processes (15), (16), (22), and (26), respectively. The FOCs of the problem (29), (30) imply that $Y_{N,t}$ depends only on exogenous shocks,

$$Y_{N,t} = \left[\frac{\exp(\eta_{a,t})^{1+\vartheta}}{[\exp(\eta_{G,t})]^{-\gamma} \exp(\eta_{L,t})} \right]^{\frac{1}{\vartheta+\gamma}}. \quad (31)$$

5.2 Summary of equilibrium conditions

We summarize the equilibrium conditions below (the derivation of the first-order conditions is provided in Appendix E):

$$S_t = \frac{\exp(\eta_{u,t} + \eta_{L,t})}{\exp(\eta_{a,t})} L_t^\vartheta Y_t + \beta \theta E_t \{ \pi_{t+1}^\varepsilon S_{t+1} \}, \quad (32)$$

$$F_t = \exp(\eta_{u,t}) C_t^{-\gamma} Y_t + \beta \theta E_t \{ \pi_{t+1}^{\varepsilon-1} F_{t+1} \}, \quad (33)$$

$$C_t^{-\gamma} = \frac{\beta \exp(\eta_{B,t}) R_t}{\exp(\eta_{u,t})} E_t \left[\frac{C_{t+1}^{-\gamma} \exp(\eta_{u,t+1})}{\pi_{t+1}} \right], \quad (34)$$

$$\frac{S_t}{F_t} = \left[\frac{1 - \theta \pi_t^{\varepsilon-1}}{1 - \theta} \right]^{\frac{1}{1-\varepsilon}}, \quad (35)$$

$$\Delta_t = \left[(1 - \theta) \left[\frac{1 - \theta \pi_t^{\varepsilon-1}}{1 - \theta} \right]^{\frac{\varepsilon}{\varepsilon-1}} + \theta \frac{\pi_t^\varepsilon}{\Delta_{t-1}} \right]^{-1}, \quad (36)$$

$$Y_t = \exp(\eta_{a,t}) L_t \Delta_t, \quad (37)$$

$$C_t = \left(1 - \frac{\bar{G}}{\exp(\eta_{G,t})} \right) Y_t, \quad (38)$$

$$R_t = \max \{ 1, \Phi_t \}, \quad (39)$$

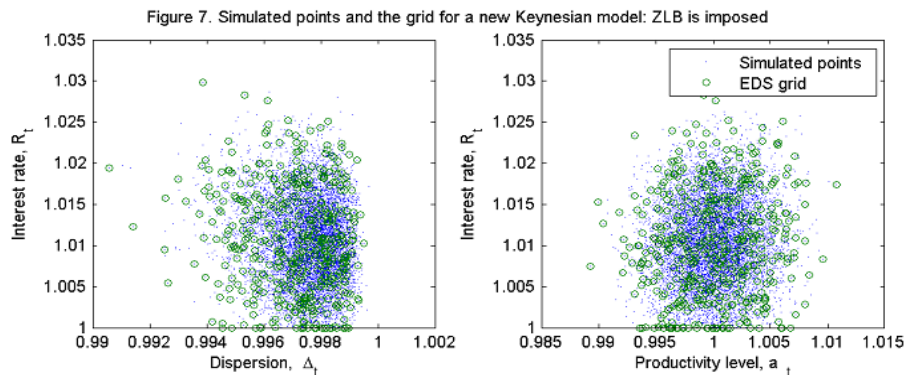
where Φ_t is given by (27); the variables S_t and F_t are introduced for a compact representation of the profit-maximization condition of the intermediate-good firm and are defined recursively; $\pi_{t+1} \equiv \frac{P_{t+1}}{P_t}$ is the gross inflation rate between t and $t+1$; Δ_t is a measure of price dispersion across firms (also referred to as efficiency distortion). The conditions (32)–(38) correspond, respectively, to (E17), (E18), (E23), (E33) and (E3) in Appendix E.

An interior equilibrium is described by 8 equilibrium conditions (32)–(39), and 6 processes for exogenous shocks, (15)–(17), (22), (28) and (26). The system of equations must be solved with respect to 8 unknowns $\{C_t, Y_t, L_t, \pi_t, \Delta_t, R_t, S_t, F_t\}$. There are 2 endogenous and 6 exogenous state variables, $\{\Delta_{t-1}, R_{t-1}\}$, and $\{\eta_{u,t}, \eta_{L,t}, \eta_{B,t}, \eta_{a,t}, \eta_{R,t}, \eta_{G,t}\}$, respectively.

5.3 Numerical analysis

Methodology. We use the estimates of Smets and Wouters (2003, 2007) and Del Negro et al. (2007) to assign values to the parameters. We approximate the equilibrium rules

$S_t = S(x_t)$, $F_t = F(x_t)$ and $C_t^{-\gamma} = MU(x_t)$ with $x_t = \{\Delta_{t-1}, R_{t-1}, \eta_{u,t}, \eta_{L,t}, \eta_{B,t}, \eta_{a,t}, \eta_{R,t}, \eta_{G,t}\}$ using the Euler equations (32), (33) and (34), respectively. We solve for the other variables analytically using the remaining equilibrium conditions. We compute the polynomial solutions of degrees 2 and 3, referred to as EDS2 and EDS3, respectively. For comparison, we also compute first- and second-order perturbation solutions, referred to as PER1 and PER2, respectively (we use Dynare 4.2.1 software). When solving the model with the ZLB by the EDS algorithm, we impose the ZLB both in the solution procedure and in subsequent simulations (accuracy checks). Perturbation methods do not allow us to impose the ZLB in the solution procedure. The conventional approach in the literature is to disregard the ZLB when computing perturbation solutions and to impose the ZLB in simulations when running accuracy checks (that is, whenever R_t happens to be smaller than 1 in simulation, we set it to 1). A detailed description of the methodology of our numerical analysis is provided in Appendix E. We illustrate the EDS grid for the model with the ZLB in Figure 7 where we plot the time-series solution and the grids in two-dimensional spaces, namely, (R_t, Δ_t) and $(R_t, \exp(\eta_{a,t}))$. We see that many points happen to be on the bound $R_t = 1$ and that the essentially ergodic set in the two figures is shaped roughly as a circle.



Accuracy and cost of solutions. Two parameters that play a special role in our analysis are the volatility of labor shocks σ_L and the target inflation rate π_* . Concerning σ_L , Del Negro et al. (2007) finds that shocks to labor must be as large as 40% to match the data, namely, they estimate the interval $\sigma_L \in [0.1821, 0.6408]$ with the average of $\sigma_L = 0.4054$. Concerning π_* , Del Negro et al. (2007) estimate the interval $\pi_* \in [1.0461, 1.0738]$ with the average of $\pi_* = 1.0598$, while Smets and Wouters (2003) use the value $\pi_* = 1$. The inflation rate affects the incidence of the ZLB: a negative net nominal interest rate is more likely to occur in a low- than in a high-inflation economy.¹⁴ In Table 4, we show how the parameters σ_L and π_* affect the quality of numerical solutions.

¹⁴Chung, Laforte, Reifschneider, and Williams (2011) provide estimates of the incidence of the ZLB in the US economy. Christiano, Eichenbaum and Rebelo (2009) study the economic significance of the ZLB in the context of a similar model. Also, Mertens and Ravn (2011) analyze the incidence of the ZLB in a model with sunspot equilibria.

Table 4: Accuracy and speed of the EDS algorithm in the multi-agent model.^a

Polyn. degree	$N = 2$			$N = 20$			$N = 40$			$N = 200$		
	Monom. $2N$ nodes			Monom. $2N$ nodes			GH 1-node			GH 1-node		
	$\bar{M} = 300$			$\bar{M} = 1000$			$\bar{M} = 4000$			$\bar{M} = 1000$		
	L_1	L_∞	CPU	L_1	L_∞	CPU	L_1	L_∞	CPU	L_1	L_∞	CPU
1st	-4.70	-3.17	0.7	-4.76	-3.05	21	-4.79	-3.09	89	-4.66	-2.90	105
2nd	-6.01	-4.06	1.9	-5.88	-4.14	282	-5.48	-4.13	1463			

^a Notes: L_1 and L_∞ are, respectively, the average and maximum of absolute residuals across optimality condition and test points (in log10 units) on a stochastic simulation of 10,000 observations; CPU is the time necessary for computing a solution (in minutes); \bar{M} is the target number of points in the EDS grid, respectively; "Monom. $2N$ " and "GH 1-node" are the monomial rule with two nodes and one-node integration rule, respectively.

In the first experiment, we neglected the ZLB: the goal of this experiment is to show that net interest rates will be occasionally negative if ZLB is not imposed. We assume $\sigma_L = 0.1821$ (which is the lower bound of the interval estimated by Smets and Wouters, 2003), set $\pi_* = 1$ and allow for a negative net interest rate. Both the perturbation and EDS methods deliver reasonably accurate solutions. The maximum size of residuals in the equilibrium conditions is about 6% and 2% for PER1 and PER2, respectively ($10^{-1.21}$ and $10^{-1.64}$ in the table), and it is less than 1% and 0.2% for EDS2 and EDS2, respectively ($10^{-2.02}$ and $10^{-2.73}$ in the table). We also report the minimum and maximum values of R_t on a stochastic simulation, as well as a percentage number of periods in which $R_t < 1$. Here, R_t falls to 0.9916, and the frequency of $R_t < 1$ is about 2%.

We design the next two experiments to separate the effect of the volatility of labor shocks σ_L and the inflation rate $\pi_* = 1$ on the quality of numerical solutions. In the second experiment, we consider a higher volatility of labor $\sigma_L = 0.4054$, and we set $\pi_* = 1.0598$, which is sufficient to preclude net nominal interest rates from being negative. The performance of the perturbation methods becomes significantly worse. The residuals in the equilibrium conditions for the PER1 solution are as large as 25% ($10^{-0.59}$), and they are even larger for the PER2 solution, namely, 38% ($10^{-0.42}$). Thus, increasing the order of perturbation does not help us increase the quality of approximation. The accuracy of the EDS solutions also decreases but less dramatically: the corresponding residuals for the EDS2 and EDS3 methods are less than 5% ($10^{-1.31}$) and 1.2% ($10^{-1.91}$), respectively. For the EDS method, high-degree polynomials do help us increase the quality of approximation.

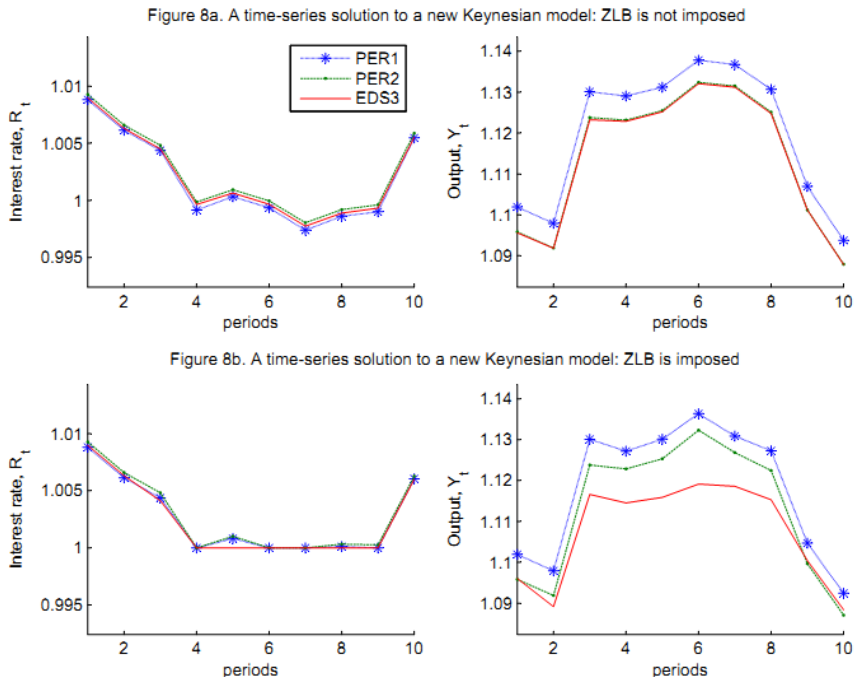
In the third experiment, we concentrate on the effect of the ZLB on equilibrium by setting $\pi_* = 1$ and by imposing the restriction $R_t \geq 1$ under the low-volatility $\sigma_L = 0.1821$ of labor shocks assumed in the first experiment.¹⁵ Again, we observe that the accuracy of the perturbation solutions decreases more than the accuracy of the global EDS solutions. In particular, the maximum residual for the PER2 solution is about 5%, while the corresponding residuals for the EDS2 and EDS3 solutions are less than 2.7% ($10^{-1.58}$) and 1.6% ($10^{-1.81}$), respectively.

To appreciate how much the equilibrium quantities differ across the methods, we report the maximum percentage differences between the variables produced by EDS3 and the other methods. The regularities are similar to those we observed for the residuals. First, the difference between the series produced by PER1 and EDS3 can be as large as 17%. Second, the difference between the series produced by PER2 and EDS3 depends on the model: it is about 1% when the ZLB is not imposed in the model with $\sigma_L = 0.1821$ but can reach 10% when the ZLB is imposed. Finally, the difference between the series produced by EDS2 and EDS3 is smaller in all cases (5.04% at most for the model with the imposed ZLB). Generally, the supplementary variables S_t and F_t differ more across

¹⁵In a model with an active ZLB, we find the convergence to be slow and fragile. This fact is possibly related to the finding of Benhabib et al. (2001) that a deterministic version of the model has multiple trajectories converging to a liquidity trap, in addition to a locally-unique equilibrium that converges to a target inflation level.

methods than economically relevant variables such as Y_t , L_t and C_t .

Economic importance of the ZLB. Figures 8a and 8b plot fragments of a stochastic simulation when the ZLB is not imposed and imposed, respectively, for the model parameterized by $\sigma_L = 0.1821$ and $\pi_* = 1$. When the ZLB is not imposed, both the perturbation and EDS methods predict 5 periods of negative (net) interest rates (see periods 4, 6-9 in Figure 8a). When the ZLB is imposed, the EDS methods, EDS2 and EDS3, predict a zero interest rate in those 5 periods, while the perturbation methods, PER1 and PER2, predict a zero interest rate just in 3 periods (see periods 4, 6 and 7 in Figure 8b).



The way we deal with the ZLB in the perturbation solution misleads the agents about the true state of the economy. To be specific, when we chop off the interest rate at zero in the simulation procedure, agents perceive the drop in the interest rate as being small and respond by an immediate recovery. In contrast, under the EDS algorithm, agents accurately perceive the drop in the interest rate as being large and respond by 5 periods of a zero net interest rate (which correspond to 5 periods of negative net interest rates predicted in the case when the ZLB is not imposed). The output differences between PER2 and EDS3 are relatively small when the ZLB is not imposed but they become quantitatively important when the ZLB is imposed and can be as large as 2%.

Piecewise local basis functions and locally adaptive EDS grids Experiments 1 and 3 in Table 4 showed that imposing the ZLB on the nominal interest rate reduces the accuracy of solutions produced by the EDS method. Increasing the degree of an

approximating polynomial function does not increase much the accuracy. In particular, the maximum residuals of the second- and third-order solutions are $10^{-2.02}$ and $10^{-2.73}$ in the model with an inactive ZLB, while these residuals are respectively $10^{-1.58}$ and $10^{-1.81}$ in the model with an active ZLB. The residuals are large in the model with an active ZLB for two reasons: first, a global polynomial function is not sufficiently flexible to accurately approximate the kink in the ZLB area and second, an evenly-spaced EDS grid does not produce sufficiently many grid points in the ZLB area. Below, we perform two additional experiments in which we analyze how the degree of flexibility of an approximating function and the specific placement of grid points affect the accuracy of solutions under the EDS method.

In the first experiment, we replace a global polynomial approximating function with piecewise local bases as described in Section 2.6. Namely, we first construct an EDS grid with M grid points, $x_1^\varepsilon, \dots, x_M^\varepsilon$, using a constant value of ε ; we then construct an 8-dimensional hypercube with the side of $10^{-3}\varepsilon$ that surrounds each grid point x_i^ε , and we populate it with a 100 low-discrepancy, Sobol, points; we then solve the model on each of the M hypercubes constructed. To simulate the model, we compute a distance from the current state x_i to all grid points, and we adopt a solution from the closest grid point - the nearest neighbor approach. In all experiments, we use second-degree ordinary polynomials as local bases. In the first row of Table 5 (see the experiment "LB"), we show the results depending on the number of local approximations used. As is seen from the

Table 5: Accuracy and speed of the EDS algorithm with local bases and locally adaptive grid in the new Keynesian model.^a

	$\bar{M}=2$			$\bar{M}=10$			$\bar{M}=100$			$\bar{M}=500$		
	L ₁	L _∞	CPU	L ₁	L _∞	CPU	L ₁	L _∞	CPU	L ₁	L _∞	CPU
LB	-3.80	-1.66	4.2	-4.15	-1.83	6.4	-4.54	-2.29	29.3	-4.85	-2.35	131
LB-LA	-4.04	-1.75	4.7	-4.14	-1.93	6.6	-4.53	-2.21	34.7	-4.90	-2.25	139

^a Notes: L₁ and L_∞ are, respectively, the average and maximum of absolute residuals across optimality condition and test points (in log10 units) on a stochastic simulation of 10,000 observations; CPU is the time necessary for computing a solution (in minutes); \bar{M} is the target number of points in the EDS grid, respectively; "LB" and "LB-LA" denote the EDS method with local bases and the EDS method with both local bases and locally-adaptive grid, respectively.

table, the accuracy of solutions increases with the number of local approximations used although the improvements become smaller as the number of approximations increases. In particular, we are able to reduce the maximum residuals to the order of $10^{-2.35}$ when constructing 500 local approximations. The running time for this experiment is about 2 hours.

In our second experiment, we study a variant of the EDS method in which we combine locally-adaptive grid points with local bases. There are many ways to construct an inverse relation between the size of the residuals and the value of ε ; the specific construction of this

relation may considerably affect the accuracy results. In our example, we use the following procedure: In each simulated point x_i , we compute $\varepsilon_i = |\ln \mathcal{R}(x_i)| \bar{\varepsilon}$, where $\mathcal{R}(x_i)$ is the maximum residual across all model's equations, and $\bar{\varepsilon}$ is a normalizing parameter. For example, if for some points x' and x'' , the residuals are $\ln \mathcal{R}(x') = -2$ and $\ln \mathcal{R}(x'') = -6$, then the corresponding value of ε differs by a factor of 3, i.e., $\varepsilon' = 2\bar{\varepsilon}$ and $\varepsilon'' = 6\bar{\varepsilon}$. We compute the normalizing parameter $\bar{\varepsilon}$ using our bisection procedure to obtain the target number of grid points, \bar{M} .

The results about the EDS method with both local bases and locally adaptive grid are shown in the second row of Table 5 (see the experiment "LB-LA"). We observe that the EDS method with locally adaptive grid points, LB-LA, has large accuracy improvements than the EDS method with uniformly spaced grid points, LB, in the first row of Table 5. In particular, when $\bar{M} = 2$, the average residuals for the two methods, LB-LA and LB, are $10^{-4.04}$ and $10^{-3.80}$, respectively, and the maximum residuals are $10^{-1.75}$ and $10^{-1.66}$, respectively. However, as the number of grid points increases, both methods arrive to similar accuracy levels. In fact, for large values of \bar{M} , the EDS method with an evenly-spaced grid, LB, slightly overperforms the locally-adaptive EDS method, LB-LA. Indeed, when we have to select just two areas, the adaptive way of selecting these areas is relatively more important than if we have to select 500 areas.

We shall finally discuss the relation between our locally-adaptive EDS solution method and those studied in Aruoba and Schorfheide (2013). To increase the accuracy of solutions in the ZLB area, Aruoba and Schorfheide (2013) implemented two modifications to the baseline cluster-grid algorithm studied in Judd et al. (2010, 2011b): first, they add grid points near the ZLB area using the actual data on the US economy; and second, they use two piecewise local bases to separately approximate the solution in the areas with active and non-active ZLB. For this special case, the locally-adaptive EDS technique with $\bar{M} = 2$ delivers a similar type of analysis if one of the two EDS grid points is selected in a neighborhood of the ZLB area, namely, we also approximate the solution on two disjoint sets of 100 Sobol points that represent the areas with an active and inactive ZLB using two piecewise local bases. The key novelty of the locally-adaptive EDS method is that it automates the construction of adaptive grid points and local basis functions in a general case of \bar{M} piecewise local approximations.

Lessons. The studied new Keynesian model is a challenging application for any numerical method. First, the dimensionality of the state space is large; second, the volatility of exogenous variables is high; and finally, there is a kink in the equilibrium rules due to the ZLB. We choose this application in order to subject the EDS method to a tight test that makes it possible to see its limitations.

Our results indicate that the EDS method is able to confront the above challenges. First, the running time for the EDS method ranges from 4 to 25 minutes; the EDS method would be tractable in much larger applications, as our results for the multi-country model suggest. Second, the EDS method produces very accurate solutions if the volatility of shocks is not excessively high, and its accuracy can be increased using

polynomial functions of higher degrees or local bases functions, unlike the accuracy of the perturbation methods. Finally, in the presence of the ZLB, the perturbation and EDS methods may produce qualitatively different results. The accuracy of the EDS projection algorithm can be increased by adapting the density of grid points to a given application and by using more flexible functional families that can accommodate kinks and strong non-linearities. This increases the computational cost but the studied EDS methods are naturally parallelizable, and the cost can be reduced.

6 Conclusion

We introduce a projection algorithm that operates on a high-probability area of the ergodic set of an economic model. The EDS algorithm is tractable in problems with much higher dimensionality than those studied in the related literature. In particular, we are able to compute accurate quadratic solutions to a multicountry growth model with up to 80 state variables. Furthermore, we are able to compute an accurate global solution to a new Keynesian model. This model is of particular interest to the literature as it is used by governments and financial institutions all over the world for the policy analysis. We find that perturbation methods are not reliable in the context of new Keynesian models, and we show examples where perturbation and global solution methods produce qualitatively different predictions. We emphasize that all the numerical results in the paper are obtained using a standard desktop computer and serial MATLAB software. The speed and accuracy of the EDS algorithm can be increased by far using more powerful hardware and software, as well as parallelization techniques.

Acknowledgements

This paper is a substantially revised version of NBER working papers 15965 and 18501 entitled "A Cluster-Grid Projection Method: Solving Problems with High Dimensionality" and "Merging Simulation and Projection Approaches to Solve High-Dimensional Problems", respectively. This paper was written jointly with Kenneth L. Judd who decided to withdraw as an author. In our view, he remains to be a co-author of the present paper. We thank Editor of Quantitative Economics Frank Schorfheide for many thoughtful comments and suggestions that led to a substantial improvement of the paper. We are also indebted to Editor of Econometrica Jean-Marc Robin and five anonymous referees for a large number of useful comments and suggestions in three rounds of the revision. Also, we benefited from the comments of seminar participants at University of London, Canadian Central Bank, Cornell University, Federal Reserve Bank of San Francisco, Penn State University, Santa Clara University, Stanford University, University Autònoma de Barcelona, University of Alicante, University of Bilbao, University of California at Berkeley, University of Chicago, University of Edinburgh, and University of Oxford, as well as from participants of the 2011 Society of Economic Dynamics meeting, the 2011, 2012 and

2013 meetings of Society of Computational Economics, the 2013 meeting of Macro Financial Modeling group at the University of Chicago, the Summer 2013 Computation in CA Workshop at Stanford University and the 2014 American Economic Association meeting in Philadelphia. Errors are ours. Lilia Maliar and Serguei Maliar acknowledge support from the Hoover Institution and Department of Economics at Stanford University, University of Alicante, Ivie, MECD and FEDER funds under the projects SEJ-2007-62656 and ECO2012-36719.

References

- [1] Adam, K. and R. Billi, (2006). Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit, and Banking* 38 (7), 1877-1905.
- [2] Adda, J. and R. Cooper, (2003). *Dynamic Economics: Quantitative Methods and Applications*. London, England: The MIT Press, Cambridge Massachusetts.
- [3] Adjemian, S., H. Bastani, M. Juillard, F. Mihoubi, G. Perendia, M. Ratto and S. Villemot, (2011). *Dynare: Reference Manual, Version 4*. Dynare Working Papers 1, CEPREMAP.
- [4] Adjemian, S. and M. Juillard, (2011). Accuracy of the extended path simulation method in a new Keynesian model with zero lower bound on the nominal interest rate. Manuscript.
- [5] Anderson, G., J. Kim and T. Yun, (2010). Using a projection method to analyze inflation bias in a micro-founded model. *Journal of Economic Dynamics and Control* 34 (9), 1572-1581.
- [6] Aruoba, S., J. Fernández-Villaverde and J. Rubio-Ramírez, (2006). Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control* 30, 2477-2508.
- [7] Aruoba, S. and F. Schorfheide, (2013). Macroeconomic dynamics near ZLB: a tale of two equilibria. NBER working paper 19248.
- [8] Baryshnikov, Yu., Eichelbacker, P., Schreiber, T., and J.E. Yukich, (2008). Moderate deviations for some point measures in geometric probability. *Annales de l'Institut Henri Poincaré - Probabilités et Statistiques* 44, 422-446.
- [9] Benhabib, J., S. Schmitt-Grohé and M. Uribe, (2001a). Monetary policy and multiple equilibria. *American Economic Review* 91(1), 167-186.
- [10] Benkard, Weintraub and van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76(6), 1375-1411.
- [11] Bertsekas, D. and J. Tsitsiklis (1996). *Neuro-Dynamic Programming*. Optimization and Neural Computation Series. Athena Scientific: Belmont, Massachusetts.
- [12] Braun, R.A., L.M.Körber, and Y. Waki (2012). Some unpleasant properties of log-linearized solutions when the nominal rate is zero. Federal Reserve Bank of Atlanta Working Paper Series.
- [13] Brumm, J. and S. Scheidegger, (2013). Using adaptive sparse grids to solve high-dimensional dynamic models. Manuscript. University of Zurich.

- [14] Christiano, L., M. Eichenbaum and C. Evans, (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113/1, 1-45.
- [15] Christiano, L., M. Eichenbaum, and S. Rebelo, (2011). When is the government spending multiplier large? *Journal of Political Economy* 119(1), 78-121.
- [16] Christiano, L. and D. Fisher, (2000). Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control* 24, 1179-1232.
- [17] Chung, H., J.-P. Laforte, D. Reifschneider and J. Williams, (2011). Have we underestimated the probability of hitting the zero lower bound? Federal Reserve Bank of San Francisco. Working paper 2011-01.
- [18] Collard, F., and M. Juillard, (2001). Accuracy of stochastic perturbation methods: the case of asset pricing models, *Journal of Economic Dynamics and Control*, 25, 979-999.
- [19] Coibion, O., Y. Gorodnichenko and J. Wieland, (2012), The optimal inflation rate in new Keynesian models: should central banks raise their inflation targets in light of the zero lower bound. Manuscript.
- [20] Del Negro, M., F. Schorfheide, F. Smets and R. Wouters, (2007). On the fit of new Keynesian models. *Journal of Business and Economic Statistics* 25 (2), 123-143.
- [21] Den Haan, W. (1990). The optimal inflation path in a Sidrauski-type model with uncertainty. *Journal of Monetary Economics* 25, 389-409.
- [22] Den Haan, W. and A. Marcet, (1990). Solving the stochastic growth model by parameterized expectations. *Journal of Business and Economic Statistics* 8, 31-34.
- [23] Den Haan, W., and A. Marcet (1994). Accuracy in Simulations. *Review of Economic Studies* 61, 3–18.
- [24] Den Haan, W., (2010), Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control* 34, 4–27.
- [25] Everitt, B., S. Landau, M. Leese and D. Stahl, (2011). *Cluster Analysis*. Wiley Series in Probability and Statistics. Wiley: Chichester, United Kingdom.
- [26] Fair, R. and J. Taylor, (1983). Solution and maximum likelihood estimation of dynamic non-linear rational expectation models. *Econometrica* 51, 1169-1185.
- [27] Feng, Z., J. Miao, A. Peralta-Alva, and M. Santos, (2009). Numerical simulation of nonoptimal dynamic equilibrium models. Working papers Federal Reserve Bank of St. Louis 018.

- [28] Fernández-Villaverde, J., G. Gordon, P. Guerrón-Quintana, and J. Rubio-Ramírez, (2012). Nonlinear adventures at the zero lower bound. NBER working paper 18058.
- [29] Fernández-Villaverde, J. and J. Rubio-Ramírez, (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies* 74, 1059-1087.
- [30] Fudenberg D. and D. Levine, (1993). Self-confirming equilibrium. *Econometrica* 61, 523-545.
- [31] Galí, J., (2008). *Monetary Policy, Inflation and the Business Cycles: An Introduction to the New Keynesian Framework*. Princeton University Press: Princeton, New Jersey.
- [32] Gaspar, J. and K. Judd, (1997). Solving large-scale rational-expectations models. *Macroeconomic Dynamics* 1, 45-75.
- [33] Gavion, W.T., B.D.Keen, A.W.Richter, and N.A. Throckmorton, (2013). Global dynamics at the zero lower bound, Federal Reserve Board St. Louis working paper 2013-007B.
- [34] Gust, C., D. Lopez-Salido and M.E. Smith, (2012). The empirical implications of the interest-rate lower bound. Manuscript, Federal Reserve Board.
- [35] Hughes, T., (1987). *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice Hall, Englewood Cliff, NJ.
- [36] Judd, K., (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory* 58, 410-452.
- [37] Judd, K., (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- [38] Judd, K. and S. Guu, (1993). Perturbation solution methods for economic growth models, in: H. Varian, (Eds.), *Economic and Financial Modeling with Mathematica*, Springer Verlag, pp. 80-103.
- [39] Judd, K., L. Maliar and S. Maliar, (2010). A cluster grid projection method: solving problems with high dimensionality. NBER working paper 15965.
- [40] Judd, K., L. Maliar and S. Maliar, (2011a). Numerically stable and accurate stochastic simulation approaches for solving dynamic models. *Quantitative Economics* 2, 173-210.
- [41] Judd, K., L. Maliar and S. Maliar, (2011b). A cluster grid projection algorithm: solving problems with high dimensionality. Manuscript available at <http://stanford.edu/~maliarl>.

- [42] Judd, K., L. Maliar and S. Maliar, (2012). Merging simulation and projection approaches to solve high-dimensional problems. NBER working paper 18501.
- [43] Judd, K., L. Maliar, S. Maliar and R. Valero, (2013). Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. NBER working paper 19326.
- [44] Juillard, M. and S. Villemot, (2011). Multi-country real business cycle models: Accuracy tests and testing bench. *Journal of Economic Dynamics and Control* 35, 178–185.
- [45] Kiefer, J., (1961). On large deviations of the empiric D.F. of vector change variables and a law of the iterated logarithm. *Pacific Journal of Mathematics* 11, 649-660.
- [46] Kollmann, R., (2002). Monetary policy rules in the open economy: effects on welfare and business cycles. *Journal of Monetary Economics* 49, 989-1015.
- [47] Kollmann, R., S. Kim and J. Kim, (2011a). Solving the multi-country real business cycle model using a perturbation method. *Journal of Economic Dynamics and Control* 35, 203-206.
- [48] Kollmann, R., S. Maliar, B. Malin and P. Pichler, (2011b). Comparison of solutions to the multi-country real business cycle model. *Journal of Economic Dynamics and Control* 35, 186-202.
- [49] Krueger, D. and F. Kubler, (2004). Computing equilibrium in OLG models with production. *Journal of Economic Dynamics and Control* 28, 1411-1436.
- [50] Ma, X. and N. Zabararas, (2009). An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations, *Journal of Computational Physics* 228, 3084–3113.
- [51] Maliar, L. and S. Maliar, (2005). Solving non-linear stochastic growth models: iterating on value function by simulations. *Economics Letters* 87, 135-140.
- [52] Maliar, L. and S. Maliar, (2013). Numerical methods for large scale dynamic economic models. in: K. Schmedders and K. Judd (Eds.), *Handbook of Computational Economics*, Amsterdam: Elsevier Science, forthcoming.
- [53] Maliar, S., L. Maliar and K. Judd, (2011). Solving the multi-country real business cycle model using ergodic set methods. *Journal of Economic Dynamic and Control* 35, 207–228.
- [54] Malin, B., D. Krueger and F. Kubler, (2011). Solving the multi-country real business cycle model using a Smolyak-collocation method. *Journal of Economic Dynamics and Control* 35, 229-239.

- [55] Marcet, A. (1988), Solving non-linear models by parameterizing expectations. Unpublished manuscript, Carnegie Mellon University, Graduate School of Industrial Administration.
- [56] Marcet, A., and G. Lorenzoni, (1999). The parameterized expectation approach: some practical issues. in: R. Marimon and A. Scott (Eds.), *Computational Methods for Study of Dynamic Economies*, Oxford University Press, New York, pp. 143-171.
- [57] Marcet, A. and T. Sargent, (1989). Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy* 97, 1306-1322.
- [58] Marimon, R. and A. Scott, (1999). *Computational Methods for Study of Dynamic Economies*, Oxford University Press, New York.
- [59] McGrattan, E., (1996). Solving the stochastic growth model with a finite element method. *Journal of Economic Dynamics and Control* 20, 19-42.
- [60] Mertens, K. and M. Ravn, (2011). Credit channels in a liquidity trap. CEPR discussion paper 8322.
- [61] Mertens, K. and M. Ravn, (2013). Fiscal policy in an expectations driven liquidity trap. Manuscript, Cornell University.
- [62] Niederreiter, H., (1992). *Random Number Generation and Quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.
- [63] Pakes, A. and P. McGuire, (2001). Stochastic algorithms, symmetric Markov perfect equilibria, and the 'curse' of dimensionality. *Econometrica* 69, 1261-1281.
- [64] Peralta-Alva, A. and M. Santos, (2005). Accuracy of simulations for stochastic dynamic models. *Econometrica* 73, 1939-1976.
- [65] Pichler, P., (2011). Solving the multi-country real business cycle model using a monomial rule Galerkin method. *Journal of Economic Dynamics and Control* 35, 240-251.
- [66] Powell W., (2011). *Approximate Dynamic Programming*. Wiley: Hoboken, New Jersey.
- [67] Ravenna, F. and C. Walsh, (2011). Welfare-based optimal monetary policy with unemployment and sticky prices: a linear-quadratic framework. *American Economic Journal: Macroeconomics* 3, 130-162.
- [68] Rényi, A., (1958). On a one-dimensional space-filling problem. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences* 3, 109-127.

- [69] Richter, A.W. and N.A. Throckmorton (2013). The zero lower bound: frequency, duration, and determinacy. Manuscript.
- [70] Romesburg, C., (1984). Cluster Analysis for Researchers. Lifetime Learning Publications: Belmont, California.
- [71] Rudebusch, G. and E. Swanson, (2008). Examining the bond premium puzzle with a DSGE model. *Journal of Monetary Economics* 55, S111-S126.
- [72] Rust, J., (1996). Numerical dynamic programming in economics, in: H. Amman, D. Kendrick and J. Rust (Eds.), *Handbook of Computational Economics*, Amsterdam: Elsevier Science, 619-722.
- [73] Rust, J., (1997). Using randomization to break the curse of dimensionality. *Econometrica* 65, 487-516.
- [74] Santos, M., (1999). Numerical solution of dynamic economic models, in: J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Amsterdam: Elsevier Science, pp. 312-382.
- [75] Santos, M., (2000). Accuracy of numerical solutions using the Euler equation residuals. *Econometrica* 68, 1377-1402.
- [76] Santos, M., (2002). On non-existence of Markov equilibria in competitive market economies. *Journal of Economic Theory* 105, 73-98.
- [77] Benhabib, J., Schmitt-Grohé, S. and M Uribe (2001). The perils of Taylor rules. *Journal of Economic Theory* 96, 40-69.
- [78] Schmitt-Grohé S. and M. Uribe, (2007). Optimal simple and implementable monetary fiscal rules. *Journal of Monetary Economics* 54, 1702-1725.
- [79] Schmitt-Grohé, S. and M. Uribe, (2012). The making of a great contraction with a liquidity trap and a jobless recovery. Manuscript, Columbia University.
- [80] Scott, D. and S. Sain, (2005). Multidimensional density estimation, in: C. Rao, E. Wegman and J. Solka (Eds.), *Handbook of Statistics*, vol. 24, Amsterdam: Elsevier B. V, pp. 229-261.
- [81] Smets, F. and R. Wouters, (2003). An estimated dynamic stochastic general equilibrium model of the Euro area. *Journal of the European Economic Association* 1(5), 1123-1175.
- [82] Smets, F. and R. Wouters, (2007). Shocks and frictions in US business cycles: a Bayesian DSGE approach. *American Economic Review* 97 (3), 586-606.

- [83] Smith, A., (1993). Estimating non-linear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics* 8, S63-S84.
- [84] Stachursky, J., (2009). *Economic Dynamics: Theory and Computations*. Cambridge, MA: MIT Press.
- [85] Stokey, N. L. and R. E. Lucas Jr. with E. Prescott, (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.
- [86] Tauchen, G. and R. Hussey, (1991). Quadrature-based methods for obtaining approximate solutions to non-linear asset pricing models. *Econometrica* 59, 371-396.
- [87] Taylor, J. and H. Uhlig, (1990). Solving non-linear stochastic growth models: a comparison of alternative solution methods. *Journal of Business and Economic Statistics* 8, 1-17.
- [88] Taylor, J., (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39, 195-214.
- [89] Temlyakov, V., (2011). *Greedy approximation*. Cambridge University Press, Cambridge.
- [90] Tsitsiklis, J., (1994). Asynchronous stochastic approximation and Q-Learning, *Machine Learning* 16, 185-202.
- [91] Winschel, V. and M. Krätzig, (2010). Solving, estimating and selecting non-linear dynamic models without the curse of dimensionality. *Econometrica* 78(2), 803-821.
- [92] Wright, B. and J. Williams, (1984). The welfare effects of the introduction of storage. *Quarterly Journal of Economics* 99, 169-192.
- [93] Woodford, M. (2003). *Interest and prices*. Princeton University Press.
- [94] Zhao, O. and M. Woodroffe, (2008). Law of the iterated logarithm for stationary processes. *The Annals of Probability* 36 (1), 127–142.