Perfect Simulation of Stationary Equilibria[☆]

Kazuo Nishimura^a, John Stachurski^{b,∗}

^a*Institute of Economic Research, Kyoto University, Yoshida honmachi, Sakyo-ku, Kyoto 606–8501, Japan*

^b*School of Economics, Australian National University, ACT 0200, Australia*

Abstract

Using a variation of the coupling from the past technique, this paper develops algorithms which generate independent observations from the stationary distributions of various dynamic economic models. These variates can be used for calibration, calculation of steady state phenomena, and simulation-based estimation. As an application, we demonstrate how to generate exact samples from the stationary distribution of an incomplete markets model routinely calibrated by macroeconomists. Our implementation generates 100,000 independent draws from the stationary distribution in less than 3 seconds.
Journal of Economic Literature Classification: C15, C61, C63

Keywords: Stationarity, coupling from the past, perfect sampling

1. Introduction

In economic modeling, dynamic outcomes are often computed using approximation techniques such as linearization. Approximation helps to represent dynamics in a tractable form, but the size of the errors introduced by the approximation are difficult to assess.[1] As a result, efforts have been made to

---

[∗]Corresponding author.
Email addresses: `nishimura@kier.kyoto-u.ac.jp` (Kazuo Nishimura),
`john.stachurski@gmail.com` (John Stachurski)

[1]Tight error bounds require significant smoothness in the laws of motion, which is often unavailable or difficult to prove. In addition, Fernández-Villaverde, Rubio-Ramírez and Santos (2006) show that second order approximation effects in the law of motion have first order effects on likelihood functions, and hence on estimates of parameters.

compute and estimate stochastic dynamic equilibrium models with greater accuracy by employing simulation. Simulation is well suited to uncovering the predictions of nonlinear models and comparing them with data. Simulation-based calibration and estimation are becoming increasingly commonplace.[2]

For stationary models, such simulation typically addresses properties of the stationary distribution and stationary time series. Provided that the model in question is both stationary and ergodic, approximate draws from the stationary distribution can be generated, since the distribution of the time t state converges to the stationary distribution for any choice of initial condition (cf., e.g., Santos and Peralta-Alva, 2005). Simulating until the distribution of the state is approximately stationary is referred to as "burn-in."

However, when the state space of the model is not finite, little guidance is available on the *rate* of convergence to the stationary solution, in which case the length of burn-in required is the subject of guesswork and heuristics. This is particularly problematic when searching over a large parameter space for a good fit to data.[3] A second and related issue is that, regardless of how much burn-in is performed, the resulting sample is never exactly stationary.[4]

In this paper we address these two problems in the context of dynamic stochastic economic models, studying algorithms which permit *exact* sampling from the stationary distribution of the model. The second problem is then solved: one need not consider errors because the sample from the stationary distribution is exact. The first problem is also solved, because the algorithm *terminates automatically* when the stationary distribution is attained.

The algorithms in question are variations of Propp and Wilson's (1996) coupling from the past (CFTP) technique. In essence, one can think of CFTP as permitting observation at time zero of time series that start in the infinite past, in which case the burn-in is infinite and the stationary distribution is attained. Despite the fact that the sequence is infinite, when certain assumptions are satisfied the limit of the process can be observed in *finite* time with probability one.

From a single observation of the stationary distribution, a stationary time series can be generated by simulating the model in the usual way, taking the observation as the initial condition. More importantly, the algorithm can be used to generate *independent* draws from the stationary distribution. Standard IID laws of large numbers and central limit theorems can then be used in Monte Carlo estimation.

---

[2]For discussion of estimation via simulation, see, for example, Lee and Ingram (1991), Duffie and Singleton (1993) or Gouriéroux and Monfort (1996).

[3]For example, burn-in time tends to vary over the parameters, causing biased estimates.

[4]As the size of the error depends on the initial condition, this deviation between computed and stationary distributions is sometimes called "initialization bias."

Propp and Wilson's technique was originally proposed for problems in statistical mechanics involving finite Markov chains. In economic dynamics, most problems of interest have uncountable state spaces. Outside of economics, the CFTP method has recently been extended to some general state settings (see, e.g. Foss and Tweedie, 1998, Murdoch and Green, 1998, Corcoran and Tweedie, 2001, or Athreya and Stenflo, 2003). We use the iterated function and reversed process notation of Athreya and Stenflo (2003). Our theoretical results in section 3 are based on similar ideas, with some additional structure that avoids the need for computing residual kernels and provides a simple exposition.

As an application of the algorithm, we generate exact random draws from a stationary distribution corresponding to cross-sectional household assets in Aiyagari's (1994) well-known model. Our implementation generated 100,000 IID draws in 2.4 seconds. From this sample, we estimated aggregate capital using Monte Carlo. The estimate had a 99% confidence interval with width of less than 0.0006.

The structure of the paper is as follows. Section 2 formulates the problem and describes our assumptions. Theoretical perfect sampling results are presented in section 3. Section 4 gives the numerical application, and section 5 concludes.

2. Formulation of the Problem

We consider a class of dynamic models which evolve on state space $\mathbb{X}$. The time t state is denoted by $X_t \in \mathbb{X}$. Typically $\mathbb{X}$ is a subset of $\mathbb{R}^n$, so that X_t is an n-dimensional vector consisting of all (endogenous and exogenous) state variables. In terms of our theory, we require only that $\mathbb{X}$ is a separable and completely metrizable topological space. The Borel subsets of $\mathbb{X}$ are denoted by $\mathscr{B}$.

All random variables are assumed to be defined on a fixed probability space $(\Omega, \mathscr{F}, \mathbf{P})$. For any random variable Y defined on this space, we use the symbol $\mathcal{L}Y$ to denote the distribution (or law) of Y.[5] As usual, a sequence of probability measures μ_n on $\mathscr{B}$ is said to converge to probability measure μ if $\int h(y)\mu_n(dy) \to \int h(y)\mu(dy)$ for all continuous bounded $h\colon \mathbb{X} \to \mathbb{R}$, and we write $\mu_n \to \mu$.

We consider a generic economic model, the dynamics of which are summarized by a stochastic (transition) kernel Q. Formally, Q is a family of probability measures $\{Q(x, dy)\}_{x \in \mathbb{X}}$ on $\mathscr{B}$ such that $x \mapsto Q(x, B)$ is measurable for each $B \in \mathscr{B}$. The term $Q(x, B)$ represents the probability that the next period state is in B, given that the current state is x.

---

[5]In other words, $\mathcal{L}Y$ is the probability measure $B \mapsto \mathbf{P}\{Y \in B\}$ on $\mathscr{B}$.

Example 2.1. Santos and Peralta-Alva (2005) consider a class of dynamic models which obey the law of motion

$$z_{t+1} = \Psi(z_t, \epsilon_{t+1}), \quad k_{t+1} = g(k_t, z_t, \epsilon'_{t+1})$$

The process $(z_t)_{t \geq 0}$ is a vector of correlated exogenous variables in $\mathbb{R}^j$, while $(k_t)_{t \geq 0}$ is a vector of endogenous variables in $\mathbb{R}^k$, updating according to the equilibrium decision rule g. The processes $(\epsilon_t)_{t \geq 1}$ and $(\epsilon'_t)_{t \geq 1}$ are IID with distributions ν and ν' respectively. For this model, the state space is $\mathbb{X} := \mathbb{R}^j \times \mathbb{R}^k$, and the corresponding stochastic kernel is defined by

$$Q((z,k), B) = \int \int \mathbb{1}_B[\Psi(z,u), g(k,z,u')]\nu(du)\nu'(du')$$

for $(z,k) \in \mathbb{X}$ and $B \in \mathscr{B}$, where $\mathbb{1}_B$ is the indicator function of B.

Returning to the general case, we consider processes which obey the dynamics encapsulated in a given kernel Q. A stochastic process $(X_t)_{t \geq 0}$ on $\mathbb{X}$ is defined to be *Q-Markov* if

$$\mathbf{P}\{X_{t+1} \in B \mid X_t = x\} = Q(x, B) \qquad (x \in \mathbb{X}, B \in \mathscr{B}) \tag{1}$$

The k-step kernel Q^k associated with Q is defined by

$$Q^1 = Q \quad \text{and} \quad Q^k(x, B) = \int Q(x, dz) Q^{k-1}(z, B)$$

If $(X_t)_{t \geq 0}$ is Q-Markov, then, for all $k \in \mathbb{N}$, $x \in \mathbb{X}$ and $B \in \mathscr{B}$, we have $\mathbf{P}\{X_{t+k} \in B \mid X_t = x\} = Q^k(x, B)$. In particular, $Q^k(x, dy)$ is the marginal distribution $\mathcal{L}X_k$ of X_k when $X_0 = x$.

We focus exclusively on stationary processes. A probability measure ψ^* on $(\mathbb{X}, \mathscr{B})$ is called *stationary* for Q if

$$\psi^*(B) = \int Q(x, B) \psi^*(dx) \quad \forall B \in \mathscr{B}$$

Q is said to be *globally stable* if there exists a unique stationary distribution ψ^* in the set of probability measures on $\mathscr{B}$, and, moreover,

$$Q^t(x, dy) \to \psi^* \text{ as } t \to \infty \text{ for all } x \in \mathbb{X} \tag{2}$$

By a *stochastic recursive sequence* (SRS) on $\mathbb{X}$ we mean a model of the form

$$X_{t+1} = F(X_t, U_{t+1}) =:: F_{U_{t+1}}(X_t) \qquad t = 0, 1, \ldots \tag{3}$$

where $(X_t)_{t \geq 0}$ takes values in $\mathbb{X}$ and $(U_t)_{t \geq 1}$ takes values in a measurable space $(\mathbb{U}, \mathscr{U})$.[6] The shock sequence is IID, with $\mathcal{L}U_t = \nu$.[7] The time t realization of

---

[6]The function $\mathbb{X} \times \mathbb{Y} \ni (x, u) \mapsto F_u(x) \in \mathbb{X}$ is assumed to be jointly measurable.

[7]Note that correlated exogenous shocks can be incorporated in this framework. Example 2.1 is a case in point, with $X_t = (k_t, z_t)$ and $U_t = (\epsilon_t, \epsilon'_t)$.

this process can be expressed as

$$X_t = F_{U_t} \circ F_{U_{t-1}} \circ \cdots \circ F_{U_1}(x) \qquad (x \in \mathbb{X}) \qquad (4)$$

where $x = X_0$ is the initial condition for the process. We say that the process (3) *represents* the kernel Q if

$$Q(x, B) = \nu\{u \in \mathbb{U} \mid F_u(x) \in B\} = \mathbf{P}\{F_{U_1}(x) \in B\} \qquad (5)$$

for $x \in \mathbb{X}$ and $B \in \mathscr{B}$. In other words, $F_{U_1}(x)$ is a draw from the distribution $Q(x, dy)$. Should this be the case, then $(X_t)_{t \geq 0}$ in (3) is Q-Markov, and

$$Q^k(x, B) = \mathbf{P}\{F_{U_k} \circ \cdots \circ F_{U_1}(x) \in B\} \qquad (x \in \mathbb{X}, \ B \in \mathscr{B}, \ k \in \mathbb{N})$$

Under our assumptions on $\mathbb{X}$, for each kernel Q there exists at least one SRS which represents it (see, e.g., Kifer, 1986).

We require some stability assumptions on the stochastic kernel Q. In the assumptions, we suppose that (3) is an SRS which represents Q.

Assumption 2.1. There exists a set $C \in \mathscr{B}$ and a measurable function $H \colon \mathbb{U} \to \mathbb{X}$ such that $F(x, u) = H(u)$ for all $x \in C$.

Assumption 2.2. There exists an $m \in \mathbb{N}$ and an $E \in \mathscr{U}$ such that $\nu(E) = \mathbf{P}\{U_1 \in E\} > 0$, and

$$u_i \in E \text{ for } i = 1, \ldots, m \implies F_{u_m} \circ \cdots \circ F_{u_1}(x) \in C \ \forall x \in \mathbb{X}$$

In other words, there exists a set C such that whenever the current state X_t is in C, the next period state is given by $H(U_{t+1})$, independent of the precise value of X_t. Second, there exists a finite sequence of shock values with positive probability such that this sequence sends the process into C, regardless of the starting point x. Under these conditions we have the following result:

Proposition 2.1. *If assumptions 2.1 and 2.2 hold, then Q is globally stable.*

PROOF. Recall that Doeblin's condition is said to hold when there exists an $t \in \mathbb{N}$, a probability measure μ on $(\mathbb{X}, \mathscr{B})$ and a constant $\rho \in (0, 1)$ such that

$$Q^t(x, B) \geq \rho\mu(B) \quad \forall x \in \mathbb{X}, \forall B \in \mathscr{B} \qquad (6)$$

Under this assumption, it is well-known that Q is globally stable in the sense defined in (2).[8] We claim that Doeblin's condition holds. To see this, let m and

---

[8]See, e.g., Meyn and Tweedie (1993, thm. 16.2.4.).

E be as in assumption 2.2. Fix $x \in \mathbb{X}$ and $B \in \mathscr{B}$. Observe that

$$
\begin{aligned}
Q^{m+1}(x, B) &= \mathbf{P}\{F_{U_1} \circ F_{U_2} \circ \ldots \circ F_{U_{m+1}}(x) \in B\} \\
&\geq \mathbf{P}\left[\cap_{i=2}^{m+1}\{U_i \in E\} \cap \{F_{U_1} \circ F_{U_2} \circ \ldots \circ F_{U_{m+1}}(x) \in B\}\right] \\
&= \mathbf{P}\left[\cap_{i=2}^{m+1}\{U_i \in E\} \cap \{H(U_1) \in B\}\right] \\
&= \mathbf{P}\cap_{i=2}^{m+1}\{U_i \in E\}\mathbf{P}\{H(U_1) \in B\} \\
&= \epsilon\, v(H^{-1}(B))
\end{aligned}
$$

where $\epsilon := v(E)^{1/m}$. We have now verified Doeblin's condition, with $t = m + 1$, $\rho = \epsilon$ and μ defined as the image measure $v \circ H^{-1}$.

Example 2.2. Recall the optimal stopping problem considered by Rust (1987) in his seminal analysis of structural estimation for Markov decision processes. The problem is one of optimal timing for engine replacement. The state variable $X_t \in \mathbb{X} := \mathbb{R}_+$ represents accumulated mileage on the engine, and it follows the transition law

$$
Q(x, B) = \int_B v[y - x(1 - g(x))]dy \tag{7}
$$

where v is the exponential density $v(u) = \lambda e^{-\lambda u}$ for some $\lambda > 0$, and g is a replacement policy.[9] The range of g is $\{0, 1\}$, with $g(x) = 0$ indicating that the current engine is kept when the state is x, and $g(x) = 1$ indicating replacement. The kernel Q defined in (7) is represented by (and best understood in terms of) the SRS

$$
X_{t+1} = X_t(1 - g(X_t)) + U_{t+1} \tag{8}
$$

where $(U_t)_{t \geq 1}$ is IID on $\mathbb{U} := \mathbb{R}_+$ with density v. The optimal policy g is chosen as the solution to

$$
\min_{g \in \Sigma} \mathbf{E} \sum_{t \geq 0} \beta^t c(X_t, g(X_t))
$$

where Σ is the measurable binary functions on $\mathbb{X}$ (i.e., the indicator functions on $\mathscr{B}$) and c is a given cost function which need not be presented here. Rust (1987) shows that the optimal policy is given by $g(x) = \mathbb{1}\{x > \gamma\}$ for some constant $\gamma > 0$. Rewriting (8) with g defined in this way gives

$$
X_{t+1} = F_{U_{t+1}}(X_t) := X_t\mathbb{1}\{X_t \leq \gamma\} + U_{t+1} \tag{9}
$$

Assumption 2.1 is satisfied with $C := \{0\} \cup (\gamma, \infty)$ and $H(u) := u$. Assumption 2.2 is satisfied with $m = 1$ and $E := (\gamma, \infty)$.

Example 2.3. The commodity pricing model of Deaton and Laroque (1992) treats a competitive market for a single commodity. At the start of time t, current "harvest" U_t is made available to the market. The sequence $(U_t)_{t \geq 1}$ is IID,

---

[9]For $u < 0$ we set $v(u) = 0$.

and supported on an interval $[0, b]$ for some $b > 0$. That is, $\nu([0, b]) = 1$, and $\nu(G) > 0$ for any nonempty open set $G \subset [0, b]$.

Demand for the commodity is from "consumers" and "speculators." Speculators purchase the commodity in the current period and sell it in the next. Storage costs are positive: Quantity I_t carried over from time t yields αI_t at $t + 1$, where $\alpha \in (0, 1)$. Aggregate supply at time t is the sum of the harvest U_t and the depreciated investment αI_{t-1}. Thus, if X_t denotes supply, then $X_t = \alpha I_{t-1} + U_t$. Let $\bar{x} := (1 - \alpha)^{-1} b$ and $\mathbb{X} := [0, \bar{x}]$. Using $I_t \leq X_t$, one can show that $X_t \in \mathbb{X}$ implies $X_{t+1} \in \mathbb{X}$.

Based on rational expectations and arbitrage conditions, Deaton and Laroque develop a concept of recursive equilibrium, and, applying a contraction mapping argument, solve for an equilibrium pricing function and carry-over function $I \colon \mathbb{X} \to \mathbb{X}$. The carry-over function I is shown to be nondecreasing and satisfies $I(x) \leq x$. Moreover, there exists an $x_b > 0$ such that $x \leq x_b$ implies $I(x) = 0$. The equilibrium process for the state is given by

$$X_{t+1} = F_{U_{t+1}}(X_t) := \alpha I(X_t) + U_{t+1}$$

Assumptions 2.1 and 2.2 both hold for this process. To see this, let $C := [0, x_b]$ and $H(u) = u$. For $x \in C$ we have $F_u(x) = \alpha I(x) + u = u = H(u)$, so assumption 2.1 is satisfied. To verify assumption 2.2, let $m \in \mathbb{N}$ and $u_0 > 0$ be chosen such that

$$\alpha^m \bar{x} + u_0 \frac{1 - \alpha^m}{1 - \alpha} \leq x_b$$

Define $E := [0, u_0]$. Since $u_0 > 0$, we have $\nu(E) > 0$. Define (x_t) recursively by $x_{t+1} = \alpha I(x_t) + u_{t+1}$. If $u_i \in E$ for $i = 1, \ldots, m$, then Then for any fixed i with $i \leq m$ we have

$$x_i = \alpha I(x_{i-1}) + u_i \leq \alpha x_{i-1} + u_0$$

Combining these m inequalities gives

$$x_m \leq \alpha^m x_0 + u_0 \frac{1 - \alpha^m}{1 - \alpha} \leq \alpha^m \bar{x} + u_0 \frac{1 - \alpha^m}{1 - \alpha} \leq x_b$$

In other words, $x_m \in C$ whenever $u_i \in E$ for $i = 1, \ldots, m$, and hence assumption 2.2 is satisfied.

3. Perfect Sampling

Let Q be a stochastic kernel satisfying assumptions 2.1 and 2.2, and let ψ^* be the (unique) stationary distribution. The kernel is represented by SRS (3). In this setting, there is an obvious way to generate random variables which are approximately distributed according to ψ^*: Pick any initial condition $x \in \mathbb{X}$, generate independent observations of the shock, and compute X_t via (4). Since

the SRS represents Q, we have $\mathcal{L}X_t = Q^t(x, dy)$, and the latter converges to ψ^* by (2).

As discussed in the introduction, however, equality cannot be attained in finite time, and the size of the error caused by stopping at finite t is difficult to assess. As an alternative, then, consider the *reversed process* (Y_t) defined by

$$Y_t := F_{U_1} \circ F_{U_2} \circ \cdots \circ F_{U_t}(x) \qquad (x \in \mathbb{X}) \tag{10}$$

Since $(U_1, \ldots U_t)$ and $(U_t, \ldots, U_1)$ have the same joint distribution, it is clear that X_t and Y_t have the same distribution. That is,

$$\mathcal{L}F_{U_1} \circ \cdots \circ F_{U_t}(x) = \mathcal{L}F_{U_t} \circ \cdots \circ F_{U_1}(x) \quad \forall x \in \mathbb{X}, \ t \in \mathbb{N} \tag{11}$$

Hence $\mathcal{L}Y_t = Q^t(x, dy)$, and one can also use Y_t as an approximate draw from ψ^* when t is large.

The advantage of using the reversed process is as follows: On one hand, the sample paths of (X_t) usually fail to converge, in the sense that $\lim_t X_t$ does not exist. On the other hand, as we shall see, (Y_t) may well converge to a limiting random variable Y_∞ almost surely as $t \to \infty$. Assuming global stability, it is easily shown (see lemma 3.1 below) that $\mathcal{L}Y_\infty = \psi^*$. Moreover, under our assumptions, the convergence to Y_∞ occurs in *finite* time. Once this convergence has occurred, there is no need to simulate any further.

To formalize these arguments, we begin with the following lemma: Pick any $x \in \mathbb{X}$, and let Y_t be defined by (10). Then

Lemma 3.1. *If $Y_\infty := \lim_{t \to \infty} Y_t$ exists $\mathbf{P}$-a.s., then $\mathcal{L}Y_\infty = \psi^*$.*

PROOF. It suffices to show that $\mathbf{E}h(Y_\infty) = \int h d\psi^*$ for any bounded continuous $h \colon \mathbb{X} \to \mathbb{R}$. This is the case because, given any such h, we have

$$\mathbf{E}h(Y_\infty) = \mathbf{E}h(\lim_t Y_t) = \mathbf{E}\lim_t h(Y_t) = \lim_t \mathbf{E}h(Y_t)$$

$$= \lim_t \mathbf{E}h(X_t) = \lim_{t \to \infty} \int h(y)Q^t(x, dy) = \int h d\psi^*$$

where the third equality is by dominated convergence, the fourth is due to the fact that $\mathcal{L}Y_t = \mathcal{L}X_t$, and the last is by global stability.

Let's say that an arbitrary sequence (x_t) in $\mathbb{X}$ *coalesces* if it is eventually constant (i.e., $\exists n \in \mathbb{N}$ such that $x_t = x_n$ for all $t \geq n$). Suppose that Y_t coalesces almost surely. That is, there exists a measurable set $\Omega_\infty \subset \Omega$ with $\mathbf{P}(\Omega_\infty) = 1$, and such that, for all $\omega \in \Omega_\infty$, there exists an $n \in \mathbb{N}$ with $Y_t(\omega) = Y_n(\omega)$ for all $t \geq n$. In this case, (Y_t) converges almost surely, and we let Y_∞ denote the limit. By lemma 3.1, we have $\mathcal{L}Y_\infty = \psi^*$. Informally, since (Y_t) converges to Y_∞ in finite time, we can draw from ψ^* in finite time. Let us now formalize this idea. As above, $x \in \mathbb{X}$ is arbitrary, and Y_t is defined by (10).

Proposition 3.1. *The reversed sequence Y_t coalesces almost surely.*

The proof is given immediately, since its logic drives our first algorithm for perfect sampling. To interpret the proof in terms of computer simulation, imagine that at the start of the simulation a particular $\omega \in \Omega$ is randomly selected, and this value ω determines the entire path of shocks $(u_t)_{t \geq 1} := (U_t(\omega))_{t \geq 1}$ that will be observed in that simulation run. The random number generator simply returns successive elements of this sequence.

PROOF (PROOF OF PROPOSITION 3.1). Pick any initial condition $x \in \mathbb{X}$. Let E and m be as in assumption 2.2, let

$$\Omega_k := \{\omega \in \Omega : U_{k+i}(\omega) \in E \text{ for } i = 1, \ldots, m\}$$

and let $\Omega_\infty = \cup_{k \geq 1} \Omega_k$. Since $\nu(E) > 0$, we have $\mathbf{P}(\Omega_\infty) = 1$. Now fix $\omega \in \Omega_\infty$, so that, in particular, $\omega \in \Omega_k$ for some $k \in \mathbb{N}$. For ease of notation, let $u_j = U_j(\omega)$ for $j \in \mathbb{N}$. For any $t \geq k + m$,

$$Y_t(\omega) := F_{u_1} \circ \cdots \circ F_{u_k} \circ F_{u_{k+1}} \circ \cdots \circ F_{u_{k+m}} \circ \cdots \circ F_{u_t}(x)$$

In view of assumption 2.2, we have

$$F_{u_{k+1}} \circ \cdots \circ F_{u_{k+m}} \circ \cdots \circ F_{u_t}(x) \in C \qquad \forall t \geq k + m$$

It now follows from assumption 2.1 that

$$Y_t(\omega) = F_{u_1} \circ \cdots \circ F_{u_{k-1}} \circ H(u_k) \qquad \forall t \geq k + m$$

The right-hand side is constant in t, so $Y_t(\omega)$ has coalesced.

To produce a perfect sample, it suffices to generate and append successive shocks u_t to a running list $L = \{u_1, \ldots, u_{t-1}\}$, continuing until m consecutive elements of this list are in E. The value $Y_t(\omega)$ computed using the shocks in L is a perfect sample. This is the justification for algorithm 1. In the algorithm, L_m denotes the last m elements of the list L.

From the proof of proposition 3.1, algorithm 1 terminates in finite time with probability one, and the variate returned is a draw from ψ^*. It should be noted, however, that the stopping condition is sufficient but not necessary for coalescence of $Y_t(\omega)$, and might be a low probability event. In specific applications, more efficient test for coalescence may be available. An example is given below.

```
    set t ← m + 1 ;
    draw u₁, ..., uₜ independently from ν and set L ← {u₁, ... uₜ} ;
    repeat
        if Lₘ ∈ Eᵐ then                    /* last m elements all in E */
        |   break ;
        end
        set t ← t + 1 ;
        draw uₜ from ν and append it to the end of list L ;
    end
    return Yₜ := F_{u₁} ∘ F_{u₂} ∘ ··· ∘ F_{u_{t-m-1}} ∘ H(u_{t-m})
```

Algorithm 1: Perfect sampling

4. Applications

In our numerical application, we consider an equilibrium model of saving and asset accumulation under borrowing-constraints studied by Aiyagari (1994). Households maximize the objective function

$$\max_{g \in \Sigma} \mathbf{E} \sum_{t \geq 0} \beta^t u(Z_t - g(Z_t))$$

where Σ is the set of measurable (policy) functions $g \colon \mathbb{R}_+ \to \mathbb{R}_+$ satisfying $0 \leq g(z) \leq z$, and the state variable $(Z_t)_{t \geq 0}$ obeys $Z_{t+1} = wU_{t+1} + (1+r)g(Z_t) - r\phi$. Here ϕ represents the borrowing constraint, w is wages, r in the interest rate, U_t is an idiosyncratic labor supply shock, and the control variable $g(Z_t)$ represents assets plus ϕ (Aiyagari, 1994, p. 666). The utility function is specified as $u(c) := c^{1-\sigma}/(1-\sigma)$. To simplify the exposition, $(U_t)_{t \geq 1}$ is taken to be IID, with common distribution ν. This corresponds to section III of Aiyagari (1994). The optimal policy satisfies

$$g(z; w, r) =$$

$$\text{argmax}_{0 \leq a \leq z} \left\{ u(z - a) + \beta \int V(wx + (1+r)a - r\phi)\nu(dx) \right\} \quad (12)$$

where V is the value function. A stationary equilibrium is a set of values, K, r and w such that

$$r = f_1(K, 1) - \delta, \quad w = f_2(K, 1), \quad \text{and} \quad K = \int g(z; w, r)\psi^*(dz) \quad (13)$$

where f is the production function and ψ^* is the stationary distribution of the state variable Z_t under the optimal law of motion

$$Z_{t+1} = wU_{t+1} + (1+r)g(Z_t; w, r) - r\phi \quad (14)$$

10

Numerical solution involves guessing the value of K, determining r and w via the first two equations in (13), solving the household problem for $g(\cdot; w, r)$, obtaining ψ^* from (14), and adjusting K to match the right-hand side of the final equality in (13). Most studies treat this problem by discretizing the state space, but we will use a continuous state space and fitted value iteration. The integral $\int g(z; w, r)\psi^*(dz)$ is then computed using Monte Carlo and perfect sampling. To focus attention on perfect sampling, we consider only computation of this integral for given w and r.

As in Aiyagari (1994), we take $\beta = 0.96$, $\sigma = 2$ and $\phi = 0$. The shock U_t is assumed to be discrete, with distribution ν chosen to match the stationary distribution of the shock process in Aiyagari. In particular, $U_t \in \{1 - d, 1, 1 + d\}$, where $d = 0.49$ is selected so that the standard deviation of U_t is 0.4. We set $w = 1.3712$ and $r = 0.0129$, which are chosen to approximate equilibrium values.

Given these parameters, an approximation to the value function is calculated using fitted value iteration. Piecewise linear interpolation is used to approximate the value functions at each step. The resulting approximate optimal policy (12) converges in value to the true optimal policy as the number of iterations and the distance between grid points converge to infinity and zero respectively (Stachurski, 2008). Here, interpolants are computed on an evenly spaced grid of 150 elements.

The approximate optimal policy $\hat{g}$ is also represented by a piecewise linear interpolant, where the interpolation values are obtained by solving (12) on the grid points. The resulting dynamics for $(Z_t)_{t \geq 0}$ are then given by

$$Z_{t+1} = F_{U_{t+1}}(Z_t) := wU_{t+1} + (1 + r)\hat{g}(Z_t) \tag{15}$$

Given our solution for $\hat{g}$, the SRS has the following features. First, $\hat{g}$ and hence $z \mapsto F_u(z)$ are monotone nondecreasing. Second, for any realization u of U, we have $F_u(\bar{z}) < \bar{z}$ when $\bar{z} := 14$, and hence $\mathbb{X} := [0, \bar{z}]$ can serve as a state space for (15). Third, $\hat{g}$ satisfies $\hat{g}(z) = 0$ for $z < z_b := 0.95274$.

Consider now the problem of generating random deviates with distribution ψ^*. Assumptions 2.1 and 2.2 are both satisfied,[10] and theorem 3.1 applies. However, we will construct an application-specific method of detecting when the reversed process has coalesced.

To begin, let $Y_t := F_{U_1} \circ \cdots \circ F_{U_t}(\bar{z})$ be the reversed process starting from $\bar{z}$, and define the intermediate values of this process by

$$Y_t^i := F_{U_i} \circ F_{U_{i+1}} \circ \cdots \circ F_{U_t}(\bar{z}) \qquad (t \geq i)$$

---

[10] Assumption 2.1 is satisfied with $C = [0, z_b)$ and $H(u) = wu$. Regarding assumption 2.2, it turns out that $F(z, 1 - d) < z$ for all $z \in [0, \bar{z}]$, so a sufficiently long sequence of bad shocks drives the state below z_b (i.e., into C), independent of the current state.

so that $Y_t := F_{U_1} \circ \cdots \circ F_{U_{i-1}}(Y_t^i)$ for each $i \leq t$. We will consider values of this process with respect to a fixed realization of the shocks $(u_t)_{t \geq 1}$. Lower case symbols such as y_t and y_t^i to represent realizations of the reversed process corresponding to this fixed shock path.

Lemma 4.1. *Given any $i \leq n \leq t$, we have $y_t^i \leq y_n^i$.*

PROOF. We can and do assume that $n < t$. Recall that $z \mapsto F_u(z)$ is nonde-creasing for any u, and that $\bar{z}$ is an upper bound for the state space, implying $y_t^{n+1} \leq \bar{z}$. It follows that

$$y_t^i := F_{u_i} \circ \cdots \circ F_{u_t}(\bar{z}) = F_{u_i} \circ \cdots \circ F_{u_n}(y_t^{n+1}) \leq F_{u_i} \circ \cdots \circ F_{u_n}(\bar{z}) =: y_n^i$$

as was to be shown.

Lemma 4.2. *If $1 < i < t$ and $y_t^i < z_b$, then $y_t = F_{u_1} \circ \cdots \circ F_{u_{i-2}}(wu_{i-1})$.*

PROOF. Since $\hat{g}(z) = 0$ whenever $z < z_b$, the assumption $y_t^i < z_b$ implies that $F_{u_{i-1}}(y_t^i) := wu_{i-1} + (1+r)\hat{g}(y_t^i) = wu_{i-1}$. But then

$$y_t = F_{u_1} \circ \cdots \circ F_{u_{i-1}}(y_t^i) = F_{u_1} \circ \cdots \circ F_{u_{i-2}}(wu_{i-1})$$

as was to be shown.

Lemmas 4.1 and 4.2 suggest the following technique for perfect sampling. Suppose that the reversed process computed from $u_1, \ldots, u_n$ falls below z_b at some $i > 1$. More precisely, there exists a i with $1 < i \leq n$ and $y_n^i < z_b$. In that case the sequence (y_t) has coalesced: $y_t = y_n$ for all $t \geq n$. To see this, pick any $t > n$, and suppose that we generate additional observations $u_{n+1}, \ldots, u_t$ of the shock and append them to the existing sequence $u_1, \ldots, u_n$ to obtain $u_1, \ldots, u_t$. Since $y_t^i \leq y_n^i$ (Lemma 4.1) and $y_n^i < z_b$, we have $y_t^i < z_b$. From lemma 4.2 it then follows that $y_t = F_{u_1} \circ \cdots \circ F_{u_{i-2}}(wu_{i-1})$, which is independent of t.

We have now justified algorithm 2 as a means of sampling from ψ^*. In the algorithm, L is a list used to store the shocks draws generated within the loop. The variable $C \in \{0,1\}$ switches to one when $y_t^i < z_b$ at some $i > 1$, indicating that (y_t) has coalesced.

Note that repeated draws from the algorithm are *independent*. Hence repetition allows one to generate IID draws from ψ^*. We can then calculate the desired integral (i.e., aggregate capital given the stationary distribution) using Monte Carlo:

$$\int \hat{g}(z)\psi^*(dz) \cong \frac{1}{n}\sum_{i=1}^n \hat{g}(Z_i)$$

```
set t ← 1, C ← 0, y ← z̄, L ← { } ;
repeat
    draw u_t ∼ ν and append to the end of list L ;
    for i in t,...,1 do                    /* compute y_t recursively */
        y ← F_{u_i}(y) ;
        if y < z_b and i > 1 then set C ← 1 ;
    end
    if C = 1 then  break ;                 /* process has coalesced */
    else
        t ← t + 1 ;
        y ← z̄ ;
    end
end
return y
```

Algorithm 2: Perfect sampling, Aiyagari model

where the sequence $(Z_i)_{i=1}^n$ is the draws from ψ^*. Moreover, since the draws are IID, and since the second moment satisfies

$$\int \hat{g}(z)^2 \psi^*(dz) \leq \hat{g}(\bar{z})^2 < \infty$$

we have

$$\sqrt{n} \left[\frac{1}{n} \sum_{i=1}^n \hat{g}(Z_i) - \int \hat{g}(z)\psi^*(dz) \right] \to N(0, \sigma^2) \tag{16}$$

in distribution, for

$$\sigma^2 := \int \left(\hat{g}(z) - \int \hat{g}(z')\psi^*(dz') \right)^2 \psi^*(dz)$$

The latter quantity can be estimated from the sample to provide an asymptotic confidence interval for the integral.

To test the algorithm we drew 100,000 independent perfect samples from ψ^*. The algorithm was implemented in C, and compiled using the gcc compiler.[11] Some modifications were made to the algorithm in order to reduce runtime, such as incrementing time in steps of 100, rather than 1. On a system running Ubuntu Linux with 2GB memory and a dual core 2.13GHz CPU, the total runtime for the program was 2.4 seconds.[12] The length of the 99% confidence

---

[11]See http://johnstachurski.net/lectures/aiyagari.html for the code and instructions on compilation.

[12]Naturally, this does not include the time required to compute the approximate optimal policy $\hat{g}$, which is not related to algorithm 2. As discussed above, $\hat{g}$ was implemented using linear interpolation. The interpolation values were computed separately via a program written in Python, and passed into the C program via file I/O in order to implement the SRS (15).

interval implied by (16) was 0.00054, indicating a very high degree of precision.

5. Conclusion

In this paper we discussed perfect sampling for a class of dynamic models. Under certain assumptions on the process, it was shown that perfect samples could be obtained in finite time with probability one. We also presented a modified algorithm for generating perfect samples in the case of Aiyagari's (1994) model. The algorithm permits rapid generation of IID variates from the stationary distribution.

Not all dynamic models of interest to economists satisfy our assumptions. For example, the standard neoclassical growth model does not generally satisfy assumptions 2.1 and 2.2. However, simple models from this class have been shown to be ergodic and Harris recurrent (cf., e.g., Nishimura and Stachurski, 2005), and algorithms for perfect sampling from ergodic Harris recurrent Markov processes were obtained by Corcoran and Tweedie (2001). Generalization of the theoretical results and implementation of the algorithms is left for future research.

[1] Aiyagari, S. R., 1994. Uninsured idiosyncratic risk and aggregate saving. Quarterly Journal of Economics 109, 659–684.

[2] Athreya, K. B., Stenflo, O., 2003. Perfect sampling for Doeblin chains. Sankhya 65, 763–777.

[3] Corcoran, J. N., Tweedie, R. L., 2001. Perfect sampling of ergodic Harris chains. Annals of Applied Probability 11, 438–451.

[4] Deaton, A., Laroque, G., 1992. On the behavior of commodity prices. Review of Economic Studies 59, 1–23.

[5] Duffie, D., Singleton, K., 1993. Simulated moments estimation of Markov models of asset prices. Econometrica 61, 929–952.

[6] Fernández-Villaverde, J., Rubio-Ramírez, J. F., Santos, M. S., 2006. Convergence properties of the likelihood of computed dynamic models. Econometrica 74, 93–119.

[7] Foss, S., Tweedie, R. L., 1998. Perfect simulation and backward coupling. Stochastic Models 14, 187–204.

[8] Gouriéroux, C., Monfort, A., 1996. Simulation-Based Econometric Methods. Oxford University Press.

[9] Kifer, Y., 1986. Ergodic Theory of Random Transformations. Birkhäuser, Boston.

[10] Lee, B., Ingram, B., 1991. Simulation estimation of time series models. Journal of Econometrics 47, 197–205.

[11] Meyn, S., Tweedie, R. L., 1993. Markov Chains and Stochastic Stability. Springer-Verlag, New York.

[12] Murdoch, D. J., Green, P. J., 1998. Exact sampling from a continuous state space. Scandinavian Journal of Statistics 25, 483–502.

[13] Nishimura, K., Stachurski, J. 2005. Stability of stochastic optimal growth models: A new approach. Journal of Economic Theory 122, 100–118.

[14] Propp, J. G., Wilson, D. B., 1996. Exact sampling with coupled Markov chains and applications to statistical mechanics. Random Structures and Algorithms 9, 223–252.

[15] Rust, J., 1987. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. Econometrica 55, 999–1033.

[16] Santos, M. S., Peralta-Alva, A., 2005. Accuracy of simulations for stochastic dynamic models. Econometrica 73, 1939–1976.

[17] Stachurski, J., 2008. Continuous state dynamic programming via nonexpansive approximation. Computational Economics 31, 141–160.