

Exact Present Solution with Consistent Future Approximation: A Gridless Algorithm to Solve Stochastic Dynamic Models

Wouter J. Den Haan, Michal L. Kobielarz and Pontus Rendahl*

January 12, 2016

Abstract

This paper proposes an algorithm that finds model solutions at a particular point in the state space by solving a simple system of equations. The key step is to characterize future behavior with a Taylor series expansion of the current period's behavior around the contemporaneous values for the state variables. Since current decisions are solved from the original model equations, the solution incorporates nonlinearities and uncertainty. The algorithm is used to solve the model considered in Coeurdacier, Rey, and Winant (2011), which is a challenging model because it has no steady state and uncertainty is necessary to keep the model well behaved. We show that our algorithm can generate accurate solutions even when the model series are quite volatile. The solutions generated by the risky-steady-state algorithm proposed in Coeurdacier, Rey, and Winant (2011), in contrast, is shown to be not accurate.

Key Words: Risky Steady State, Solution Methods

JEL Classification: C63, E10, E23, F41

*Den Haan: Centre for Macroeconomics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK and CEPR, London, UK. E-mail: wjdenhaan@gmail.com. Kobielarz: CentER, Tilburg University. E-mail: michal.kobielarz@gmail.com. Rendahl: Centre for Macroeconomics, University of Cambridge. E-mail: pontus.rendahl@gmail.com.

1 Introduction

Carpe diem, quam minimum credula postero

(seize the present, trust tomorrow e'en as little as you may)

The difficulty in solving stochastic dynamic rational expectations models is that agents are forward looking, which means that one cannot determine this period's behavior unless one knows next period's behavior. It is standard practice to find a recursive solution, which means that the model's solution (as a function of the model's state variables) are the same each period. A model solution is then a set of functions with the following property: If this set of functions is used to describe future model outcomes, then this period's model outcomes are described by the same set of functions. That is, a solution is a fixed point in function space.¹

By contrast, if future model outcomes are known (as a function of next period's state variables), then the problem is much simpler. If the model has n endogenous variables, then finding the solution for a particular set of state variables would require solving a system of n equations in n unknowns. That is, instead of finding a solution in function space, one only has to find a solution in much smaller n -dimensional Euclidian space. Building on this logic, Den Haan and De Wind (2012) propose to describe future behavior using a simple perturbation approximation and solve for this period's behavior from the original model equations. Since the original model equations are used, this solution does take into account nonlinearities, uncertainty, and any possible interaction between the two. Unfortunately, describing future behavior with a simple perturbation solution can be so inaccurate that this period's choices are inaccurate as well.

The algorithm that we propose eliminates the ad hoc choice of describing future behavior, but still finds the solution by solving a simple system of equations. The idea is the following. Given future behavior, one can solve for this period's behavior. If one can solve for this period's solution, then one can also solve for the derivatives of this period's solution. The key step of our algorithm is to use these derivatives to construct a Taylor series

¹The arguments of the functions are the model's state variables.

expansion, which is then used to describe next period's behavior. More formally, we solve for model outcomes in each period using a small system of equations that contains the original model equations as well as some additional equations which ensure that next period's behavior is a Taylor series expansion of this period's behavior, around this period's state variables. Since this period's outcomes are solved using the original model equations, they incorporate any possible nonlinearities, uncertainty, and interactions thereof. We refer to our algorithm as the Exact-Today (ET) algorithm since this period's outcomes are an exact solution of the model equations and we only approximate next period's outcomes.

In contrast to perturbation methods, our Taylor series expansion is only used to characterize next period's behavior in the model equations. Each period, actual behavior is solved from the original model equations and incorporates any possible consequences of nonlinearities and/or uncertainty. The algorithm's advantage relative to projection methods is that it does not require constructing a grid. In more complex models, it may be difficult to construct a grid such that all calculations make sense at all nodes on the grid.² Another problem with grid-based methods is the curse of dimensionality, in which the complexity of the problem increases exponentially with the number of state variables. In contrast, the complexity of our algorithm only increases linearly with the number of state variables.

There are already quite a few algorithms to solve stochastic dynamic models. To document the usefulness of our algorithm, we implement it using a challenging model. This is the model considered in Coeurdacier, Rey, and Winant (2011) (CRW). The model is cast in partial equilibrium in which an agent faced with stochastic income and stochastic returns decides how much to save and how much to consume. This type of model is often used in open economy macroeconomics to describe small open economies. The difficulty of this model lies not in its size but in that uncertainty is key in keeping the model well behaved. There exist no steady state and savings diverge absent a sufficient amount of uncertainty. Moreover, as documented in this paper savings also diverge if uncertainty is

²As discussed in section 2.3, this matters in practice. For example, particular combinations of state variables may make no sense. The model would never end up at those points, but one may not be aware of this when constructing the grid.

too large.

CRW propose a modified perturbation method to solve this model, which entails taking a second-order approximation of the Euler equation and finding a consistent perturbation solution. We document that their solution is actually very inaccurate. By contrast, the solution generated by our algorithm is shown to be very close to an accurate projection algorithm that solves the problem using more than eight million nodes.³

Section 2 describes and motivates the algorithm. Section 3 describes the model of Coeurdacier, Rey, and Winant (2011), which we use as a test case for our algorithm. Section 4 documents the accuracy of our algorithm and the algorithm proposed in Coeurdacier, Rey, and Winant (2011).

2 The ET algorithm

In this section, we describe the method, which is followed by a short subsection describing what one would actually have to program. The method’s merits are discussed in the last subsection.

We focus on a class of models that can be represented by the following system of equations:

$$0 = \mathbb{E}_t [f(x_{t-1}, x_t, x_{t+1}, z_t, z_{t+1})], \quad (1)$$

where x_t contains the endogenous variables, z_t the exogenous random variables, $\mathbb{E}_t[\cdot]$ denotes the expectation operator, and $f(\cdot)$ is a known, typically nonlinear, function.⁴ To simplify notation, we focus on the case when both x_t and z_t are scalars, but the method easily generalizes to the multidimensional case. For the same reason, we assume that z_t is

³Eight million nodes may seem overly prudent. However, the projection method using “only” 441,000 actually performs worse than our proposed algorithm (!), and a very fine grid is therefore necessary in order to adequately assess the accuracy of the various approaches.

⁴When solving stochastic dynamic models, this system of equations is typically the set of first-order and equilibrium conditions. The variables could be transformations of the underlying economic variables such as the logarithm.

a simple AR(1) process. That is,

$$z_{t+1} = \rho z_t + \varepsilon_{t+1}, \quad (2)$$

$$\mathbb{E}_t [\varepsilon_{t+1}] = 0. \quad (3)$$

2.1 The method

The objective is to find a recursive solution to equation (1) of the following form:

$$x_t = g(x_{t-1}, z_t). \quad (4)$$

Projection methods posit an approximating function, $p(x_{t-1}, z_t; \Phi)$, where Φ is a vector containing the coefficients of the approximating function, and find Φ such that equation (1) holds exactly, or approximately, on a grid for (x_{t-1}, z_t) . Perturbation methods also specify a particular functional form and choose the coefficients Φ such that the derivatives of $p(\cdot; \Phi)$ are consistent with the implicit solution of equation (1) at the steady state. By contrast, our algorithm, does not restrict the solution of the *current period's* decisions to be of a particular functional form. Instead, it directly solves for x_t *one point at the time*, by solving a simple system of equations including equation (1). By doing so, it preserves the nonlinearities in $f(\cdot)$ as well as any interaction between nonlinearities in $f(\cdot)$ and uncertainty about ε_{t+1} .

The reader may think that it is not possible to directly solve for x_t using equation (1), because equation (1) clearly indicates that to solve for x_t one needs to know how x_{t+1} is determined. The underlying principle of ET that makes this possible is that the relationship between x_{t+1} and *next period's* state variables is imposed to be equal to an approximation of the relationship between x_t and *this period's* state variables. When solving for x_t , this approximation for x_{t+1} is only used to describe the behavior at the relevant points in the state space, namely the points where the economy could find itself next period. Moreover, it is important to realize that the approximation we use to describe next period's behavior is not fixed, but depends on the period- t state variables. This process is repeated each period, so that the actual outcome in period t is not given by

this approximation, but by the solution to the "ET-system" of equations that includes equation (1).

The order of approximation for ET may vary depending on the desired degree of accuracy. With first-order ET, the value of next period's choice, x_{t+1} , is approximated by the following linear function:⁵

$$x_{t+1} \approx \tilde{h}(x_t, z_{t+1}; \Phi_t) = \phi_{0,t} + \phi_{x,t}(x_t - \tilde{x}_t) + \phi_{z,t}(z_{t+1} - \tilde{z}_{t+1}), \quad (5)$$

where Φ_t contains the coefficients of the approximation. That is,⁶

$$\Phi_t = \Phi(x_{t-1}, z_t) \equiv [\tilde{x}_t, \tilde{z}_{t+1}, \phi_{0,t}, \phi_{x,t}, \phi_{z,t+1}]. \quad (6)$$

The five elements of Φ_t are functions of the state variables, x_{t-1} and z_t , and are, thus, time varying. For the ET algorithm it does not matter whether x_{t-1} and z_t are observations in a simulated time series or an arbitrary point in the state space. Therefore, we focus on the recursive representation of the system of equations, that is,

$$0 = \mathbb{E}[f(x_{-1}, x, x_{+1}, z, z_{+1})], \quad (7a)$$

$$z_{+1} = \rho z + \varepsilon_{+1}, \quad (7b)$$

$$\mathbb{E}[\varepsilon_{+1}] = 0. \quad (7c)$$

Substituting the approximation for x_{+1} given by equation (5) into equation (7a) gives

$$\begin{aligned} 0 &= \mathbb{E}\left[f\left(x_{-1}, x, \tilde{h}(x, z_{+1}; \Phi(x_{-1}, z)), z, z_{+1}\right)\right] \\ &= \mathbb{E}\left[f\left(x_{-1}, x, \tilde{h}(x, \rho z + \varepsilon_{+1}; \Phi(x_{-1}, z)), z, \rho z + \varepsilon_{+1}\right)\right] \\ &= \mathbb{E}[F(x_{-1}, x, z, \varepsilon_{+1}; \Phi(x_{-1}, z))]. \end{aligned} \quad (8)$$

Since ε_{+1} is integrated out, equation (8) implicitly defines x as a function of x_{-1} and z for a given $\Phi(x_{-1}, z)$. We denote this function by $h(x_{-1}, z)$. Thus,

$$0 = \mathbb{E}[F(x_{-1}, h(x_{-1}, z), z, \varepsilon_{+1}; \Phi(x_{-1}, z))]. \quad (9)$$

⁵Higher-order ET is discussed in appendix D.

⁶It is important to distinguish between x_t and \tilde{x}_t and between z_{t+1} and \tilde{z}_{t+1} . x_t and z_{t+1} are next period's state variables whereas \tilde{x}_t and \tilde{z}_{t+1} are coefficients of $\tilde{h}(\cdot)$.

The idea of ET is to choose $\Phi(x_{-1}, z)$ such that the function that determines x_{+1} , i.e., $\tilde{h}(x, z_{+1}; \Phi(x_{-1}, z))$, is the Taylor series expansion of the (implicit) function that determines this period's choice, $h(x_{-1}, z)$, around this period's state variables. The approximation $\tilde{h}(x, z_{+1}; \Phi(x_{-1}, z))$ will be accurate if next period's state variables are not that different from this period's state variables. This is a much weaker condition than the one that will ensure accuracy of the perturbation solution, which is that both this period's and next period's state variables are close enough to their steady state values *and* uncertainty is close to zero. Moreover, the Taylor series approximation, plays a less important role in the ET algorithm relative to perturbation methods, since it is only used to describe possible outcomes for next period's behavior and not to solve for each period's actual outcomes.

Since $\tilde{h}(x, z_{+1}; \Phi(x_{-1}, z))$ is a Taylor series expansion of $h(\cdot)$ around x_{-1} and z , it follows that

$$\tilde{x} = x_{-1}, \quad (10a)$$

$$\tilde{z}_{+1} = z. \quad (10b)$$

That is, \tilde{x} and \tilde{z}_{+1} are equal to x_{-1} and z respectively. Of course, the values of the coefficients ϕ_0 , ϕ_x , and ϕ_z also depend on x_{-1} and z . To understand how these coefficients are determined, notice that

$$\tilde{h}(\tilde{x}, \tilde{z}; \Phi(x_{-1}, z)) = \phi_0, \quad (11a)$$

$$\frac{\partial \tilde{h}(x, z_{+1}; \Phi(x_{-1}, z))}{\partial x} = \phi_x, \quad (11b)$$

$$\frac{\partial \tilde{h}(x, z_{+1}; \Phi(x_{-1}, z))}{\partial z_{+1}} = \phi_z. \quad (11c)$$

Differentiating equation (9) with respect to x_{-1} and z , gives

$$0 = \mathbb{E} \left[\begin{array}{c} \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi(x_{-1}, z))}{\partial x_{-1}} \\ + \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi(x_{-1}, z))}{\partial x} \frac{\partial h(x_{-1}, z)}{\partial x_{-1}} \end{array} \right], \quad (12a)$$

$$0 = \mathbb{E} \left[\begin{array}{c} \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi(x_{-1}, z))}{\partial z} \\ + \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi(x_{-1}, z))}{\partial x} \frac{\partial h(x_{-1}, z)}{\partial z} \end{array} \right]. \quad (12b)$$

From this equation, we obtain expressions for $\partial h(x_{-1}, z)/\partial x_{-1}$ and $\partial h(x_{-1}, z)/\partial z$. From now on, we suppress the dependence of $\Phi(x_{-1}, z)$ on x_{-1} and z , although this dependence is a key feature of ET. The key approximation step of ET is to assume that the coefficients of the approximation for next period's choices are set as follows:

$$\phi_0 = \tilde{h}(\tilde{x}, \tilde{z}; \Phi) = h(x_{-1}, z), \quad (13a)$$

$$\phi_x = \left. \frac{\partial \tilde{h}(x, z_{+1}; \Phi)}{\partial x} \right| = \left. \frac{\partial h(\hat{x}_{-1}, \hat{z})}{\partial \hat{x}_{-1}} \right|_{\substack{\hat{x}_{-1}=x_{-1}, \\ \hat{z}=z}}, \quad (13b)$$

$$\phi_z = \left. \frac{\partial \tilde{h}(x, z_{+1}; \Phi)}{\partial z_{+1}} \right| = \left. \frac{\partial h(\hat{x}_{-1}, \hat{z})}{\partial \hat{z}} \right|_{\substack{\hat{x}_{-1}=x_{-1}, \\ \hat{z}=z}}. \quad (13c)$$

In words, the function determining *next period's* choice, $x_{+1} = \tilde{h}(x, z_{+1}; \Phi)$, is equal to a local approximation of *this period's* choice, $x = h(x_{-1}, z)$, around this period's values of the state variables. Equation (13a) immediately implies that $\phi_0 = x$. This leaves us with the following system of three equations to solve for x , ϕ_x , and ϕ_z :⁷

$$0 = \mathbb{E}[F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)], \quad (14a)$$

$$0 = \mathbb{E}\left[\frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)}{\partial x_{-1}} + \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)}{\partial x} \phi_x\right], \quad (14b)$$

$$0 = \mathbb{E}\left[\frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)}{\partial z} + \frac{\partial F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)}{\partial x} \phi_z\right]. \quad (14c)$$

Lastly, we have to deal with the expectations operator. If ε_{+1} has discrete support, then the expectations operator can be replaced by a sum of the outcomes corresponding to the possible realizations of ε_j multiplied with the associated weights, π_j , with $j \in \{1, \dots, J\}$. If ε_{+1} has continuous support, then one can use a numerical integration procedure, which boils down to doing the same. For example, if ε_{+1} has a Normal distribution with mean zero and standard deviation equal to σ , then using Gauss-Hermite quadrature implies that

$$\varepsilon_j = \sqrt{2}\sigma\zeta_j, \text{ and} \quad (15)$$

$$p_j = \frac{\omega_j}{\sqrt{\pi}}, \quad (16)$$

where the ζ_j 's and the ω_j 's are the Gauss-Hermite nodes and weights, respectively.⁸

⁷Although $\tilde{x} = x_{-1}$ and $\tilde{z}_{+1} = z_t$, these variables should be treated as constant coefficients when differentiating $F(x_{-1}, x, z, \varepsilon_{+1}; \Phi)$.

⁸Intuitively, one can think of the $\sqrt{2}\sigma\zeta_j$ terms as the possibly outcomes of ε_{+1} and of the $\omega_j/\sqrt{\pi}$ terms as

Relationship to Den Haan and De Wind (2012) Den Haan and De Wind (2012) propose to set $\tilde{h}(x, z_{+1}; \Phi)$ equal to a standard perturbation solution. Equation (8) can then also be used to solve for x for given values of x_{-1} and z . In this approach, the value of Φ is time invariant and independent of x_{-1} and z . The solution for x would still take into account the nonlinearity of $f(\cdot)$ and the uncertainty about ε_{+1} . However, the standard perturbation solution may not be a reliable representation of the solution for x_{+1} , because it is a local approximation around the steady state values of the state variables when there is no uncertainty. By contrast, ET sets the choice for x_{+1} equal to the behavior according to a local approximation around this period's state variables and uncertainty is not reduced to zero.⁹

2.2 What actually needs to be programmed

The key step is to construct a function that defines the three error terms associated with equation (14) as a function of the arguments, x , ϕ_x , and ϕ_z for given values of x_{-1} and z . Those error terms are given by

$$\begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^J F(x_{-1}, x, z, \varepsilon_j; \Phi) p_j \\ \sum_{j=1}^J \left(\frac{\partial F(x_{-1}, x, z, \varepsilon_j; \Phi)}{\partial x_{-1}} + \frac{\partial F(x_{-1}, x, z, \varepsilon_j; \Phi)}{\partial x} \phi_x \right) p_j \\ \sum_{j=1}^J \left(\frac{\partial F(x_{-1}, x, z, \varepsilon_j; \Phi)}{\partial z} + \frac{\partial F(x_{-1}, x, z, \varepsilon_j; \Phi)}{\partial x} \phi_z \right) p_j \end{bmatrix}, \quad (17)$$

the associated probabilities. Simple subroutines exist to generate the Gauss-Hermite quadrature nodes and weights. Numerical integration techniques are very powerful. For example, if ε_{+1} is distributed according to a Normal distribution, and one wants to calculate $\mathbb{E}[F(\varepsilon_{+1})]$, then the Gauss-Hermite quadrature approximation with J nodes gives the exact for *all* $(2J - 1)^{\text{th}}$ -order polynomials, and an accurate answer when $F(\cdot)$ is well approximated by a $(2J - 1)^{\text{th}}$ order polynomial.

⁹The procedure used in Den Haan and De Wind (2012) improves substantially upon standard perturbation solutions for some models. However, it is shown to be only slightly better than the regular perturbation solution – and still far from accurate – for others.

whereas before

$$F(x_{-1}, x, z, \varepsilon_j; \Phi) = f\left(x_{-1}, x, \tilde{h}(x, \rho z + \varepsilon_j; \Phi), z, \rho z + \varepsilon_j\right), \quad (18)$$

$$\tilde{h}(x, \rho z + \varepsilon_j; \Phi) = \phi_0 + \phi_x(x - \tilde{x}) + \phi_z(z_{+1} - \tilde{z}), \quad (19)$$

and

$$\Phi = [\tilde{x}, \tilde{z}, \phi_0, \phi_x, \phi_z] = [x_{-1}, z, x, \phi_x, \phi_z]. \quad (20)$$

Defining these error terms requires taking derivatives. As shown in appendices D and F, this is a very simple programming step if one has access to a symbolic toolbox even if one considers higher-order approximations. For quite a few problems this could also be done by hand.

The next and final step is to use an equation solver to find the values of x , ϕ_x , and ϕ_z that set the three error terms equal to zero for given values of x_{-1} and z .

2.3 Merits of the ET algorithm

The ET algorithm shares with projection methods the property that the numerical solution incorporates the uncertainty and nonlinearity underlying the original system of equations without modifying $f(\cdot)$ and with either no modification of the expectations operator (in case ε_{+1} has discrete support) or with a minor modification (in case ε_t has continuous support). This is even true for first-order ET, since the period- t choice variables are always solved from the original nonlinear set of first-order conditions. The approximation aspect of the ET solution *only* affects the characterization of *next period's* policy functions. Moreover, the derivatives of the Taylor-series expansion used to characterize next period's behavior vary with the value of this period's state variables. By contrast, the traditional perturbation method characterizes this period's behavior with a particular *fixed* Taylor series expansion using the state variables *and* the amount of uncertainty as its arguments. The derivatives used in this Taylor-series approximation correspond to the derivatives at the steady state. This means that the accuracy of the standard perturbation solution is only guaranteed if the state variables are close to their steady state values and the model's volatility parameters are close to zero. Higher-order perturbation does incorporate the

effect of uncertainty, but imposes that the effect of uncertainty on model outcomes is equal to an extrapolation of the effect of uncertainty on behavior at a situation when there is no uncertainty to begin with.

Projection methods, on the other hand, solve a much larger set of simultaneous equations, namely the model equations at *all* grid points.¹⁰ By contrast, ET can solve for the choices at a particular point in the state space without having to solve for the solution at other points in the state space. For the application considered in this paper, which is quite challenging, the projections algorithm has to solve a simultaneous system of more than eight million equations (one for each node), whereas second-order ET needs to solve a system of ten equations and ten unknowns.¹¹ ET can do so because it focuses on behavior at the *relevant* points in the state space, namely those where the economy could be in the next period, and it does so by imposing that next period's choices are determined by a Taylor-series expansion of this period's choices. It is true that ET has to solve this system many times, but it is a fixed system of equations, so it only has to be defined once.

It is important to realize that the comparison of standard projection methods and ET is not only an issue of comparing the speed of solving a large system of equations once versus solving a smaller system of equations many times. There are two potential advantages of ET. The first is that ET is not subject to the curse of dimensionality the way it affects regular grid-based methods. Whereas the number of elements of a regular grid increases exponentially with the number of state variables, the number of unknowns the ET algorithm has to solve for increases at a slower rate. In particular, it increases at a linear rate if a first-order approximation is used to describe next period's behavior.¹² A second advantage of ET is related to the points considered in calculating the numerical approximation. To apply standard projection methods, one has to construct a grid and the algorithm must be able to find a sensible solution at *all* grid points. But there may be points in the state space where some calculations are not well defined.¹³ This is not a

¹⁰If the number of nodes exceeds the number of coefficients, then a minimization routine would be used.

¹¹First-order ET requires solving a system of three equations in three unknowns.

¹²Recall that even the first-order ET solution takes into account the full nonlinearity and uncertainty of the underlying system.

¹³One possibility is that the model simply has no solution at these nodes. Another possibility is that

problem if the economy never reaches those points. Unfortunately, when constructing the grid one typically would not know which points in the state space one has to exclude for this reason. This is a relevant problem in practice and has motivated some researchers to develop methods that – like ET – use simulated time paths instead of pre-specified grids.¹⁴

3 The model of Coeurdacier, Rey, and Winant (CRW)

We use the algorithm to solve the model considered in Coeurdacier, Rey, and Winant (2011). This is a challenging model to solve, since the model has no steady state and is not well defined when there is no or not a sufficient amount of uncertainty.¹⁵ Moreover, as discussed below, the model can also generate non-stationary behavior if there is too much uncertainty and there are some important nonlinearities.

In the CRW model, an infinitely-lived agent decides how much to consume, c_t , and how much funds to invest and carry over to the next period, w_t . Both the agent’s income, y_t , and the rate of return, r_t , are exogenous stochastic processes. The solution to the agent’s problem is characterized by the following set of equations:

$$c_t^{-\gamma} = \beta \mathbb{E}_t \left[c_{t+1}^{-\gamma} r_{t+1} \right], \quad (21)$$

$$c_t + w_t = y_t + w_{t-1} r_t, \quad (22)$$

$$y_t = (1 - \rho_y) \bar{y} + \rho_y y_{t-1} + \varepsilon_{y,t}, \quad \varepsilon_{y,t} \sim N(0, \sigma_y^2), \quad (23)$$

$$r_t = (1 - \rho_r) \bar{r} + \rho_r r_{t-1} + \varepsilon_{r,t}, \quad \varepsilon_{r,t} \sim N(0, \sigma_r^2). \quad (24)$$

there is no solution for the particular functional form chosen for the approximation. But even if neither is the case, then it still may be the case that some sets of state variables are fatally problematic at some point *in the process of* finding the numerical solution.

¹⁴Simulation-projection methods are not new. An example is the Parameterized Expectations Algorithm of Den Haan and Marcet (1990). As pointed out in Judd (1992) and documented in Den Haan (1995), those earlier simulation approaches can be quite inefficient in that accuracy of the projection’s part of the algorithm, i.e., the regression phase, requires a very large number of observations. Judd, Maliar, and Maliar (2011) develop a simulation-based method with a more efficient projection’s element.

¹⁵Standard projection methods can be used to solve this model, but it is not a trivial problem for these methods either. Since uncertainty is essential, one cannot follow standard practice of first solving a model with little uncertainty and then gradually moving to the case with the desired levels of uncertainty.

If there would be no uncertainty, the Euler equation reduces to

$$c_{t+1} = (\beta\bar{r})^{1/\gamma} c_t, \quad (25)$$

which implies that consumption would either increase without bound or go to zero unless $\beta = 1/\bar{r}$. That is, there is no steady-state solution when $\sigma_y = \sigma_r = 0$.¹⁶ In the CRW model, there are no constraints on w_t and short positions of any size are allowed. Thus, uncertainty is key for the ability of the model to generate a stationary solution.¹⁷

Natural borrowing constraint. For the projections algorithm to work smoothly, we found it necessary to impose the natural borrowing constraint. That is, the maximum amount that an agent can borrow is such that their income will always be enough to cover interest payments, even in the worst possible circumstances. In particular, when solving for the projections solution, we impose that

$$w \geq -\frac{y_{\text{low}}}{r_{\text{high}} - 1}, \quad (26)$$

where y_{low} is the lowest possible value for y_t and r_{high} the highest possible value for r_t .¹⁸ For the parameters considered, this constraint turns out to be never binding, even in the very long simulations considered. As discussed below, this does not imply that it does not affect the behavior of the agent's choices. Neither the CRW algorithm nor the ET algorithm impose the natural borrowing constraint. Nevertheless, as discussed below, the ET solution is characterized by nonlinear behavior that resembles the nonlinearity induced by the natural borrowing constraint.

Parameter values. The parameter values are reported in table 1. The two key parameters are the standard deviations of the innovations to output and the rate of return. As

¹⁶When $\beta = 1/\bar{r}$, then there would be a steady state with $c_t = \bar{y} + (r - 1)w_0$, but any amount of uncertainty would cause the solution to be non-stationary.

¹⁷In appendix A, we provide some intuition for this statement. A more formal analysis of the model's properties can be found in Chamberlain and Wilson (2000).

¹⁸We set y_{low} equal to \bar{y} minus four standard deviations of the unconditional standard deviation of y_t and we set r_{high} equal to \bar{r} plus four standard deviations of the unconditional standard deviation of r_t .

documented below, the model variables are quite volatile at the chosen values. Nevertheless, the values for σ_r and σ_y used here are lower than those chosen in Coeurdacier, Rey, and Winant (2011). We do not use the original CRW values for our comparison, because our accurate projection algorithm indicates that the model solution is not well behaved at those values. As documented in appendix E, it displays "escape dynamics". That is, the simulated time paths for wealth will occasionally – but for sustained time periods – diverge to *extremely* high values.¹⁹ The CRW method imposes stationarity even when it is not correct to do so. Although less restrictive, the ET algorithm also does not seem capable of capturing such explosive behavior.²⁰ Therefore, we chose more moderate volatility levels at which the generated series do not display such divergent behavior. The case with the original CRW parameter values is discussed in more detail in appendix E. We follow CRW and choose the value of \bar{r} such that the risky-steady-state value for wealth is equal to 0.²¹

4 Evaluation of solution methods

In this section, we compare the solutions obtained with the CRW and ET algorithms with an accurate projections-method solution.²² The CRW procedure is identical to the one used in Coeurdacier, Rey, and Winant (2011). The procedure consist of deriving a second-order approximation of the Euler equation and a perturbation solution that are consistent with each other.²³ To implement the ET procedure, we use a second-order approximation

¹⁹Our projection method imposes an upper bound on wealth. Without imposing such an upper bound, the solution diverges to infinity. Outside the (finite) grid, solutions can be obtained by extrapolation, but these cannot be relied on to be accurate. Our experience indicates that the simulated wealth series reach this upper bound *even* when it is set at extremely high values for a sufficiently long sample. The stochastic innovations for r_t and y_t have a normal distribution and could in principle take on extremely high values. This is not the cause of the divergence, since the same explosive behavior is observed when the values of r_t and y_t are restricted to be in a bounded set.

²⁰With the second-order version, the method fails to find a solution at some point during the simulation. It might still be possible that a higher-order version of the algorithm would be successful.

²¹Although, the original system of equations does not have a steady state, the second-order approximation used by CRW does have a steady state, which is referred to as the risky steady state.

²²The accuracy of the projections-method solution is documented in appendix B.

²³A description of the CRW procedure is given in appendix C.

to describe next period's choices.²⁴ The reason we do not use a first-order approximation is that the ET system of equations does not have a solution at some points in the sample when only a first-order approximation is used.

Policy functions. Figure 1 plots the policy functions for the change in wealth, $w - w_{-1}$, as a function of beginning of period wealth, w_{-1} , for different values of the realized rate of return and one particular value of income.²⁵ The top panel displays the results for both the ET and the accurate projection solution. The solutions of these two algorithms are very close for most values of w_{-1} . The only noticeable differences occur at values of beginning-of-period wealth close the natural borrowing constraint for higher values of the realized rate of return. The projection solution indicates that $w - w_{-1}$ gets less negative rapidly as w_{-1} is close to the natural borrowing constraint. The ET solution only uses local information and it does not impose the natural borrowing constraint at all. It is, therefore, remarkable that the ET solution displays a similar nonlinearity close to the natural borrowing constraint.²⁶

The figure documents another – more important – nonlinearity: The slope of the policy function for $w - w_{-1}$ as a function of w_{-1} depends on the level of this period's realized rate of return, r . The slope is positive at high rates of return and negative at low rates of return. The level of r does not only affect this period's resources through rw_{-1} , it also affects the expected rate of return. Another way to think of the nonlinearity associated with r is the following: If r increases then $w - w_{-1}$ increases for higher levels of wealth and decreases for lower levels of beginning-of-period wealth, w_{-1} .²⁷

The bottom panel displays the results for the CRW solution. Although the CRW

²⁴Details of the implementation of the ET algorithm are given in appendix D.

²⁵The results are very similar for other income values.

²⁶Although ET does not impose the natural borrowing constraint, it does avoid choosing wealth levels which could imply very low consumption levels next period, which becomes more likely at very low wealth levels.

²⁷If r increases, then a saver will have more funds in the current period and a borrower will have less resources. The realization of r also affects the expected return. The associated income effect also differs for savers and lenders.

solution is quite accurate for the middle value of r it predicts that $\partial w/\partial w_{-1}$ does not depend on r . As discussed above, even the sign of $\partial w/\partial w_{-1}$ depends on the level of r according to the accurate projection solution (and the ET solution).

Figure 2 repeats the exercise for the consumption policy function. The ET policy function accurately captures that the slope of the consumption function depends on the level of r . The CRW policy function does not. Moreover, it overpredicts the consumption choice at high rates of return and underpredicts the consumption choice at low rates of return, especially at lower wealth levels. The nonlinearity associated with the natural borrowing constraint is less pronounced in the consumption policy than in the wealth policy function.

Simulated time paths. Next we compare simulated time paths. These will reveal possible differences and similarities in exactly those parts of the state space that matter for model properties. The model variables are quite persistent, so we use a long simulation of 80,000 observations. Panels A and B of figure 3 plot the generated values for wealth and consumption, respectively. The figure shows that the ET solutions for consumption and wealth follow the accurate projection solutions closely. That is clearly not the case for the CRW solution. These figures also document that the series are quite volatile even though the chosen standard deviations are less than those used in Coeurdacier, Rey, and Winant (2011).²⁸

Figures 4 and 5 display wealth and consumption time paths for four shorter samples taken from the full sample. These figures show in greater detail that the behavior of the CRW solution can be quite different from the accurate projection solution. For example, in the top right panel of figure 4, the wealth series of all three algorithms display a similar downward trend in the beginning of the sample. While this downward trend continues according to the projection and ET solutions, the CRW solution displays a remarkable recovery. In the top left panel, the projection and ET solutions display a sharp increase which is not present in the perturbation solution.

²⁸To put the volatility in perspective, note that the mean income level is equal to 1.

Comovement of CRW & ET time paths with accurate projection solution.

Table 2 reports how close the ET and CRW solutions are to the accurate projection solution using average and maximum absolute deviations as well as correlation coefficients. Whereas the correlation between the time series generated with the ET algorithm and those generated by the accurate projection solution are virtually equal to 1, they are considerably lower for the CRW series. For example, the correlation between the CRW wealth series and the projection wealth series is only 0.729. Despite the almost perfect correlation, the ET data are not exactly equal to their projection equivalent. In particular, there is an average absolute difference of 0.83% for wealth and 0.21% for consumption. These compare very favorable to the outcomes for CRW, which are 52.3% and 13%. The results are similar for the maximum differences from the accurate projection solution. For the ET data, the maximum differences are 3.62% for wealth and 0.92% for consumption, whereas the corresponding numbers are 268.2% and 68.4% for the CRW solution.

Comparison of generated moments. Tables 3 and 4 report moments according to the three different solutions for the generated wealth and consumption data, respectively. The statistics based on the projection and the ET algorithm are very close to each other. This is not true for the CRW statistics. In particular, the data generated by the CRW algorithm are substantially less volatile; the standard deviation of wealth is only 53% of the projection analogue and the standard deviation of consumption is only 54% of the projection analogue.

A The role of uncertainty for stationarity

The following second-order approximation of the Euler equation provides some intuition for the result that the solutions for c_t and w_t could be stationary because of uncertainty.

$$\begin{aligned}
 c_t^{-\gamma} &\approx \beta \mathbb{E}_t [r_{t+1}] \mathbb{E}_t [c_{t+1}]^{-\gamma} \\
 &+ \\
 &\gamma (\gamma + 1) \beta \mathbb{E}_t [r_{t+1}] \mathbb{E}_t [c_{t+1}]^{-\gamma-2} \mathbb{V}_t (c_{t+1}) \\
 &- \beta \gamma \mathbb{E}_t [c_{t+1}]^{-\gamma-1} \mathbb{C}_t (c_{t+1}, r_{t+1}),
 \end{aligned} \tag{27}$$

where

$$\mathbb{V}_t (c_{t+1}) \equiv \mathbb{E}_t [(c_{t+1} - \mathbb{E}_t [c_{t+1}])^2], \tag{28}$$

$$\mathbb{C}_t (c_{t+1}, r_{t+1}) \equiv \mathbb{E}_t [(c_{t+1} - \mathbb{E}_t [c_{t+1}]) (r_{t+1} - \mathbb{E}_t [r_{t+1}])]. \tag{29}$$

First, consider the case when the rate of return is constant, that is, $\varepsilon_{r,t} = 0 \forall t$. This implies that $\mathbb{C}_t (c_{t+1}, r_{t+1}) = 0$. Also, suppose that the discount rate exceeds the average rate of return, which – by itself – would induce agents to run down their wealth and then take on debt. Uncertainty increases the right-hand side of the Euler equation and dampens the desire of agents to consume more in this period than the next and to reduce wealth levels. Would this mean that the inequality is reversed and that it is optimal to save a lot and increase future consumption levels. The answer is no. The magnitude of the uncertainty effect depends on the value of $\mathbb{E}_t [c_{t+1}]^{-\gamma-2}$. Consequently, expected consumption is high enough and $\mathbb{E}_t [c_{t+1}]^{-\gamma-2}$ low enough, then the uncertainty effect is less important.

If the covariance between unexpected changes in c_{t+1} and r_{t+1} is negative—that is, bonds act as a hedge—then the covariance has a similar dampening role. However, if this covariance is positive, then stochastic rates of return would reinforce the desire of the agent to keep on reducing consumption over time.

B Accuracy of the projections solution

In this section, we document that the solution generated by the projection algorithm is accurate. It is important to establish this, since the projection solution is used as the

benchmark against which the other solutions are compared.

The projection solution uses 8,405,000 grid points; 41 for both r and y and 5,000 for w_{-1} . We use the Dynamic Euler-equation accuracy test, described in Den Haan (2010). The test compares the generated time series with the series that are obtained by explicitly solving the first-order conditions, using the approximation only to calculate next period's wealth choice.²⁹ This test is more stringent than the standard Euler-equation accuracy test, since it would detect if small errors accumulate in a simulation. Figure 7 documents that the two series are very close to each other. Table 5 reports some key statistics. Both the figure and the table indicate that the solution is very accurate.

C The CRW procedure

The CRW procedure starts out with the second-order approximation of the Euler equation, which is given by

$$c_t^{-\gamma} = \beta \begin{bmatrix} \mathbb{E}_t[r_{t+1}] \mathbb{E}_t[c_{t+1}]^{-\gamma} \\ +\gamma(\gamma+1)\beta \mathbb{E}_t[r_{t+1}] \mathbb{E}_t[c_{t+1}]^{-\gamma-2} \mathbb{V}_t(c_{t+1}) \\ -\beta\gamma \mathbb{E}_t[c_{t+1}]^{-\gamma-1} \mathbb{C}_t(c_{t+1}, r_{t+1}) \end{bmatrix}, \quad (30)$$

where $\mathbb{V}_t(c_{t+1})$ and $\mathbb{C}_t(c_{t+1}, r_{t+1})$ are the variance of c_{t+1} and the covariance of c_{t+1} and r_{t+1} , respectively, both conditional on period- t information. In addition, the budget constraint, and the laws of motion for y_t and r_t are needed. That is

$$c_t + w_t = y_t + w_{t-1}r_t, \quad (31)$$

$$y_t = (1 - \rho_y)\bar{y} + \rho_y y_{t-1} + \varepsilon_{y,t}, \quad \varepsilon_{y,t} \sim N(0, \sigma_y^2), \quad (32a)$$

$$r_t = (1 - \rho_r)\bar{r} + \rho_r r_{t-1} + \varepsilon_{r,t}, \quad \varepsilon_{r,t} \sim N(0, \sigma_r^2). \quad (32b)$$

²⁹In particular, the numerical solution is not used to determine the amount of wealth carried over into the next period.

The objective of the CRW procedure is to find an approximation to the savings decision that is linear in levels. The latter is given by

$$w_t = \bar{w} + W_w \hat{w}_{t-1} + W_y \hat{y}_t + W_r \hat{r}_t. \quad (33)$$

Using the budget constraint and this linear approximation for w_t , one can derive the second-order terms in the approximated Euler equation. The results are as follows:

$$\mathbb{V}_t(c_{t+1}) = \left((1 - W_y)^2 \sigma_y^2 \right) + (W_r^2 + w_t^2 - 2w_t W_r) \sigma_r^2, \text{ and} \quad (34)$$

$$\mathbb{C}_t(c_{t+1}, r_{t+1}) = (w_t - W_r) \sigma_r^2. \quad (35)$$

The idea of the CRW algorithm is to find a first-order perturbation solution to the system of equations consisting of equations (30), (31), (32), (34), and (35), which is consistent with the definitions of $\mathbb{V}_t(c_{t+1})$ and $\mathbb{C}_t(c_{t+1}, r_{t+1})$ given in equations (34) and (35). One could find the values for \bar{w} , W_r , W_y using an equation solver or with the following iterative procedure. Start with initial values for W_r and W_y and denote these by \widetilde{W}_r and \widetilde{W}_y . If we use these values to calculate $\mathbb{V}_t(c_{t+1})$ and $\mathbb{C}_t(c_{t+1}, r_{t+1})$ and substitute the expressions for $\mathbb{V}_t(c_{t+1})$ and $\mathbb{C}_t(c_{t+1}, r_{t+1})$ into equation (30), we get the CRW system of equations:

$$\begin{aligned} & c_t^{-\gamma} \\ & = \\ & \left[\begin{array}{c} \beta \mathbb{E}_t[r_{t+1}] \mathbb{E}_t[c_{t+1}]^{-\gamma} \\ +\gamma(\gamma+1)\beta \mathbb{E}_t[r_{t+1}] \mathbb{E}_t[c_{t+1}]^{-\gamma-2} \left(\begin{array}{c} \left((1 - \widetilde{W}_y)^2 \sigma_y^2 \right) \\ + \left(\widetilde{W}_r^2 + w_t^2 - 2w_t \widetilde{W}_r \right) \sigma_r^2 \end{array} \right) \\ -\beta\gamma \mathbb{E}_t[c_{t+1}]^{-\gamma-1} \left(w_t - \widetilde{W}_r \right) \sigma_r^2 \end{array} \right], \end{aligned} \quad (36)$$

$$c_t + w_t = y_t + r_t w_{t-1}, \quad (37)$$

$$y_t = (1 - \rho_y) \bar{y} + \rho_y y_{t-1} + \varepsilon_{y,t}, \quad \varepsilon_{y,t} \sim N(0, \sigma_y^2), \quad (38)$$

$$r_t = (1 - \rho_r) \bar{r} + \rho_r r_{t-1} + \varepsilon_{r,t}, \quad \varepsilon_{r,t} \sim N(0, \sigma_r^2). \quad (39)$$

For given values of \widetilde{W}_r and \widetilde{W}_y , this system of equations has a steady state and regular perturbation techniques can be used to calculate \widetilde{w} , W_r , and W_y . The values for W_r and W_y can then be used to update equation (36). This process is continued until convergence.

D Details of the ET procedure

In this section, we give detailed information about the implementation of the ET algorithm. The discussion follows the Matlab program closely.³⁰ To increase transparency, we replace some of the names used in the program with more informative mathematical symbols.³¹

Define key variables and preliminaries

- The Gauss-Hermite nodes and weights (scaled by $\sqrt{\pi}$) are stored in the $(N \times 1)$ vectors X and P , respectively. We assume that the number of nodes is the same for both random variables.
- The following list of variables are declared to be symbolic variables: w , r , w_{-1} , y , \widetilde{r} , \widetilde{w} , \widetilde{y} , ϕ_0 , ϕ_r , ϕ_w , ϕ_y , ϕ_{rr} , ϕ_{ww} , ϕ_{yy} , ϕ_{rw} , ϕ_{ry} , ϕ_{wy} . The last six are only used when implementing second-order ET.
- Define next period's realizations of the two random variables, that is, r_{+1} and y_{+1} , as functions of the possible outcomes, i.e., the Gauss-Hermite nodes:

$$r_{+1} = @ (X) \quad r_{+1} = (1 - \rho_r) \bar{r} + \rho_r r + \sqrt{2} \sigma_r X, \quad (40)$$

$$y_{+1} = @ (X) \quad y_{+1} = (1 - \rho_y) \widetilde{y} + \rho_y y + \sqrt{2} \sigma_y X. \quad (41)$$

Note that r_{+1} and y_{+1} are not just functions of X , but also functions of r and y .

³⁰The program is available at www.wouterdenhaan.com/software.htm.

³¹In particular, `lhs` in the code stands for L , i.e., the left-hand side of the Euler equation; `wlag` in the code stands for wealth chosen previous period, that is, w_{-1} ; `wconstm` stands for ϕ_0 ; `rm` stands for \widetilde{r} ; `wm` stands for \widetilde{w} ; `ym` stands for \widetilde{y} ; `phi_i` stands for ϕ_i with $i \in \{\mathbf{r}, \mathbf{w}, \mathbf{y}\}$; and `phi_i_j` stands for $\phi_{i,j}$ with $i, j \in \{\mathbf{r}, \mathbf{w}, \mathbf{y}\}$; `eq` stands for $Q(\cdot)$, the Euler-equation error term; `eq_x` stands for $\partial Q / \partial x$ with $\mathbf{x} \in \{\mathbf{w}, \mathbf{r}, \mathbf{w}, \mathbf{y}\}$; `eq_xz` stands for $\partial^2 Q / \partial^2 xz$ with $\mathbf{x}, \mathbf{z} \in \{\mathbf{w}, \mathbf{r}, \mathbf{w}, \mathbf{y}\}$. In both the first and second order derivatives, the modifier w stands for previous period wealth, w_{-1} , and the modifier W stands for this period wealth, w .

- Define the left-hand side of the Euler equation, that is, $c^{-\gamma}$:

$$L = (y + rw_{-1} - w)^{-\gamma}. \quad (42)$$

- The key variable is the Euler equation error, which is equal to

$$\begin{aligned} Q_0(w, w_{-1}, r, y; \Phi) \\ = \\ -1 + \frac{\beta}{L} \sum_{i=1}^N \sum_{j=1}^N r_{+1}(X) \begin{bmatrix} y_{+1}(X(j)) + r_{+1}(X(i))w \\ -\tilde{h}(r_{+1}, w, y_{+1}; \Phi) \end{bmatrix}^{-\gamma} P(i)P(j), \end{aligned} \quad (43)$$

where Φ contains all the coefficients of the approximating functional form. In the program, $Q(\cdot)$ is calculated using a double "for loop".

First-order ET. When using first-order ET, the function $\tilde{h}(\cdot)$ is given by

$$\tilde{h}(r_{+1}, w, y_{+1}; \Phi) = \phi_0 + \phi_r(r_{+1}(X) - \tilde{r}) + \phi_w(w - \tilde{w}) + \phi_y(y_{+1}(X) - \tilde{y}) \quad (44)$$

where

$$\Phi = [\phi_0, \phi_r, \phi_w, \phi_y, \tilde{r}, \tilde{w}, \tilde{y}]. \quad (45)$$

A key aspect of ET is that all elements of Φ depend on the current values of r , w_{-1} , and y .

The model says that the error term defined in equation (43) should be equal to zero for all r , w_{-1} , and y , that is,

$$Q_0(w, w_{-1}, r, y; \Phi) = 0. \quad (46)$$

Differentiating this equation with respect to r , w_{-1} , and y gives

$$\frac{\partial Q_0}{\partial w} \frac{\partial w}{\partial r} + \frac{\partial Q_0}{\partial r} = 0, \quad (47a)$$

$$\frac{\partial Q_0}{\partial w} \frac{\partial w}{\partial w_{-1}} + \frac{\partial Q_0}{\partial w_{-1}} = 0, \quad (47b)$$

$$\frac{\partial Q_0}{\partial w} \frac{\partial w}{\partial y} + \frac{\partial Q_0}{\partial y} = 0, \quad (47c)$$

With a symbolic toolbox, one defines Q_0 as a function, exactly as it is defined in equation (43). Next, one obtains the partial derivatives as functions of w , w_{-1} , r , and y by using the symbolic toolbox' differentiation command.³² After the derivatives have been calculated, one would impose that

$$\phi_0 = w, \quad (48)$$

$$\tilde{r} = r, \quad (49)$$

$$\tilde{w} = w_{-1}, \quad (50)$$

$$\tilde{y} = y. \quad (51)$$

That is, the Taylor series expansion to characterize next period's choice is around today's state variables and the function value at those outcomes is, of course, equal to this period's choice. The derivatives of this Taylor series expansion are equal to the derivatives of this period's choice. That is,

$$\phi_r = \frac{\partial w}{\partial r}, \quad (52a)$$

$$\phi_w = \frac{\partial w}{\partial w_{-1}}. \quad (52b)$$

$$\phi_y = \frac{\partial w}{\partial y}. \quad (52c)$$

After these substitutions, we get the following system of four equations in the four unknowns – w , ϕ_r , ϕ_w , and ϕ_y – for given values of r , w_{-1} , and y .³³

$$0 = Q_0(w, w_{-1}, r, y; \tilde{\Phi}), \quad (53a)$$

$$0 = Q_r(w, w_{-1}, r, y; \tilde{\Phi}), \quad (53b)$$

$$0 = Q_w(w, w_{-1}, r, y; \tilde{\Phi}), \quad (53c)$$

$$0 = Q_y(w, w_{-1}, r, y; \tilde{\Phi}), \quad (53d)$$

³²In Matlab, `Q_r=diff(Q_0,r)` generates the derivative of Q_0 with respect to r .

³³For this system of equations, the nonlinear-equation solver is keen to set ϕ_w to values close to r and ϕ_y to values close to 1. Since $r > 1$, this implies an explosive solution. It is easy to show that $\phi_w = r$ and $\phi_y = 1$ are exact solutions to equations (53c) and (53d) when $\sigma_r = 0$. To avoid the nonlinear-equation solver choosing or moving towards this possibility, we divide equation (53c) by $r - \phi_w$ and equation (53d) by $1 - \phi_y$.

where

$$\tilde{\Phi} = [\phi_r, \phi_w, \phi_y]. \quad (54)$$

Second-order ET. When using second-order ET, the function $\tilde{h}(\cdot)$ is given by

$$\begin{aligned} & \tilde{h}(r_{+1}, w, y_{+1}; \Phi) \\ & = \\ & \phi_0 + \phi_r(r_{+1}(X) - \tilde{r}) + \phi_w(w - \tilde{w}) + \phi_y(y_{+1}(X) - \tilde{y}) \\ & + \frac{1}{2}\phi_{rr}(r_{+1}(X) - \tilde{r})^2 + \frac{1}{2}\phi_{ww}(w - \tilde{w})^2 + \frac{1}{2}\phi_{yy}(y_{+1}(X) - \tilde{y})^2 \\ & + \phi_{rw}(r_{+1}(X) - \tilde{r})(w - \tilde{w}) + \phi_{ry}(r_{+1}(X) - \tilde{r})(y_{+1}(X) - \tilde{y}) \\ & + \phi_{wy}(w - \tilde{w})(y_{+1}(X) - \tilde{y}) \end{aligned} \quad (55)$$

where

$$\Phi = [\phi_0, \phi_r, \phi_w, \phi_y, \phi_{rr}, \phi_{ww}, \phi_{yy}, \phi_{rw}, \phi_{ry}, \phi_{wy}, \tilde{r}, \tilde{w}, \tilde{y}]. \quad (56)$$

The additional six equations and associated error terms are given by

$$0 = Q_{ij}(w, w_{-1}, r, y; \Phi) \quad i, j \in \{r, w, y\} \quad (57)$$

$$Q_{rr}(w, w_{-1}, r, y; \Phi) = \frac{\partial Q}{\partial w} \phi_{rr} + 2 \frac{\partial^2 Q}{\partial w \partial r} \phi_r + \frac{\partial^2 Q}{\partial w^2} \phi_r^2 + \frac{\partial^2 Q}{\partial r^2}, \quad (58)$$

$$Q_{ww}(w, w_{-1}, r, y; \Phi) = \frac{\partial Q}{\partial w} \phi_{ww} + 2 \frac{\partial^2 Q}{\partial w \partial w_{-1}} \phi_w + \frac{\partial^2 Q}{\partial w^2} \phi_w^2 + \frac{\partial^2 Q}{\partial w_{-1}^2}, \quad (59)$$

$$Q_{yy}(w, w_{-1}, r, y; \Phi) = \frac{\partial Q}{\partial w} \phi_{yy} + 2 \frac{\partial^2 Q}{\partial w \partial y} \phi_y + \frac{\partial^2 Q}{\partial w^2} \phi_y^2 + \frac{\partial^2 Q}{\partial y^2}, \quad (60)$$

$$Q_{rw}(w, w_{-1}, r, y; \Phi) = \left[\begin{array}{c} \frac{\partial Q}{\partial w} \phi_{rw} + \frac{\partial^2 Q}{\partial w \partial w_{-1}} \phi_r + \frac{\partial^2 Q}{\partial w \partial r} \phi_w \\ + \frac{\partial^2 Q}{\partial w^2} \phi_r \phi_w + \frac{\partial^2 Q}{\partial r \partial w_{-1}} \end{array} \right], \quad (61)$$

$$Q_{ry}(w, w_{-1}, r, y; \Phi) = \left[\begin{array}{c} \frac{\partial Q}{\partial w} \phi_{ry} + \frac{\partial^2 Q}{\partial w \partial y} \phi_r + \frac{\partial^2 Q}{\partial w \partial r} \phi_y \\ + \frac{\partial^2 Q}{\partial w^2} \phi_r \phi_y + \frac{\partial^2 Q}{\partial r \partial y} \end{array} \right], \quad (62)$$

$$Q_{wy}(w, w_{-1}, r, y; \Phi) = \left[\begin{array}{c} \frac{\partial Q}{\partial w} \phi_{wy} + \frac{\partial^2 Q}{\partial w \partial y} \phi_w + \frac{\partial^2 Q}{\partial w \partial w_{-1}} \phi_y \\ + \frac{\partial^2 Q}{\partial w^2} \phi_w \phi_y + \frac{\partial^2 Q}{\partial w_{-1} \partial y} \end{array} \right]. \quad (63)$$

These equations are obtained by differentiating equation (43) twice and imposing that

$$\begin{aligned} \phi_{rr} &= \frac{\partial^2 w}{\partial r^2}, \phi_{ww} = \frac{\partial^2 w}{\partial w_{-1}^2}, \phi_{yy} = \frac{\partial^2 w}{\partial y^2}, \\ \phi_{rw} &= \frac{\partial^2 w}{\partial r \partial w_{-1}}, \phi_{ry} = \frac{\partial^2 w}{\partial r \partial y}, \phi_{wy} = \frac{\partial^2 w}{\partial w_{-1} \partial y}. \end{aligned} \quad (64)$$

E Results for other parameter values

The analysis in the main text is based on values for the standard deviations of the innovations that are lower than the ones used in Coeurdacier, Rey, and Winant (2011). The reason for our choice of parameter values is that the model generates very complex non-linear "escape" dynamics at the original CRW parameter values. More precisely, we find that the solution for wealth generated by an accurate projections method hits the imposed upper bound on wealth when the sample is long enough even when that upper bound is extremely high.³⁴ This is documented in figure 6, which plots parts of the projection and CRW time paths at the original CRW parameter values. The figure has three panels, each corresponding with a different upper bound on wealth imposed by the projection algorithm. The three values considered are 35, 100, and 1000. The figure documents that increasing the upper bound is not effective in making it not binding. In fact, if the upper bound is increased from 100 to 1000 it actually becomes binding more often in this part of the state space.

Escape dynamics occur at high values of the rate of return, but it is not due to a few extreme values for the rate of return, since the same pattern of results is found when the realizations of r_t are constrained not to be further away from their mean than two standard deviations. As documented in figure 6, the CRW solution does not capture these escape dynamics. The ET algorithm also has difficulty obtaining solutions at many points in the state space.

This type of behavior indicates that the model may not have a stationary solution at those parameter values unless one does impose an upper bound.³⁵ Projection methods can incorporate such a bound, but this is not possible for CRW and ET. Therefore, we prefer not to do a formal comparison of different solution methods at parameter values

³⁴At higher values for σ_r and σ_y , it is important for the stability of the projection algorithm that an upper bound is imposed.

³⁵It is possible that there is a stationary solution despite the escape dynamics, but that the upper bounds considered here are still not high enough. Imposing the upper bound flattens the policy function for wealth at high interest rates even at values of wealth which are not that close to the upper bound, which makes it even more surprising that the upper bound is reached during simulations.

where we observe this type of escape dynamics.

F The neoclassical growth model

In this section, we make clear how easy it is to program the ET algorithm by providing the Matlab code that uses first-order ET to solve the neoclassical growth model, a model that most readers will be familiar with. The program uses the Matlab symbolic toolbox to calculate derivatives. The user also needs a subroutine to generate the Gauss-Hermite quadrature nodes and weights. These are readily available and we provide one at www.wouterdenhaan.com/software.htm. This website also contains the code displayed here and the program to solve the same model with second-order ET.

```
clear;
% This program solves the stochastic Ramsey growth model using the
% ET-algorithm by Den Haan, Kobielarz, and Rendahl (2015).
% Declare parameters for the model

alpha = 1/3;          % Capital share of income: 1/3.
beta = 1.04^(-1/4);  % Discount factor: 4% annual interest rate.
delta = 0.025;       % Depreciation rate.
gamma = 10;          % Coefficient of relative risk aversion.
sigma = 0.03;        % Standard deviation of productivity shock.
rho = 0.95;          % Persistence of technology shock.

% Implied steady state

kss = ((1/beta+delta-1)/alpha)^(1/(alpha-1));
css = kss^(alpha)-delta*kss;

% Information regarding the shock

N_quad = 5;          % Number of nodes for the quadrature (5 goes a long way).
[Z W] = hernodes(N_quad); % Generate nodes and weights.
```

```

W = pi^(-1/2)*W;      % Normalize weights.
Z = Z*sqrt(2)*sigma;  % Normalize nodes.

% Declar symbolic variables

syms k c z phi_k phi_z kt ct zt

kp = exp(z)*k^(alpha)+(1-delta)*k-c;
zp = rho*z+Z;

% Define the Euler Equation

EE = c-(W'*( beta*(1+exp(zp))*alpha*kp^(alpha-1)-delta) ...
      *(ct+phi_k*(kp-kt)+phi_z*(zp-zt)).^(-gamma) ))^(-1/gamma);

% Take derivatives

dk = diff(EE,k);
dz = diff(EE,z);
dc = diff(EE,c);

% The implicit function theorem reveals that
% the following equations should hold:

Ek = phi_k+dk./dc;
Ez = phi_z+dz./dc;

% The collection of equations to be solved is therefore

E = [EE;Ek;Ez];
E = subs(E,[kt zt ct],[k z c]);

% Convert into a matlab function

E = matlabFunction(E,'vars',{[c;phi_k;phi_z],k,z});

```

```

% Solve the problem along a stochastic simulation.

T = 500;
k0 = zeros(1,T+1);
z0 = zeros(1,T+1);
k0(1) = kss;
z0(1) = 0;
% Initial guess
X = [css;0.9;0.9];
options = optimset('Display','off','TolFun',1e-14,'TolX',1e-14);
rng(20150215,'twister');
e = sigma*randn(T,1);

for t = 1:T
    disp(t)
    E1 = @(x) E(x,k0(t),z0(t));
    X(:,t+1) = fsolve(E1,X(:,t),options);
    k0(t+1) = exp(z0(t))*k0(t)^(alpha)+(1-delta)*k0(t)-X(1,t+1);
    z0(t+1) = rho*z0(t)+e(t);
end
c0 = X(1,:);
% Done.

```

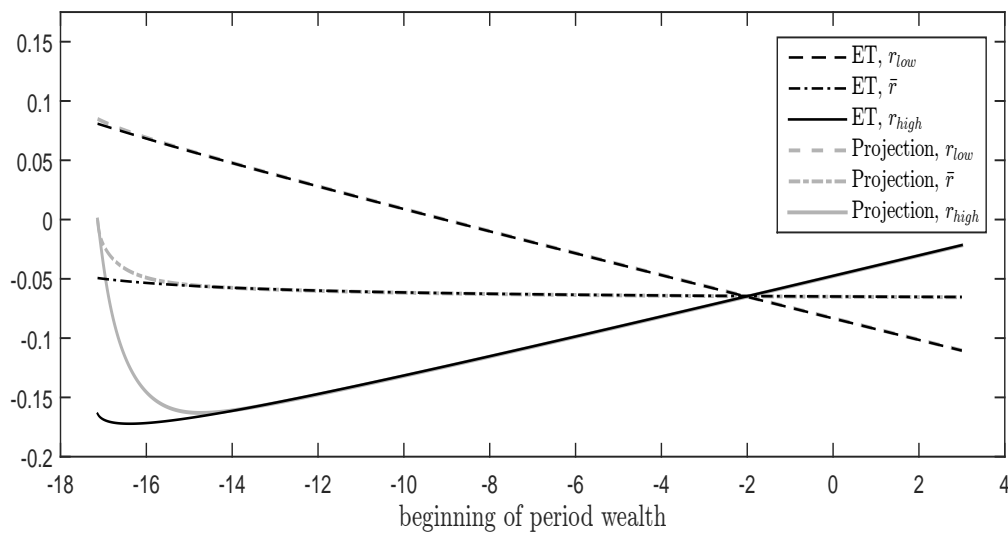
References

- CHAMBERLAIN, G., AND C. WILSON (2000): "Optimal Intertemporal Consumption Under Uncertainty," *Review of Economic Dynamics*, 3, 365–395.
- COEURDACIER, N., H. REY, AND P. WINANT (2011): "The Risky Steady State," *American Economic Review: Papers and Proceedings*, 101, 398–401.

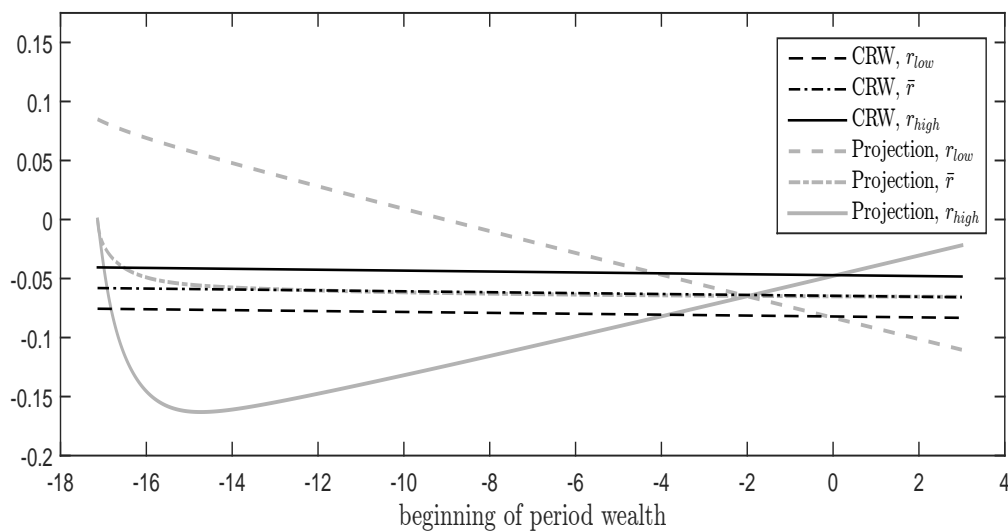
- DEN HAAN, W. J. (1995): “The Term Structure of Interest Rates in Real and Monetary Economies,” *Journal of Economic Dynamics and Control*, 19, 909–940.
- (2010): “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty,” *Journal of Economic Dynamics and Control*, 34, 4–27.
- DEN HAAN, W. J., AND J. DE WIND (2012): “Nonlinear and Stable Perturbation-based Approximations,” *Journal of Economic Dynamics and Control*, 36, 1477–1797.
- DEN HAAN, W. J., AND A. MARCET (1990): “Solving the Stochastic Growth Model by Parameterizing Expectations,” *Journal of Business and Economic Statistics*, 8, 31–34.
- JUDD, K. J., S. MALIAR, AND L. MALIAR (2011): “Numerically Stable and Accurate Stochastic Simulation Approaches for Solving Dynamic Economic Models,” *Quantitative Economics*, 2, 173–210.
- JUDD, K. L. (1992): “Projection Methods for Solving Aggregate Growth Models,” *Journal of Economic Theory*, 58, 410–452.

Figure 1: Wealth Policy Function (low y value)

A: ET and accurate projection solution

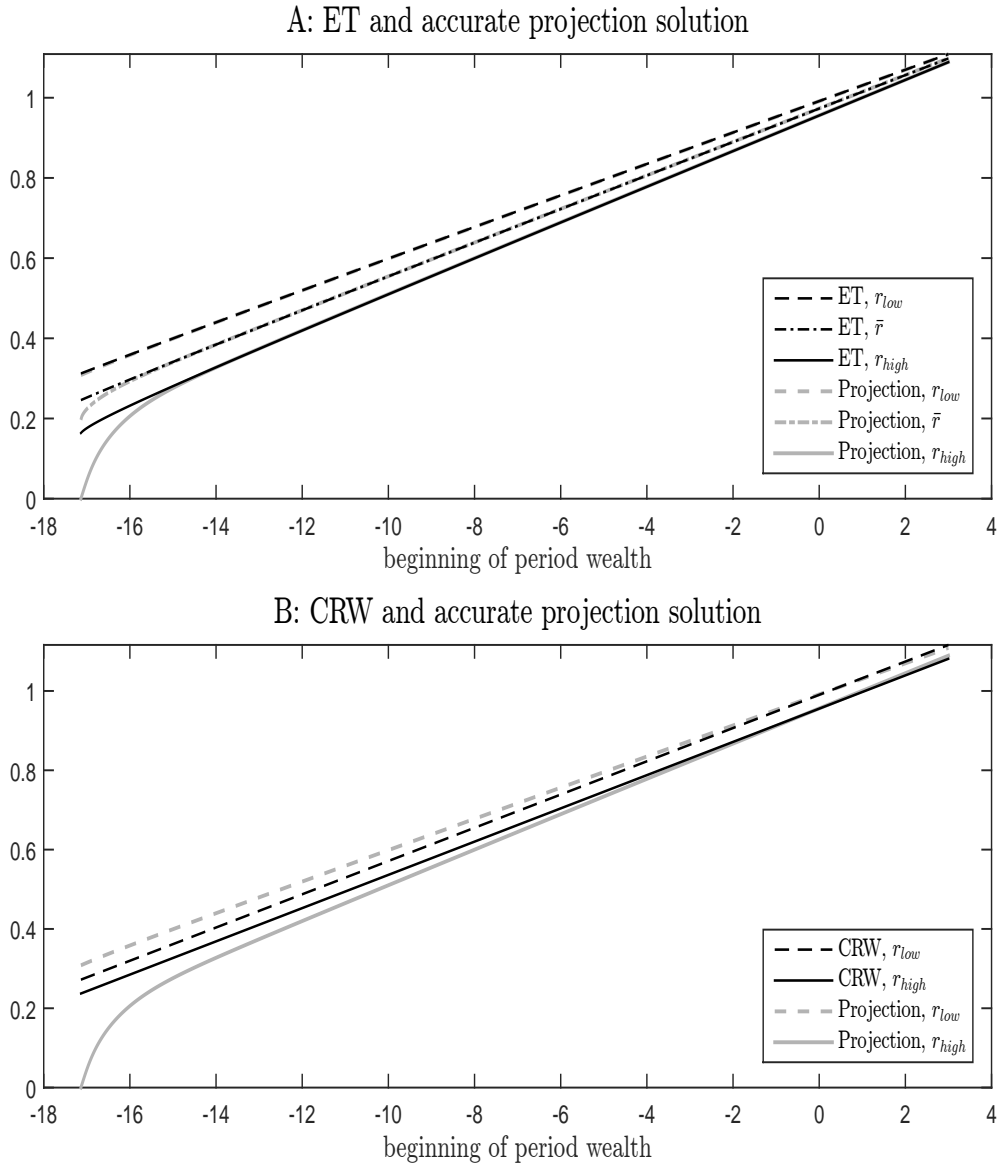


B: CRW and accurate projection solution



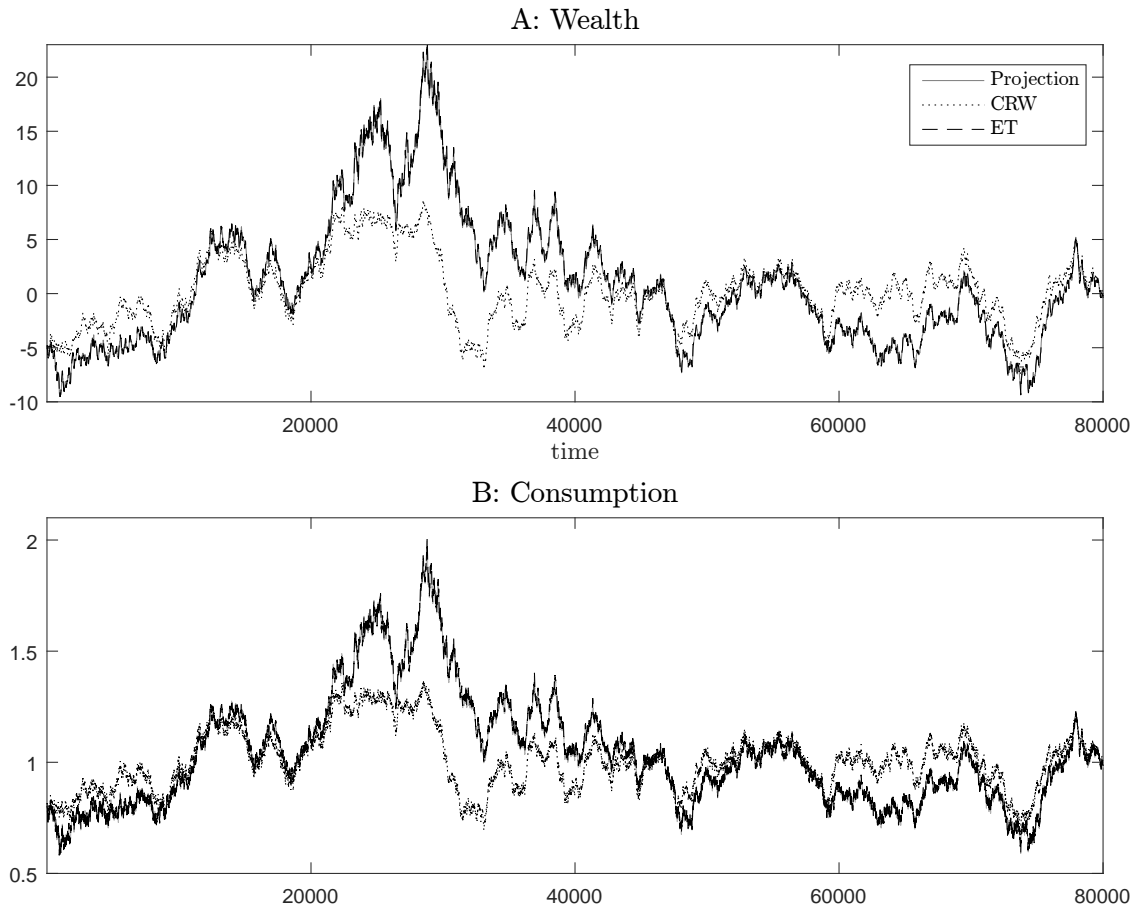
Notes: These graphs plots the values of the change in wealth, $w - w_{-1}$, according to the indicated algorithm as a function of beginning-of-period wealth, w_{-1} , when income, y , is equal to 0.9082. The three values of the realized gross rate of return, r , are equal to 1.0301, 10358, and 1.0415.

Figure 2: Consumption Policy Function (low y value)



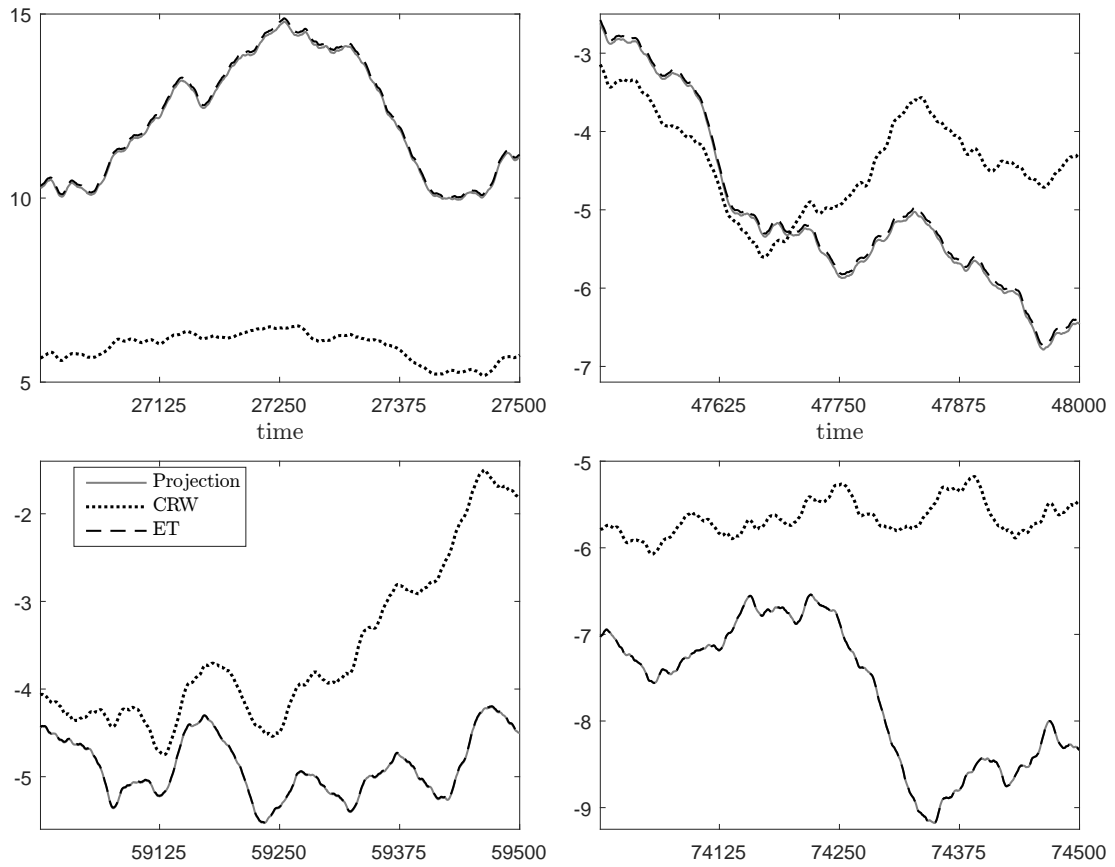
Notes: These graphs plots the values of consumption, c , according to the indicated algorithm as a function of beginning-of-period wealth, w_{-1} , when income, y , is equal to 0.9082. The three values of the realized gross rate of return, r , are equal to 1.0301, 10358, and 1.0415.

Figure 3: Simulated Time Paths (whole sample)



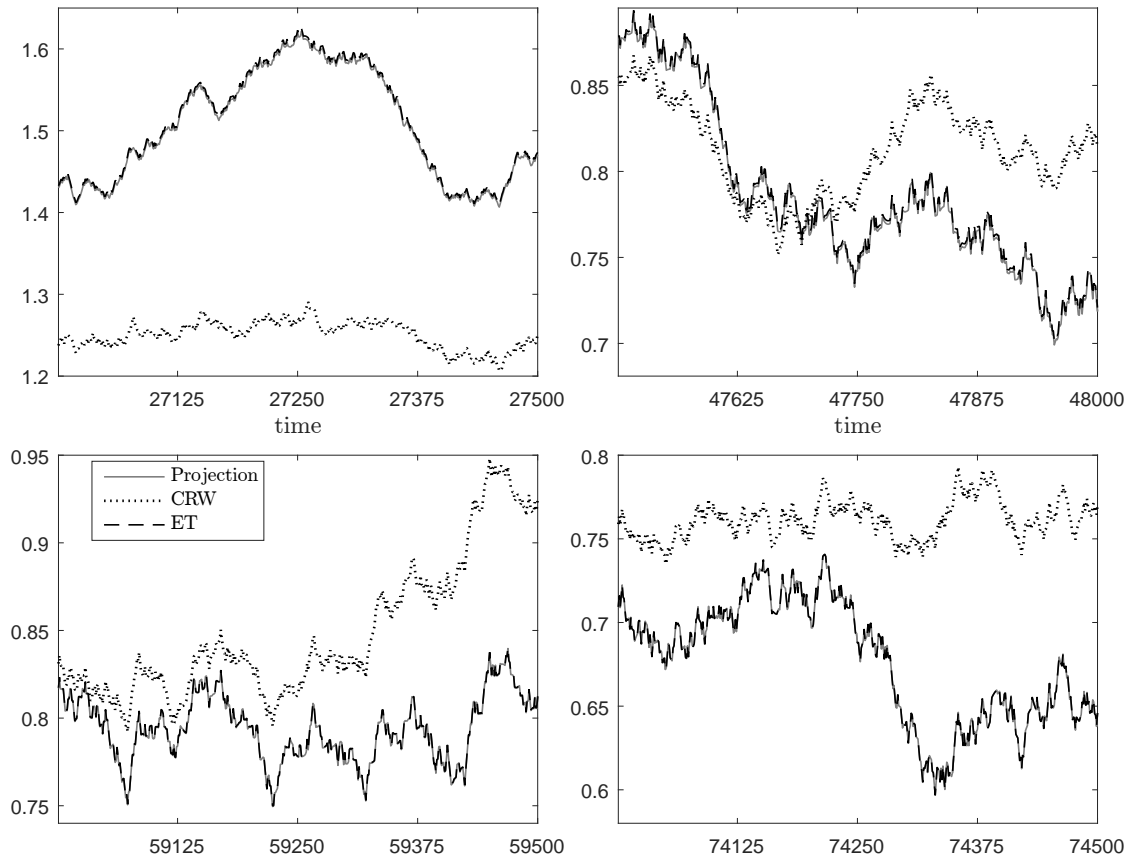
Notes: This graph plots simulated time series generated with the three algorithms using the exact same initial conditions and same draws for the exogenous random variables. Whereas the projection and the ET solution are difficult to distinguish, the CRW solution is clearly different.

Figure 4: Selected Short Samples: Wealth



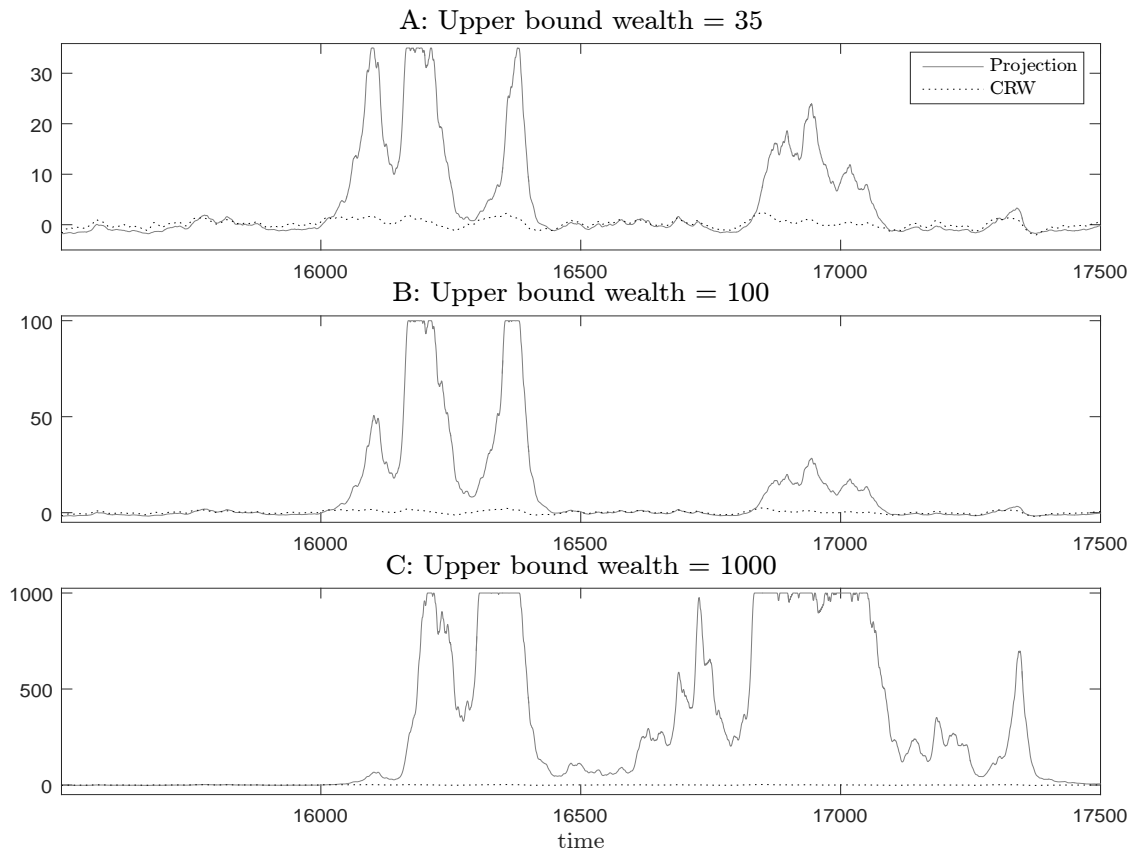
Notes: This graph plots parts of the simulated time series for wealth generated with the three algorithms using the exact same initial conditions and same draws for the exogenous random variables. Whereas the projection and the ET solution are difficult to distinguish, the CRW solution is clearly different.

Figure 5: Selected Short Samples: Consumption



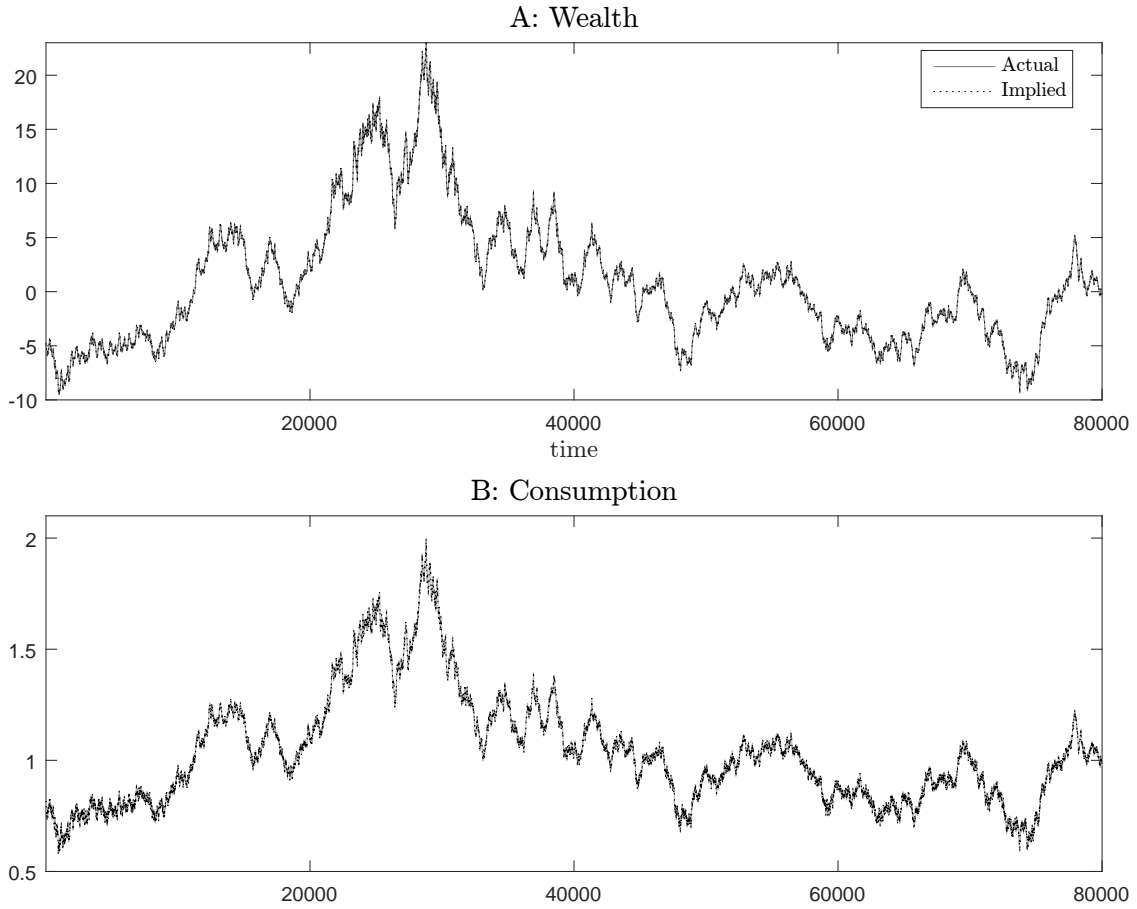
Notes: This graph plots parts of the simulated time series for consumption generated with the three algorithms using the exact same initial conditions and same draws for the exogenous random variables. Whereas the projection and the ET solution are difficult to distinguish, the CRW solution is clearly different.

Figure 6: Simulated Wealth Time Series - Original CRW Parameter Values



Notes: This graph plots simulated time series for wealth generated with the projection and the CRW algorithm using the original CRW parameter values. The projection algorithm imposes the indicated upper bound for wealth.

Figure 7: Dynamic Euler-Equation Accuracy Test



Notes: These graphs plot the actual time series for wealth and consumption and the time series that are the explicit solutions to the first-order conditions using the approximation to calculate next period's choice for wealth. The fact that the two series are difficult to distinguish indicates that the solution is accurate.

Table 1: Parameter Values

σ_r	0.00125	ρ_r	0.9	\bar{r}	1.04152878685	γ	4
σ_y	0.01	ρ_y	0.9	\bar{y}	1		

Table 2: Distance of ET and CRW Outcomes Relative to Accurate Projection Solution

	ET	CRW
correlation of w_t with projection solution	1.000	0.729
correlation of Δw_t with projection solution	1.000	0.798
correlation of HP-filtered w_t with projection solution	1.000	0.790
mean abs. deviation of w_t from projection as fraction of $SD(w_t)$	0.83%	52.3%
max. abs. deviation of w_t from projection as fraction of $SD(w_t)$	3.62%	268.2%
correlation of c_t with projection solution	1.000	0.730
correlation of Δc_t with projection solution	1.000	0.891
correlation of HP-filtered c_t with projection solution	1.000	0.885
mean abs. deviation of c_t from projection as fraction of \bar{y}	0.21%	13.0%
max abs. deviation of c_t from projection as fraction of \bar{y}	0.92%	68.4%

Notes: This table reports different measures to indicate the similarity between data generated by the CRW or the ET algorithm and data generated by an accurate projection algorithm. The statistics are based on a simulation of 80,000 observations. $SD(w_t)$ stands for the standard deviation of wealth.

Table 3: Model Properties According to the Three Algorithms: Wealth

	Projection	ET	CRW
mean of w_t	0.934	0.983	0.198
standard deviation of w_t	5.99	6.02	3.18
standard deviation of Δw_t	0.022	0.022	0.017
standard deviation of HP-filtered w_t	0.041	0.041	0.031
correlation of w_t and y_t	0.033	0.032	0.076
correlation of w_t and r_t	-0.004	-0.004	-0.004
correlation of w_t and c_t	0.999	0.999	0.998
correlation of Δw_t and Δy_t	0.168	0.168	0.216
correlation of Δw_t and Δr_t	0.068	0.069	0.060
correlation of Δw_t and Δc_t	0.375	0.376	0.327
correlation of HP-filtered w_t and y_t	0.118	0.117	0.144
correlation of HP-filtered w_t and r_t	0.048	0.049	0.045
correlation of HP-filtered w_t and c_t	0.432	0.433	0.370

Notes: This table reports model properties for data generated by the indicated algorithm. The results are based on a simulation of 80,000 observations.

Table 4: Model Properties According to the Three Algorithms: Consumption

	Projection	ET	CRW
mean of c_t	1.039	1.041	1.008
standard deviation of c_t	0.250	0.251	0.134
standard deviation of Δc_t	0.0038	0.0038	0.0036
standard deviation of HP-filtered c_t	0.0050	0.0050	0.0046
correlation of c_t and y_t	0.057	0.056	0.121
correlation of c_t and r_t	-0.020	-0.020	-0.037
correlation of Δc_t and Δy_t	0.747	0.747	0.808
correlation of Δc_t and Δr_t	-0.455	-0.452	-0.560
correlation of HP-filtered c_t and y_t	0.723	0.723	0.788
correlation of HP-filtered c_t and r_t	-0.441	-0.439	-0.546

Notes: This table reports model properties for data generated by the indicated algorithm. The results are based on a simulation of 80,000 observations.

Table 5: Accuracy Test

	mean absolute error	maximum absolute error
consumption		
as fraction of \bar{y}	$9.2e - 5$	$4.0e - 4$
as fraction of $SD(w_t)$	$1.5e - 5$	$6.7e - 5$
wealth		
as fraction of \bar{y}	$2.2e - 3$	$9.4e - 3$
as fraction of $SD(w_t)$	$3.7e - 3$	$1.6e - 3$

Notes: This table reports the errors of the dynamic Euler equation accuracy test using a simulation of 80,000 observations.