Supplementary Material

Supplement to: "Altruistically motivated transfers under uncertainty": Computational Appendix

(*Quantitative Economics*, Vol. 5, No. 3, November 2014, 705–749)

Daniel Barczyk
McGill University and CIREQ

Matthias Kredler
Universidad Carlos III

This appendix first illustrates how to use the Markov-chain approximation method to solve a standard consumption–savings problem in continuous time. We then show how to use the method to find the equilibrium of the dynamic Markov game in the paper, and finally show how to extend the model in various dimensions.

In particular, Section S.1 shows how to compute value and policy functions for a single household and how to find the distributions of households over the state space resulting from the optimal policies. We consider both transitory shocks to income (Brownian motion) and persistent risk (a Poisson process) so as to illustrate the method. Section S.2 augments Section S.1 with the elements needed to compute the equilibrium of the game in the paper, in which there are two decision makers who are imperfectly altruistic for each other. Section S.3 then shows how to adapt the altruism model to an overlapping-generations framework and a finite-horizon economy, as well as how to augment the model by an endogenous risk-taking decision or other choices.

`Matlab` code for the single-agent consumption–savings problem as well as the altruism model (and some of its extensions) is available on the journal website; see the respective passages of this appendix for references to the code. The code uses an object with a large number of generalized routines for continuous-time finite-element methods; see `Doc_FinElObj.pdf` for documentation.

S.1. A consumption–savings problem

Consider the problem of a consumer with initial wealth $W_0 \geq 0$, and a stochastic income stream with time-dependent mean y_t and standard deviation σ, who has access to a safe

Daniel Barczyk: daniel.barczyk@mcgill.ca
Matthias Kredler: matthias.kredler@uc3m.es

DOI: 10.3982/QE353

asset.[1] The consumer chooses a consumption function $c(t, W_t)$ to maximize

$$\begin{aligned} &\mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t)\,\mathrm{d}t \\ &\text{s.t.}\quad \mathrm{d}W_t = (rW_t + y_t - c_t)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t, \\ &\qquad c \geq 0, \qquad W \geq 0, \qquad W_0 \text{ given,} \end{aligned}$$

where W stands for wealth, r stands for the return on a risk-free asset, y stands for flow labor income, c stands for consumption, and σ is the per-period standard deviation of income and B_t is standard Brownian motion.[2] Per-period utility is constant relative risk aversion (CRRA), that is, $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$.

The code for the consumption–savings problem can be found in `Bewley_main.m` in the folder `Bewley`.

S.1.1 *Deriving the HJB*

We first heuristically derive the Hamilton–Jacobi–Bellman equation (HJB). This derivation will highlight the connection of the numerical approximation method to the underlying continuous-time problem.

We start by thinking about a discrete-time problem in which time is chopped up into small intervals of size Δt. Using Bellman's principle, we separate the consumption–savings problem into a trade-off between consuming today and saving for tomorrow (a horizon Δt), taking the continuation value V of the problem as of time $t+\Delta t$ as given,

$$\begin{aligned} V(t, W_t) &= \max_{c\geq 0}\left\{u(c)\Delta t + e^{-\rho\Delta t}\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t})\right\} \\ \text{s.t.}\quad W_{t+\Delta t} - W_t &= \underbrace{(rW_t + y_t - c)}_{\equiv a(t,W_t)}\Delta t + \sigma\Delta B_t, \end{aligned} \tag{S.1}$$

where $a(\cdot)$ is the *drift* of wealth and $\Delta B_t \equiv B_{t+\Delta t} - B_t$ is the increment of the Brownian motion (read: the i.i.d. shock to income). We assume for now that $W_t > 0$ is large enough so that the constraint $W \geq 0$ will not bind over $[t, t+\Delta t]$.

[1] We work with a time-dependent problem because it highlights the workings of the Markov-chain method. We choose a Brownian specification for income shocks in this section so as to explain how to deal with such shocks computationally; in Section S.2, we switch to a Poisson process for labor earnings (as in the paper), which is easier to deal with. The income process we choose is the continuous-time analogon to an independent and identically distributed (i.i.d.) income process $\tilde{y}_t$ with time-dependent mean $\mathbb{E}[\tilde{y}_t] = y_t$ and variance $\mathbb{E}[\tilde{y}_t - y_t]^2 = \sigma^2$. In continuous time, the household receives income $y_t\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$ in each instant. Over a unit of time, say a year, this adds up to $\int_0^1 y_t\,\mathrm{d}t + \int_0^1 \sigma\,\mathrm{d}B_t = \int_0^1 y_t\,\mathrm{d}t + \sigma(B_1 - B_0)$. Since Brownian motion has standardized variance, $\mathbb{E}_0[B_1 - B_0]^2 = 1$, and independent increments, yearly income has standard deviation σ and serially uncorrelated shocks.

[2] Note that in practice we have to ensure that the variance of the Brownian motion becomes small as wealth goes to zero so as to ensure that wealth remains nonnegative. This can be done by having the variance depend on wealth (e.g., σW).

Now, take a second-order Taylor approximation of $V(t+\Delta t, W_{t+\Delta t})$ in Δt around (t, W_t),

$$\begin{aligned}
&V(t+\Delta t, W_{t+\Delta t}) - V(t, W_t) \\
&\quad \approx V_t(t, W_t)\Delta t + V_W(t, W_t)(W_{t+\Delta t} - W_t) + V_{tt}\Delta t^2 \\
&\qquad + \frac{1}{2}V_{WW}(t, W_t)(W_{t+\Delta t} - W_t)^2 + V_{tW}(t, W_t)\Delta t(W_{t+\Delta t} - W_t),
\end{aligned} \tag{S.2}$$

where subscripts denote partial derivatives. Now, apply the rules of stochastic calculus to the cross-products:

$$\begin{aligned}
&(W_{t+\Delta t} - W_t)^2 \approx \sigma^2 \Delta t, \\
&\Delta t(W_{t+\Delta t} - W_t) \approx \sigma \Delta t \Delta B_t.
\end{aligned}$$

Substitute these into (S.2), take the expected value, divide by Δt, and drop terms of order lower than Δt to obtain

$$\begin{aligned}
&\mathbb{E}_t\left[\frac{V(t+\Delta t, W_{t+\Delta t}) - V(t, W_t)}{\Delta t}\right] \\
&\quad \approx V_t(t, W_t) + a(t, W_t)V_W(t, W_t) + \frac{\sigma^2}{2}V_{WW}(t, W_t).
\end{aligned} \tag{S.3}$$

Now, approximate $e^{-\rho\Delta t}$ by $1 - \rho\Delta t$ in (S.1) and subtract $V(t, W_t)$ on both sides to obtain

$$\rho\mathbb{E}_t\big[V(t+\Delta t, W_{t+\Delta t})\big]\Delta t = \max_{c\geq 0}\big\{u(c)\Delta t + \mathbb{E}_t\big[V(t+\Delta t, W_{t+\Delta t}) - V(t, W_t)\big]\big\}.$$

Finally, divide by Δt, use (S.3), and take the limit as $\Delta t \to 0$ to obtain the HJB

$$\begin{aligned}
-V_t(t, W_t) &= -\rho V(t, W_t) \\
&\quad + \max_{c\geq 0}\left\{u(c) + V_W(t, W_t)a(t, W_t) + \frac{\sigma^2}{2}V_{WW}(t, W_t)\right\},
\end{aligned} \tag{S.4}$$

where the drift $a(\cdot)$ is defined in (S.1).

S.1.2 *The Δt problem*

We describe the numerical solution method to (S.1) in three stages: first for $\sigma = 0$ (no risk), second for $\sigma > 0$ (transitory risk), and, finally, we introduce persistent earnings risk by letting y_t follow a discrete-state Markov process (Poisson process). Throughout, the grid for the state variable W is uniformly spaced with mesh size denoted by ΔW. The solution method consists of iterating backward on the Hamilton–Jacobi–Bellman equation. This is analogous to value function iteration in discrete time.[3]

[3] As in discrete time, it is also possible to iterate on the Euler equation in continuous time; see Barczyk and Kredler (2014), who use this method. For our altruism problem, however, this is not practical because consumption policies have discontinuities.

S.1.3 *Backward equation:* $\sigma = 0$

We start from Bellman's principle given in (S.1), supposing that we know the value function V as of time $t + \Delta t$ on the discretized state space. For all points off the grid, the values are linearly interpolated and denoted by $\tilde{V}$, using interpolation weights ω^+ and ω^-, as

$$\begin{aligned}\tilde{V}(t+\Delta t, W_{t+\Delta t}) = V(t+\Delta t, W) + \omega^+\big[V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)\big] \\ - \omega^-\big[V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)\big],\end{aligned} \tag{S.5}$$

where W is a grid point and $W_{t+\Delta t}$ results from the budget constraint. The interpolation weights contain the influence of the decision maker on the direction of the state's law of motion and are given by

$$\omega^+ = \max\left\{\frac{a\Delta t}{\Delta W}, 0\right\} \quad \text{and} \quad \omega^- = \max\left\{-\frac{a\Delta t}{\Delta W}, 0\right\}.$$

If the drift a is positive, the household is saving so that the state moves in an upward $(+)$ direction. Vice versa, if the household consumes more than her current income, wealth decreases and the state moves downward $(-)$.

Also, note that for given ΔW and a given drift a, we need to choose Δt small enough so that the weights ω^+ and ω^- do not exceed the bounds $[0, 1]$ for the interpolation in (S.5) to be valid. This yields the *Courant–Friedrichs–Lax* (CFL) condition, which is well known from the theory on finite-difference solution methods for partial differential equations (PDEs):

$$\frac{|a(t, W)|\Delta t}{\Delta W} \leq 1. \tag{S.6}$$

Whenever one encounters problems with the solution algorithm, one should first check if this condition is met for all points of the state space and for all t.

To link the discussion in this section to the stochastic case later on, we now point out that we can also interpret the interpolation Equation (S.5) in a stochastic sense. Suppose the household can only hold wealth levels that equal the points on the W grid, and that it jumps up (down) one grid point with probability ω^+ (ω^-) from t to $t+\Delta t$ or stays at the same grid point with probability $(1 - \omega^+ - \omega^-)$. Then this household has an expected continuation value as given by (S.5), and the three probabilities are well defined if and only if the CFL (S.6) holds.

We now substitute the linear interpolation (S.5) for the continuation value in (S.1). Furthermore, we approximate $e^{-\rho\Delta t}$ by $1 - \rho\Delta t$. Ignoring the max operator for now, the Δt problem is then given by

$$\begin{aligned}V(t, W) \approx\ & u(c)\Delta t + (1-\rho\Delta t)V(t+\Delta t, W) \\ & + (1-\rho\Delta t)\omega^+\big[V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)\big] \\ & - (1-\rho\Delta t)\omega^-\big[V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)\big].\end{aligned}$$

Using this equation, we could already update the value function backward in time. But we will now bring it into a form that is closer to the partial differential equation (PDE) given by the HJB (S.4).

We can approximate $V(\cdot, W+\Delta W) - V(\cdot, W) \approx V_W(\cdot, W)\Delta W$. Using this approximation, we see that terms like $\rho\Delta t[V(\cdot, W+\Delta W) - V(\cdot, W)]$ are of second order and can be neglected.[4] With this in mind, we can write

$$-\big[V(t+\Delta t, W) - V(t, W)\big] \approx -\rho\Delta t V(t+\Delta t, W) + u(c)\Delta t$$
$$+ \omega^+\big[V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)\big]$$
$$- \omega^-\big[V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)\big].$$

Finally, substitute the expressions for ω^+ and ω^-, and divide through by Δt,

$$\begin{aligned} &-\left[\frac{V(t+\Delta t, W) - V(t, W)}{\Delta t}\right] \\ &\approx -\rho V(t+\Delta t, W) + u(c) \\ &\quad + \left[\frac{V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)}{\Delta W}\right] a^+(t, W) \\ &\quad - \left[\frac{V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)}{\Delta W}\right] a^-(t, W), \end{aligned} \tag{S.7}$$

where we define the positive and negative parts of drift a as

$$a^+(t, W) = \max\big\{a(t, W), 0\big\}, \qquad a^-(t, W) = \max\big\{-a(t, W), 0\big\}.$$

In (S.7), we recognize a finite-difference approximation of the continuous-time HJB (S.4), leaving the max operator aside for now.

Equation (S.7) is of central importance since it tells us how the value function changes along the time dimension, specifically, *backward* in time. Note that the left-hand side is a numerical approximation to $-V_t(t, W)$ and the right-hand side is known. Using this information, we update the value function by moving a Δt period in the backward direction:

$$V(t, W) \approx V(t+\Delta t, W) - V_t(t, W)\Delta t.$$

If we are interested in computing a stationary solution, we continue to iterate backward until convergence, that is, until V_t is close to zero everywhere on the grid.

To fully understand Equation (S.7), note that the quotients on the right-hand side in the square brackets correspond to the numerical approximation of the partial derivative V_W. That is,

$$V_W \approx \frac{V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)}{\Delta W} \quad \text{and}$$
$$V_W \approx \frac{V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)}{\Delta W}$$

[4] In the algorithm, Δt linearly decreases as ΔW gets smaller, so we see that the term in question is of order ΔW^2. Terms of second order can, in general, be neglected.

are the forward and backward first-order difference quotients in W, respectively. This way of computing partial derivatives numerically is referred to as *upwind differencing* since differences are taken in the direction the system is moving. We can see that whether the forward or backward difference is used to approximate V_W depends on the drift of the economy. If $a \geq 0$, then the forward difference is used. When $a < 0$, we use the backward difference.[5]

S.1.4 *Backward equation:* $\sigma > 0$

We now proceed analogously for the case $\sigma > 0$. Once again, the value function V is assumed to be known as of time $t + \Delta t$ on the discretized state space.

When Brownian motion is present, in principle the state $W_{t+\Delta t}$ can take on values far away from the current level of wealth W_t. Fortunately, it has been shown that the true process can be well approximated with a three-state Markov chain that takes on values on the grid. The approximating process has the property that it either stays at the same grid point with probability π_m, jumps up to the next-higher grid point with probability π_u, or jumps down to the next-lower grid point with probability π_d. The probabilities are pinned down by requiring it to share the first and second (conditional) moments with the true process.[6]

The probabilities are then used to approximate the continuation value as

$$\begin{aligned}
&\tilde{\mathbb{E}}_t V(t+\Delta t, W_{t+\Delta t}) \\
&\quad = \pi_m V(t+\Delta t, W) + \pi_u V(t+\Delta t, W+\Delta W) + \pi_d V(t+\Delta t, W-\Delta W) \\
&\quad = V(t+\Delta t, W) + \pi_u \big[V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)\big] \\
&\qquad - \pi_d \big[V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)\big],
\end{aligned} \tag{S.8}$$

where we denote the expectation operator by $\tilde{\mathbb{E}}$ to indicate that this constitutes an approximation. To pin down the transition probabilities, we proceed in two steps: (i) set the drift to zero and (ii) add back in the drift. Finally, we discuss the method of centered differencing instead of upward differencing.

S.1.4.1 *Approximating Brownian motion* Suppose for now that the drift is zero. The law of motion consists only of the Brownian shocks $\sigma\, \mathrm{d}B_t$. To approximate this process,we choose probabilities π_u, π_d, and π_m so as to match the first and second (conditional) moments of this Brownian process:

$$\begin{aligned}
&\mathbb{E}(W_{t+\Delta t} - W_t) = \pi_u \Delta W - \pi_d \Delta W = 0, \\
&\mathbb{E}\big[(W_{t+\Delta t} - W_t)^2\big] = \pi_u (\Delta W)^2 + \pi_d (\Delta W)^2 = \sigma^2 \Delta t.
\end{aligned}$$

[5] Upwind differencing is often advocated as the preferred way to compute derivatives since information enters from the direction the system is moving to. Below we briefly outline *centered differencing* as another way to compute derivatives.

[6] This method is in spirit comparable to Tauchen (1986) and is studied by, for example, Kushner and Dupuis (2001). They show that for the approximating process to converge to the true process as $\Delta t \to 0$, it is enough to ensure that the Markov transition probabilities given by π_m, π_u, and π_d, are *locally consistent* (i.e., they have the same first and second conditional moments as the underlying process).

Solving these two equations for the two unknowns yields

$$\pi_u = \pi_d = \frac{\sigma^2 \Delta t}{2(\Delta W)^2} \equiv p.$$

For the probabilities to be positive, again a stability condition (the CFL) has to hold. That is, to ensure that $\pi_m \geq 0$, we need to have that

$$\frac{\sigma^2 \Delta t}{(\Delta W)^2} \leq 1.$$

Note that the *binomial* approximation method arises as a special case when Δt is set such that $\pi_m = 0$. In this case, $p = \frac{1}{2}$.

S.1.4.2 *Adding back the drift* Now suppose that $\sigma = 0$. Looking back at (S.5) and interpreting the interpolation weights as transition probabilities shows that we have already done the work. Suppose, for example, that the drift equals zero. Then $\pi_m = 1$ and the household remains at the same level of wealth a Δt period later. When the drift is positive, we take away probability mass from the middle and shift it upward. Vice versa, probability mass is shifted downward when the drift is negative.

We now add the drift to the Brownian-motion process, which itself has no drift. To do this, we take the jump probabilities found before and then shift probability mass from the middle up- or downward. The transition probabilities then become

$$\begin{aligned}
\pi(W, W + \Delta W) &\equiv \pi_u = p + \omega^+ = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma^2}{2} + \Delta W a^+(t, W)\right), \\
\pi(W, W - \Delta W) &\equiv \pi_d = p + \omega^- = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma^2}{2} + \Delta W a^-(t, W)\right), \\
\pi(W, W) &\equiv \pi_m = 1 - \pi_u - \pi_d.
\end{aligned} \tag{S.9}$$

Of course, we have to ensure again that all probabilities are positive. Again, Δt has to be chosen small enough such that $\pi_m \geq 0$.

We will now show again that using these probabilities leads to a finite-difference interpretation of the HJB. Substituting p and the ω's into (S.8), we obtain

$$\begin{aligned}
&\tilde{\mathbb{E}}_t V(t + \Delta t, W_{t+\Delta t}) \\
&\quad = V(\cdot, W) + p\big[V(\cdot, W + \Delta W) + V(\cdot, W - \Delta W) - 2V(\cdot, W)\big] \\
&\qquad + \omega^+\big[V(\cdot, W + \Delta W) - V(\cdot, W)\big] - \omega^-\big[V(\cdot, W) - V(\cdot, W - \Delta W)\big],
\end{aligned}$$

where we have replaced the argument $t + \Delta t$ by a dot for the sake of better readability. The expression

$$\begin{aligned}
&p\big[V(\cdot, W + \Delta W) + V(\cdot, W - \Delta W) - 2V(\cdot, W)\big] \\
&\quad = \frac{\sigma^2 \Delta t}{2}\left[\frac{V(\cdot, W + \Delta W) + V(\cdot, W - \Delta W) - 2V(\cdot, W)}{(\Delta W)^2}\right]
\end{aligned}$$

involves the second-order difference quotient in W (the fraction in brackets), which approximates the second derivative of V with respect to W. The terms related to the ω's are the same forward and backward first-order difference quotients in W as for the deterministic case in (S.7).

We now get the analogous equation to (S.7) by following the same steps that led up to (S.7), except that a term related to uncertainty is now included:

$$\begin{aligned}
&-\left[\frac{V(t+\Delta t, W) - V(t, W)}{\Delta t}\right] \\
&\quad \approx -\rho V(t+\Delta t, W) + u(c) \\
&\qquad + \left[\frac{V(t+\Delta t, W+\Delta W) - V(t+\Delta t, W)}{\Delta W}\right] a^{+}(t, W) \\
&\qquad - \left[\frac{V(t+\Delta t, W) - V(t+\Delta t, W-\Delta W)}{\Delta W}\right] a^{-}(t, W) \\
&\qquad + \frac{\sigma^2}{2}\left[\frac{V(t+\Delta t, W+\Delta W) + V(t+\Delta t, W-\Delta W) - 2V(t+\Delta t, W)}{(\Delta W)^2}\right].
\end{aligned} \tag{S.10}$$

We again recognize a finite-difference version of the continuous-time HJB (S.4), again leaving the max operator for c aside for later.

The significance of (S.10), just like that of (S.7), is that it provides information on how the value function changes going backward in time. The backward-iteration steps for the value function are again given by

$$V(t, W) \approx V(t+\Delta t, W_t) - V_t(t, W_t)\Delta t,$$

where the approximation for $-V_t(t, W_t)$ is given by the left-hand side of (S.10). To compute a stationary solution, we iterate until V_t is close to zero on the entire grid.

S.1.4.3 *Centered differencing* We now briefly describe a different method of dealing with the drift term. Note that we can directly approximate $\mathrm{d}W_t = a(t, W)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$ on the trinomial lattice, pinning down the probabilities by requiring a three-state Markov chain to have the correct mean and variance, an approach similar to Tauchen (1986):

$$\begin{aligned}
&\mathbb{E}(W_{t+\Delta t} - W_t) = \pi_u \Delta W - \pi_d \Delta W = a(t, W)\Delta t, \\
&\mathbb{E}\big[(W_{t+\Delta t} - W_t)^2\big] = \pi_u(\Delta W)^2 + \pi_d(\Delta W)^2 = \sigma^2 \Delta t.
\end{aligned}$$

Here, we have neglected to take into account that $\mathbb{E}_t(W_{t+\Delta t}) \neq W_t$ if $a(t, W_t) \neq 0$ in the calculation of the variance. This is not important: it can easily be show that the terms arising from this are of second order and may thus be neglected.

From the first equation, we get $\pi_d = \pi_u - (a(t, W)\Delta t)/\Delta W$, and from the second equation, we obtain $\pi_d = (\sigma^2 \Delta t)/(\Delta W)^2 - \pi_u$. From this it follows that

$$\begin{aligned}
\pi_u &= \frac{\Delta t}{2(\Delta W)^2}\big[\sigma^2 + a(t, W)\Delta W\big], \\
\pi_d &= \frac{\Delta t}{2(\Delta W)^2}\big[\sigma^2 - a(t, W)\Delta W\big].
\end{aligned}$$

So this approach shifts probability away from the bottom (and not from the middle) and shifts it to the top if the drift a is positive. For a negative drift, probability mass is shifted from the top to the bottom. Again, we have $\pi_m = 1 - \pi_u - \pi_d$ and we have to choose Δt small enough to respect the obvious stability condition arising from the inequality $\pi_m \geq 0$.

Following the same steps as above, we find a version of (S.10) (not shown here) in which forward and backward first-order difference quotients in W are replaced by the first-order *centered* difference quotient $[V(t, W + \Delta W) - V(t, W - \Delta W)]/2 \approx V_W(t, W)$.[7]

S.1.5 *Optimal consumption*

We now turn toward the resolution of the max operator in the HJB and answer the question, "How is optimal consumption computed (using the method of upwind differencing)?"

First, we group the terms of (S.10) in which the control variable c enters into the following discretized version of the Hamiltonian:

$$\begin{aligned}\mathcal{H}(c) \equiv u(c) + a^+(t, W)\left[\frac{V(\cdot, W + \Delta W) - V(\cdot, W)}{\Delta W}\right] \\ - a^-(t, W)\left[\frac{V(\cdot, W) - V(\cdot, W - \Delta W)}{\Delta W}\right].\end{aligned}$$

The household maximizes $\mathcal{H}$ by choosing $c \geq 0$. Note that when varying c, flow utility changes as well as the drift of wealth (i.e., the savings rate), giving rise to the familiar consumption–savings trade-off.

If the household saves, then $a^- = 0$ and the household's marginal value of saving is evaluated using the forward difference in W. If it dis-saves, then $a^+ = 0$ and the household's marginal value of saving is evaluated using the backward difference. Finally, if $a^- = 0 = a^+$, consumption equals income and the wealth position of the household remains unchanged. To find optimal consumption, let c vary on $(0, \infty)$ and choose the one that leads to the largest value of $\mathcal{H}$.

To find a simple formula, we define the smooth, strictly concave functions in c:

$$\begin{aligned}\mathcal{H}^+(c) &\equiv u(c) + (rW + y - c)\left[\frac{V(\cdot, W + \Delta W) - V(\cdot, W)}{\Delta W}\right], \\ \mathcal{H}^-(c) &\equiv u(c) + (rW + y - c)\left[\frac{V(\cdot, W) - V(\cdot, W - \Delta W)}{\Delta W}\right].\end{aligned}$$

Compute the unconstrained maximizers c^{*+} for $\mathcal{H}^+$ and c^{*-} for $\mathcal{H}^-$, respectively, using the first-order condition. If $c^{*+} \leq rW + y$, then $c^* \equiv c^{*+}$ is a forward maximizer. If $c^{*-} > rW + y$, then $c^{*-} \equiv c^{*-}$ is a backward maximizer. If there exists only a forward

[7]If the true value function is differentiable, then both are of course the same in the limit as $\Delta W \to 0$. If there are kinks in the true value function, then the upwind method is preferable since it takes into account the relevant directional derivatives, whereas the finite-difference method takes a (usually meaningless) average between the two directional derivatives.

maximizer (or only a backward maximizer), then this maximizer constitutes the optimal consumption rate. If none of the two exists, then $c^* = rW + y$, that is, the consumer should stick with the current wealth level. If both maximizers exist, choose the one that yields a higher $\mathcal{H}$. This last case does not arise in the case that $V(t + \Delta t, \cdot)$ is concave in wealth.

S.1.6 *Constraints and grid boundaries*

The method described so far works for points that are inside the W grid. We now turn to procedures for what to do on the margins of the grid.

At the uppermost grid point for W (call it $\bar{W}$), we face the difficulty that we do not know the value function at $\bar{W} + \Delta W$. There are two ways to deal with this. First, one can extrapolate the value function (say, by requiring the third derivative to be constant, as we do in our code). Second, one can reflect agents back into the grid in case they receive a positive shock and jump out. This amounts to adding our previous value for π_u to π_m and setting $\pi_u = 0$ at the uppermost grid point. In our computations, we find that extrapolation leads to better approximations for policies and value functions close to the grid boundary; this is because reflection rules out upside risk at $\bar{W}$ and thus makes this grid point radically different from interior points. However, reflection turns out to be the more stable method, especially once altruism is added into the setting. This is presumably the case because reflection induces a meaningful economic game on a discrete, bounded grid, whereas extrapolation does not. When using the reflection method, we have to make sure that agents are very unlikely to reach the regions in the state space where the influence of the boundary is apparent. This can always be ensured by choosing a large-enough grid for W.

At the bottom grid point $W = 0$, we have to check whether agents are constrained. To do this, we compute the optimal consumption policy given the forward derivative $[V(t + \Delta t, \Delta W) - V(t + \Delta t, 0)]/\Delta W$. If this consumption policy is feasible, then it constitutes the optimal consumption policy: the agent is unconstrained and saves. If it is not feasible, then the optimal consumption policy is to consume the entire income and stay at $W = 0$. As mentioned before, it is important in practice to impose an income-shock process that vanishes as W goes to zero (such as $\sigma W \, \mathrm{d}B_t$) so as to avoid complications at the boundary.

S.1.7 *Persistent income shocks: Poisson process for y*

We now allow y_t to follow a discrete-state Markov chain. For simplicity, consider a two-state Markov chain with $y \in \{y^b, y^g\}$, where $y^b < y^g$, and Poisson rate η. This means that the probability of switching from one income state to the other over a short time interval Δt is approximately $\eta \Delta t$. The value function is now also a function of y_t. We maintain the value function's dependence on t since we use backward iteration again.

Suppose again that we know $V(t+\Delta t, W_{t+\Delta t}, y)$. The Δt problem for a household that currently has an income endowment y^g and wealth $W_t > 0$ is given by

$$\begin{aligned} V(t, W_t, y^g) &= \max_{c\geq 0}\big\{u(c)\Delta t + e^{-\rho\Delta t}\big[\eta\Delta t\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^b) \\ &\quad + (1-\eta\Delta t)\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^g)\big]\big\} \\ \text{s.t.} \quad W_{t+\Delta t} - W_t &= \underbrace{(rW_t + y^g - c)}_{\equiv a(t,W_t,y^g)}\Delta t + \sigma\Delta B_t. \end{aligned}$$

We now derive the HJB, which will show us that the numerical approximation for this case is practically the same as in the cases before.

Using a Taylor approximation for $e^{-\rho\Delta t} \approx 1-\rho\Delta t$ and canceling terms in $(\Delta t)^2$, we find that

$$\begin{aligned} V(t, W_t, y^g) &= u(c_t)\Delta t + \eta\Delta t\big[\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^b) - \mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^g)\big] \\ &\quad + (1-\rho\Delta t)\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^g), \end{aligned}$$

where as before we neglect the max operator for now. We then expand $V(t+\Delta t, W_{t+\Delta t}, y) = V(t+\Delta t, W_t, y) + a(t, W_t, y)\Delta t V_W(t+\Delta t, W_t, y)$ and drop terms in $(\Delta t)^2$ to find

$$\begin{aligned} V(t, W_t, y^g) &= u(c_t)\Delta t + \eta\Delta t\big[V(t+\Delta t, W_t, y^b) - V(t+\Delta t, W_t, y^g)\big] \\ &\quad + (1-\rho\Delta t)\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^g). \end{aligned}$$

Here, we recognize Bellman's principle from the case with transitory income risk: when crossing out the term in $\eta\Delta t$, we are back to Bellman's principle for constant y. Thus, the above expression gives us the recipe to adjust our algorithm from before (for $\sigma > 0$) to a persistent income process: we just have to introduce a new jump probability of $\eta\Delta t$ to the point $(t+\Delta t, W_t, y_b)$, which takes into account the risk of a change in y. We then have to subtract $\eta\Delta t$ from π_m—recall that π_m is the probability of ending up at $(t+\Delta t, W_t, y_g)$. The probabilities π_u and π_d for moving up or down in the W direction are not affected. Notice that the above derivation shows that we do not have to be concerned about interactions between y and W as Δt gets small, that is, we do not have to include jumps to the point $(t+\Delta t, W+\Delta W, y_b)$ and the like, since they are of second order.

We now proceed to derive the continuous-time HJB. Rearranging the above, dividing by Δt, and then letting $\Delta t \to 0$, we find

$$\begin{aligned} \rho V(t, W_t, y^g) &= u(c_t) + \eta\big[V(t, W_t, y^b) - V(t, W_t, y^g)\big] \\ &\quad + \lim_{\Delta t\to 0}\left[\frac{\mathbb{E}_t V(t+\Delta t, W_{t+\Delta t}, y^g) - V(t, W_t, y^g)}{\Delta t}\right]. \end{aligned}$$

We see that the last term is the same as in the previous derivation of the HJB shown by (S.4). Thus, bringing back in the max operator, the HJB for the persistent income process is given by

$$
\begin{aligned}
&-V_t\big(t, W_t, y^g\big) \\
&\quad= -\rho V\big(t, W_t, y^g\big) + \eta\big[V\big(t, W_t, y^b\big) - V\big(t, W_t, y^g\big)\big] \\
&\qquad + \max_{c\geq 0}\left\{u(c) + a\big(t, W_t, y^g\big)V_W\big(t, W_t, y^g\big) + \frac{\sigma^2}{2}V_{WW}\big(t, W_t, y^g\big)\right\}.
\end{aligned}
\tag{S.11}
$$

Thus, computing the change of the value function along the time dimension, V_t, follows (S.10) but includes the additional term due to the possibility of transiting from income state y^g to y^b, which is taken into account through the difference in the respective value functions, multiplied by the Poisson rate η.

To solve the problem, we iterate backward on the two value functions, which again shows us that the Markov-chain method has an interpretation as a finite-difference method:

$$
\begin{aligned}
V\big(t-\Delta t, W_{t-\Delta t}, y^g\big) &\approx V\big(t, W_t, y^g\big) - V_t\big(t, W_t, y^g\big)\Delta t, \\
V\big(t-\Delta t, W_{t-\Delta t}, y^b\big) &\approx V\big(t, W_t, y^b\big) - V_t\big(t, W_t, y^b\big)\Delta t.
\end{aligned}
$$

S.1.8 *Computing densities: The forward equation*

In this section, we briefly explain how we compute the density of agents over the state space, here for the case where y is constant. We then illustrate how the method relates to the *Kolmogorov forward equation* that governs the evolution of densities in a continuous-time setting with Brownian shocks.

S.1.8.1 *Computing densities* Suppose we know the distribution of a mass of households at time 0, given by a density function $N(0, W)$. Furthermore, the optimal policies computed above provide a law of motion for the households:

$$
\mathrm{d}W_t = a(t, W_t)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t.
$$

The goal is to find the density function $N(t, W)$ of households over the state space at each point in time $t > 0$ that results from the initial density and the law of motion.

To do this, we will just map forward the density on our grid using the transition probabilities we derived before. So the method amounts to computing the density evolution for a discrete-state Markov chain in discrete time, which may be represented by

$$
N'_{t+\Delta t} = N'_t P_t, \tag{S.12}
$$

where P_t is an $k \times k$ transition matrix and N_t is a $k \times 1$ vector (approximating the density) whose ith element is the probability of being in state i at time t. The matrix P_t contains three nonzero elements per column, which are given by the transition probabilities (π_u, π_m, π_d) for the respective (W, t) position and are in the three positions around the diagonal.[8]

[8]There is an interesting connection to value-function iteration here. On the discretized state space, value-function iteration may be expressed as $V_t = u(c_t)\Delta t + P_t V_{t+\Delta t}$ in matrix form. So we see that the for-

If we have time-invariant policy functions, then P_t is time-invariant. We can then find the stationary distribution by iterating on the above until $N_{t+\Delta t} \simeq N_t$ (according to some convergence criterion).

S.1.8.2 *The Kolmogorov forward equation* We now state the *Kolmogorov forward equation* for our problem, which is the continuous-time analogon to the matrix operation in (S.12). It is a partial differential equation that the density function $N(t, W)$ has to obey:

$$N_t(t, W) = -a_W(t, W)N(t, W) - a(t, W)N_W(t, W) + \frac{\sigma^2}{2}N_{WW}(t, W). \tag{S.13}$$

There are further terms in the Kolmogorov forward equation if volatility σ is a function of W, but these are zero in our example. We first develop some intuition for this equation before showing the connection to the computational algorithm given by (S.12).

To do this, we first set $\sigma = 0$. In this case, the forward equation can be written as

$$N_t(t, W) + a(t, W)N_W(t, W) = -a_W(t, W)N(t, W). \tag{S.14}$$

To understand what this equation says, consider the so-called characteristic curve that a single household follows in (t, W) space. Parameterize the household's level of wealth by t:

$$W(t) = W(0) + \int_0^t a\big(s, W(s)\big)\, \mathrm{d}s.$$

Now take the total derivative of the density function N with respect to time t, following a household's path

$$\frac{\mathrm{d}N(t, W(t))}{\mathrm{d}t} = N_t\big(t, W(t)\big) + a\big(t, W(t)\big)N_W\big(t, W(t)\big),$$

which is the left-hand side of the forward equation (S.14) in the case without Brownian motion. So on the left-hand side of (S.14), we have the change in the density when following an agent on his optimal path. Equation (S.14) says that the growth in the density along this "characteristic curve" depends on what the drift looks like in a neighborhood, which we see on the right-hand side. For example, if $a_W = 0$, then the density does not change along the characteristic curve. But if $a_W > 0$, the drift is increasing in wealth and agents' paths diverges along the characteristic curve, that is, the density thins out over time.

Next, consider the case in which there is only Brownian motion but no drift. In this case, the Kolmogorov forward equation is given by

$$N_t(t, W) = \frac{\sigma^2}{2}N_{WW}(t, W),$$

which also goes by the name *heat equation*, since it provides us with a description of how heat spreads in a material. If the density function is convex, it will tend to increase over time. In our model, if there are more agents both above and below the current W, then

ward operator is the transpose of the backward operator. This translates to the forward partial-differentiable operator being the *adjoint* operator (a generalized transpose) of the backward partial-differential operator.

purely random movements will tend to even out the density at this point. On the other hand, if the density is linear, then $N_{WW} = 0$ so that $N_t = 0$ and the density is stationary.

S.1.8.3 *Connection between the computational algorithm and the forward equation*
We will now see how the continuous-time forward equation (S.13) is linked to the approximating algorithm given by (S.12). Over a short time interval Δt, a household can jump to the grid point $(t + \Delta t, W)$ coming from three possible points on the grid $((t, W), (t, W + \Delta W), (t, W - \Delta W))$ with Markov transition probabilities given by π_m, π_d, and π_u, respectively:

$$N(t+\Delta t, W) = \pi_d(t, W+\Delta W)N(t, W+\Delta W) + \pi_m(t, W)N(t, W) + \pi_u(t, W-\Delta W)N(t, W-\Delta W). \tag{S.15}$$

Note that the probabilities are written as a function of the grid point the economy is coming from.

To see that (S.15) approximates the Kolmogorov forward equation (S.13), begin by substituting

$$\pi_m(t, W) = 1 - \pi_u(t, W) - \pi_d(t, W)$$

in (S.15), then add and subtract the terms $\pi_d(t, W+\Delta W)N(t, W)$ and $\pi_u(t, W-\Delta W) \times N(t, W)$ on the right-hand side, and move $N(t, W)$ to the left-hand side to obtain

$$\begin{aligned} N(t+\Delta t, W) - N(t, W) = {} & \pi_d(t, W+\Delta W)\big[N(t, W+\Delta W) - N(t, W)\big] \\ & + \pi_u(t, W-\Delta W)\big[N(t, W-\Delta W) - N(t, W)\big] \\ & + N(t, W)\big[\pi_d(t, W+\Delta W) - \pi_d(t, W)\big] \\ & + N(t, W)\big[\pi_u(t, W-\Delta W) - \pi_u(t, W)\big]. \end{aligned}$$

We can already see that changes in the probabilities π_u and π_d along the W dimension play an important role. Finally, to make the connection to the drift a and volatility σ, substitute in for π_u and π_d from (S.9), and divide through by Δt to obtain

$$\begin{aligned} & \left[\frac{N(t+\Delta t, W) - N(t, W)}{\Delta t}\right] \\ & = -\left[\frac{N(t, W) - N(t, W-\Delta W)}{\Delta W}\right]a^+(t, W-\Delta W) \\ & \quad + \left[\frac{N(t, W+\Delta W) - N(t, W)}{\Delta W}\right]a^-(t, W+\Delta W) \\ & \quad - \left[\frac{a^+(t, W) - a^+(t, W-\Delta W)}{\Delta W}\right]N(t, W) \\ & \quad + \left[\frac{a^-(t, W+\Delta W) - a^-(t, W)}{\Delta W}\right]N(t, W) \\ & \quad + \frac{\sigma^2}{2}\left[\frac{N(t, W+\Delta W) + N(t, W-\Delta W) - 2N(t, W)}{(\Delta W)^2}\right], \end{aligned}$$

which is a numerical approximation to the Kolmogorov forward equation (S.13).[9]

Finally, recognizing that the right-hand side gives us a discretized expression for $N_t(t,W)$, we see that this gives us another, equivalent, way to compute the density a Δt period in the future:

$$N(t+\Delta t,W)\approx N(t,W)+N_t(t,W)\Delta t.$$

S.2. Altruism model

In our Markov-perfect environment, computing a solution to a consumption-smoothing problem when agents are imperfectly altruistic is very similar to the previous examples, which we consider an important strength of our approach. The following discussion shows how to adapt the numerical algorithm to the altruistic framework. The code can be found in `AltruismUncert_main.m` in the folder `Altruism`.

S.2.1 *The* Δt *problem*

To keep things simple, let us assume for now that there are no shocks to (y,y'). The payoff-relevant state is then (t,W,W'). The Δt problem, given a Markov strategy $\{c',g'\}$ of the other player, is

$$\begin{aligned}
V\big(t,W_t,W'_t\big) &= \max_{c,g}\big\{u(c)\Delta t+\alpha u\big(c'\big)\Delta t+e^{-\rho\Delta t}\mathbb{E}_t V\big(t+\Delta t,W_{t+\Delta t},W'_{t+\Delta t}\big)\big\}\\
\text{s.t.}\quad W_{t+\Delta t}-W_t &= \underbrace{\big(rW_t+y_t+g'_t-c_t-g_t\big)}_{\equiv a(t,W_t,W'_t)}\Delta t+\sigma\Delta B_t,\\
W'_{t+\Delta t}-W'_t &= \underbrace{\big(rW'_t+y'_t+g_t-c'_t-g'_t\big)}_{\equiv a'(t,W_t,W'_t)}\Delta t+\sigma'\Delta B'_t,\\
c\geq 0,\qquad g\geq 0,&\qquad W\geq 0,\qquad W_t,W'_t\text{ given.}
\end{aligned}\tag{S.16}$$

Suppose again that we have $V(t+\Delta t,W_{t+\Delta t},W'_{t+\Delta t})$ and $V'(t+\Delta t,W_{t+\Delta t},W'_{t+\Delta t})$ given, that is, the values are known on the grid. It turns out that it is enough to put together jumps in the W direction for her and in the W' direction for him.[10] Specifically, a household that is currently in state (t,W,W') can jump over a small time interval Δt to five possible points on the grid at $t+\Delta t$ (time argument omitted),

$$\begin{aligned}
&\big(\cdot,W,W'\big),\qquad \big(\cdot,W+\Delta W,W'\big),\qquad \big(\cdot,W-\Delta W,W'\big),\\
&\big(\cdot,W,W'+\Delta W\big),\qquad \big(\cdot,W,W'-\Delta W\big),
\end{aligned}$$

[9] To obtain an expression on the right-hand side that is even closer to (S.13), in the second line, we can approximate $a^+(t,W-\Delta W)\approx a^+(t,W)-a^+_W(t,W)\Delta W$ and $a^-(t,W+\Delta W)\approx a^-(t,W)+a^-_W(t,W)\Delta W$, and then drop the lowest-order terms so that $a^+(\cdot)$ and $a^-(\cdot)$ are evaluated at (t,W).

[10] This is because the two Brownian motions are uncorrelated. If this was not the case, then also jumps to points such as $(W+\Delta W,W'+\Delta W)$ would have to be considered.

and these occur with Markov transition probabilities π_m, π_u, π_d, π'_u, and π'_d, respectively. Analogously to (S.8), these probabilities are used to approximate values off the grid,

$$\begin{aligned}
&\tilde{\mathbb{E}}_t V\big(t+\Delta t, W_{t+\Delta t}, W'_{t+\Delta t}\big) \\
&\quad = \pi_m V\big(\cdot, W, W'\big) + \pi_u V\big(\cdot, W+\Delta W, W'\big) \\
&\qquad + \pi_d V\big(\cdot, W-\Delta W, W'\big) + \pi'_u V\big(\cdot, W, W'+\Delta W\big) \\
&\qquad + \pi'_d V\big(\cdot, W, W'-\Delta W\big) \\
&\quad = V\big(\cdot, W, W'\big) + \pi_u\big[V\big(\cdot, W+\Delta W, W'\big) - V\big(\cdot, W, W'\big)\big] \\
&\qquad - \pi_d\big[V\big(\cdot, W, W'\big) - V\big(\cdot, W-\Delta W, W'\big)\big] \\
&\qquad + \pi'_u\big[V\big(\cdot, W, W'+\Delta W\big) - V\big(\cdot, W, W'\big)\big] \\
&\qquad - \pi'_d\big[V\big(\cdot, W, W'\big) - V\big(\cdot, W, W'-\Delta W\big)\big],
\end{aligned} \tag{S.17}$$

where the time argument $t+\Delta t$ is omitted in the value functions on the right-hand side for better readability.

Analogously to (S.9), the transition probabilities are given by

$$\begin{aligned}
&\pi(W, W+\Delta W) \equiv \pi_u = p + \omega^+ = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma^2}{2} + a^+\big(t, W, W'\big)\Delta W\right), \\
&\pi(W, W-\Delta W) \equiv \pi_d = p + \omega^- = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma^2}{2} + a^-\big(t, W, W'\big)\Delta W\right), \\
&\pi\big(W', W'+\Delta W\big) \equiv \pi'_u = p' + \omega'^+ = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma'^2}{2} + a'^+\big(t, W, W'\big)\Delta W\right), \\
&\pi\big(W', W'-\Delta W\big) \equiv \pi'_d = p' + \omega'^- = \frac{\Delta t}{(\Delta W)^2}\left(\frac{\sigma'^2}{2} + a'^-\big(t, W, W'\big)\Delta W\right), \\
&\pi_m = 1 - \pi_u - \pi_d - \pi'_u - \pi'_d,
\end{aligned}$$

where as before $p = \frac{1}{2}\sigma^2\Delta t/(\Delta W)^2$ and analogously $p' = \frac{1}{2}\sigma'^2\Delta t/(\Delta W)^2$. Note that again we have $\pi_m = 1 - p - p'$ in the case when drifts are zero for both agents.

Substitute the p's and the ω's into (S.17). As in (S.10), terms related to p and p' arise from the stochastic parts of the laws of motion and correspond to second-order difference quotients in W and W', that is,

$$\begin{aligned}
&p\big[V\big(\cdot, W+\Delta W, W'\big) + V\big(\cdot, W-\Delta W, W'\big) - 2V\big(\cdot, W, W'\big)\big] \\
&\quad + p'\big[V\big(\cdot, W, W'+\Delta W\big) + V\big(\cdot, W, W'-\Delta W\big) - 2V\big(\cdot, W, W'\big)\big],
\end{aligned}$$

but are not subject to control by the decision maker. Terms related to the ω's,

$$\begin{aligned}
&\omega^+\big[V\big(\cdot, W+\Delta W, W'\big) - V\big(\cdot, W, W'\big)\big] - \omega^-\big[V\big(\cdot, W, W'\big) - V\big(\cdot, W-\Delta W, W'\big)\big] \\
&\quad + \omega'^+\big[V\big(\cdot, W, W'+\Delta W\big) - V\big(\cdot, W, W'\big)\big] \\
&\quad - \omega'^-\big[V\big(\cdot, W, W'\big) - V\big(\cdot, W, W'-\Delta W\big)\big],
\end{aligned}$$

are first-order forward and backward difference quotients in W and W' that arise from the drifts $a(t, W, W')$ and $a'(t, W, W')$, and *are* subject to the control variables (c, g) by the decision maker.

Now proceed as before. Substitute the p's and ω's into the Δt problem in (S.16), ignore the max operator for now, and approximate $e^{-\rho\Delta t} \approx 1 - \rho\Delta t$ to obtain

$$\begin{aligned} &-\left[\frac{V(t+\Delta t, W, W') - V(t, W, W')}{\Delta t}\right] \\ &\approx -\rho V + u(c) + \alpha u(c') \\ &\quad + a^{+}[V^{+}] - a^{-}[V^{-}] + a'^{+}[V^{+\prime}] - a'^{-}[V^{-\prime}] \\ &\quad + \frac{\sigma^2}{2}[V^{+-}] + \frac{\sigma'^2}{2}[V^{+-\prime}], \end{aligned} \tag{S.18}$$

where the arguments of the functions on the right-hand side have been suppressed. The following abbreviations are used: $[V^{+}]$ stands for the first-order forward difference quotient in W; $[V^{+\prime}]$ stands for the first-order forward difference quotient in W', analogously $[V^{-}]$ and $[V^{-\prime}]$; $[V^{+-}]$ and $[V^{+-\prime}]$ stand for the second-order centered difference quotients in W and W', respectively. We again see that this is a discrete approximation of the HJB for the game.

Shocks to y can be added in a way entirely analogous to the single-agent case described in Section S.1.7.

S.2.2 *Optimal consumption and transfers*

Group the terms in which the control variables c and g enter into the function

$$\mathcal{H}(c, g) \equiv u(c) + a^{+}[V^{+}] - a^{-}[V^{-}] + a'^{+}[V^{+\prime}] - a'^{-}[V^{-\prime}].$$

As pointed out in the paper, the other agent's contemporaneous decisions do not influence an agent's choice over Δt. So the optimal value for consumption c^* can be determined exactly as in the single-agent case (see Section S.1.5 of this computational appendix) *only* by looking at the agent's marginal value of saving, which is encoded in V^{+} and V^{-}. This makes the algorithm simple and fast.

As for transfers, these are set to zero whenever the recipient has nonzero wealth (here $W' > 0$) according to our guess for equilibrium. When the other player is broke, we compute optimal transfers according to the procedure laid out in Section 3.3 and Appendix A.2 in the paper.

S.2.3 *Computational issues*

We find that reflecting agents back into the grid at the top levels of wealth makes the algorithm more stable than extrapolating the value function (this choice can be made by choosing the variable `Extrap` in our code). When reflecting on the top boundary, we have to make sure that agents are very unlikely (under the ergodic distribution) to reach the regions in the state space where the influence of the boundary on consumption policies is apparent. This can always be ensured by choosing a large-enough grid for W.

We experimented with nonlinear grids, which require slight modifications in the trinomial-grid method. However, it turned out that the algorithm was less stable than with a linear grid. We also tried the Howard improvement algorithm, but again stability was an issue. The same was true for leaving out policy updating and using the same policy rules for more than one Δt: There were slight gains in speed, but stability was lost.

Finally, we found that centered differencing worked better than upwind differencing. The results of the two algorithms were essentially the same, but centered differentiating led to faster convergence.

S.3. Extending the algorithm

We will now demonstrate how to adapt our framework to overlapping-generations and finite-horizon settings, and how to introduce endogenous risk-taking as well as other choice variables.

S.3.1 *Overlapping generations*

Our baseline model can be used as a building block for an overlapping-generations economy. Here, the two altruistically linked agents are a parent household and a child household; together they form a *family*. We will designate *her* to be the parent household (plain variables) and *him* to be the child household (variables with primes).

An especially simple way to handle the demography is as follows. The parent household faces a mortality hazard given by a Poisson rate δ. The child household becomes a parent household upon the death of the current parent household and a new child household is born (which we will refer to as the grandchild), resulting in a new family.[11]

When the parent household dies, the remaining wealth is bequeathed to the child household. The grandchild's first income realization is a random variable that is allowed to depend on the old child household's current income. This is a simple way to allow for labor-market skill to be inheritable, that is, a household is more likely to enter the economy with a high income realization when its parent household is a high-income earner. We denote by π_{ij} the probability that the grandchild household obtains the initial income realization y'_j, given that the child household has the current income realization y'_i.

The parent household's HJB is given by

$$\begin{aligned}\rho v = {} & \xi\big[v(\cdot,\tilde{y}) - v(\cdot,y)\big] + \xi'\big[v(\cdot,\tilde{y}') - v(\cdot,y')\big] + \delta\big(\alpha v^e - v\big) \\ & + \alpha u\big(c'\big) + \big(rw' + y' - g' - c'\big)v_{w'} + \frac{\sigma^2}{2}\big(w^2 v_{ww} + w'^2 v_{w'w'}\big) \\ & + \max_{c\geq 0}\big\{u(c) + \big(rw + y + g' - c\big)v_w\big\} + \max_{g\geq 0}\big\{g[v_{w'} - v_w]\big\},\end{aligned}$$

[11] Combining the OLG framework here with the time- or age-dependent finite-horizon case in the next subsection, one can model demography more realistically, as is done in Barczyk and Kredler (2013).

where v^e is the expected value for the child household when the parent dies and they become a parent themselves, which will be discussed in more detail later. This HJB is the same as the one presented in the main paper (see Equation (8)) except that the additional term $\delta(\alpha v^e - v)$ enters. In the child household's HJB, this term shows up as $\delta(v^e - v')$. This is because the child household becomes a parent household upon the death of its parent household. The reason this term also shows up in the parent household's HJB is due to altruism. When $\alpha = 0$, this term equals $-\delta v$ so that the child household's welfare is disregarded; when $\alpha = 1$, the parent household fully internalizes the continuation value of the child household.

As mentioned before, v^e is the expected value of the child household to become a parent household. The child household has to form an expectation about the initial income realization of the grandchild household, and it has to take into account bequest from the dying parents. Given that the child has income realization y_i, this expected value is given by

$$v^e(W, y_i) = \sum_j \pi_{ij} v\big(W, w_b\big(y'_j\big), y_i, y'_j\big),$$

where $W = w + w'$, that is, the new parent's wealth is the sum of its own savings w' and the bequest w'. We assume that the grandchild enters the economy with an initial wealth that is a function of its initial income realization, which we denote by $w_b(y'_j)$.

Computing the equilibrium is similar to the computation of our baseline model. The only difference is that the terms in δ are fed in to account for the risk of death/aging; to do this, we have to compute v^e at each iteration. A good initial guess is given by computing the value functions of a "final" pair of overlapping generations. These are the ones obtained from computing our baseline model with an increased discount rate due to the probability of death.

One specific problem we need to deal with is to obtain the continuation value for very large bequests. When such bequests are made, we often have to obtain v^e for levels of wealth that lie far outside the grid. An excellent method to extrapolate the function is by exploiting homogeneity. For families with large levels of wealth, we can safely neglect the income dimension, and thus assume that value functions and policies are homogeneous in wealth. Consumption and transfer policies are roughly linear in wealth for rich families, which translates into value functions being of form $W^{1-\gamma}$ in total family wealth (for details, see the homogeneous-altruism setting in Barczyk and Kredler (2014)). The old household's value function is then given by

$$v\big(w, w', y, y'\big) = \tilde{v}(P)W^{1-\gamma}, \quad \text{where } P = \frac{w}{W} \text{ and } W = w + w'.$$

The function $\tilde{v}$ can be calculated from the outermost grid points[12] from

$$\tilde{v}(P) = v\big(w, w', y, y'\big)W^{\gamma-1}.$$

[12]That is, the grid points where either the parent or the child household (or both) hold the maximal wealth $\bar{W}$ on the grid.

This gives us $\tilde{v}(P)$ on a finite grid; intermediate values can be approximated by linear interpolation. The P that realizes upon death of the parent household is given by

$$P = \frac{w + w'}{w_b(y'') + w + w'},$$

where y'' is the grandchild's first income realization.

Barczyk (2012) implements this OLG setting with a constant death/aging hazard to study the response of consumption to a deficit-financed tax cut.

S.3.2 *Finite horizon*

We will now show how to compute our baseline model for the case where the horizon is finite, that is, $t \in [0, T]$; the code for this example can be found in `AltruismUncert_FinHorizon.m` in the folder `Altruism`.

When the horizon is finite, value functions depend on time. Her HJB is

$$\begin{aligned} -v_t = {} & -\rho v + u(c^*) + \alpha u(c'^*) + \dot{w} v_w + \dot{w}' v_{w'} + \xi\big[v(\cdot, \tilde{y}) - v(\cdot, y)\big] \\ & + \xi\big[v(\cdot, \tilde{y}') - v(\cdot, y')\big] + \frac{\sigma^2}{2}\big(w^2 v_{ww} w'^2 v_{w'w'}\big), \end{aligned}$$

where we avoid the max operators by writing the optimal consumption rules and where we restrict attention to states where no player is broke so that transfers are zero. A tricky feature of computing a finite-horizon equilibrium here (as in many other continuous-time models) is how to proceed close to T. In our model, as in any consumption–savings model, consumption rates go to infinity when $t \to T$. This makes it hard to compute value functions. To deal with this problem, we make an assumption on what happens during the last Δt stage of the game. We assume that agents' income is no longer subject to risk and that agents together consume the resources left to them at a constant rate. So we have

$$c_{T-\Delta t}\Delta t + c'_{T-\Delta t}\Delta t = \underbrace{w_{T-\Delta t} + w'_{T-\Delta t} + y_{T-\Delta t}\Delta t + y'_{T-\Delta t}\Delta t}_{=W_{T-\Delta t}}.$$

When one agent owns all resources, we assume that she obtains her preferred allocation over Δt. So she consumes a fraction $A_1 = 1/(1 + \alpha^{1/\gamma})$ of $W_{T-\Delta t}$ when she owns all resources, and she obtains what is left from his preferred fraction, $A_0 = \alpha'^{1/\gamma}/(1 + \alpha'^{1/\gamma})$ of $W_{T-\Delta t}$, when he owns everything. In all situations in between, we linearly interpolate between these two extreme allocations using the fraction $P_{T-\Delta t} = (w_{T-\Delta t} + y_{T-\Delta t}\Delta t)/W_{T-\Delta t}$ she owns out of total resources:

$$\begin{aligned} c_{T-\Delta t} &= \big[(1 - P_{T-\Delta t})A_0 + P_{T-\Delta t}A_1\big]W_{T-\Delta t}, \\ c'_{T-\Delta t} &= \big[(1 - P_{T-\Delta t})(1 - A_0) + P_{T-\Delta t}(1 - A_1)\big]W_{T-\Delta t}. \end{aligned}$$

Using these consumption levels and for some given Δt, we compute the value functions as $v_{T-\Delta t} = [u(c_{T-\Delta t}) + \alpha u(c'_{T-\Delta t})]\Delta t$ and $v'_{T-\Delta t} = [u(c'_{T-\Delta t}) + \alpha' u(c_{T-\Delta t})]\Delta t$.

Other assumptions for the final Δt period are possible. We found that it is crucial, however, that value functions are strictly increasing in the agent's own asset share (i.e., v must be increasing in $P_{T-\Delta t}$). Specifically, one very natural assumption for the final Δt period leads to problems: if we let the agents play a static altruism game at $T - \Delta t$ (i.e., agents give transfers at $T - \Delta t$ and then consume what they have left), numerical instability arises: consumption functions become locally decreasing in agents' own assets. The problems is that the value functions at $T - \Delta t$ have strong convexities at the points where transfers start to flow.

Given value functions at $T - \Delta t$, we can backward iterate on the HJBs as in the baseline case, only this time we keep the results on the way. It is very important to make adjustments to the time increments of the algorithm so as to fulfill the stability conditions. Since consumption rates are very high close to $T - \Delta t$, the time increment required to keep the Markov chain's transition probabilities positive is very small.

Computationally, we find that, as expected, policies and value functions smoothly converge to their time-invariant counterparts. Transfer motives in the problematic region (between the overconsumption region and the SS region) are lowest close to T and then rise as we go back in time. So the equilibrium is even more stable in the finite-horizon than in the infinite-horizon case. This is in the sense that equilibrium exists for a larger range of σ when fixing a $(\alpha, \alpha', \gamma)$ combination.

For an application that uses this finite-horizon setting (in combination with OLG), see Barczyk and Kredler (2013), who study the macroeconomic effects of long-term-care policy.

S.3.3 *Endogenous risk-taking*

It is straightforward to adapt the Markov-chain method to the case with a risky asset: we determine the optimal risk-taking decision $z^* = I(v_{ww} > 0)$ from the finite approximation of the second derivative and then set the variance of the chain to $z^{*2} k^2 \sigma^2 / 2$. The rest is as before.

To implement this in the code `AltruismUncert_main.m`, the risk parameters have to be set to `SigmaNorm=0` and `SigmaRisky=`σ. It turns out that the following adjustments make the code more stable.

First, instead of using the discontinuous indicator function $I(\cdot)$ directly, we approximate it by a logistic smoother. This makes the law of motion continuous in the state. We take $h(x) = 1/(1 + e^{-\kappa x})$ and apply it to our measure of risk-lovingness: the first derivative of the optimal consumption function c^* with respect to the player's own assets, that is, we set $z^* = h(c^*_w)$. Note that c^*_w is a sign-preserving transformation of v_{ww}, but unlike v_{ww}, it is roughly of the same magnitude across the grid, which makes it better suited for our purposes. For the smoothing coefficient κ, we choose values around 10. The risk-taking decision is coded in the function `AltruismRiskTaking.m`.

Second, we only allow the poorer agent to take risks in the algorithm—this is what happens in equilibrium in the end. If we allow both agents to take risk along the way to convergence, it can happen that agents engage in local "risk-taking battles" that destabilize the algorithm.

S.3.4 *Adding more choice variables*

Adding more choice variables should generally not present large difficulties in our setting. The example with the portfolio decision in the previous subsection is one example. We present here the Bellman equation for a setting where she makes a continuous labor-supply decision. The wage rate is fixed at z. We assume that his income is still exogenous. The HJB is then

$$v = \max_{c,l}\big\{u(c,l) + (rw + zl - c)v_w\big\} + \cdots,$$

where we omit the rest of the terms because they are the same as before. The first-order condition for labor is

$$u_l(c,l) + zv_w = 0.$$

If utility is separable in consumption and labor, the computation of labor supply is especially simple. Even in the nonseparable case, the adaptation of the Markov-chain method is straightforward: once the optimal (c, l) pair is found for the Δt stage game, value-function updating works the same as before.

References

Barczyk, D. and M. Kredler (2014), "The ABC of incomplete-market economies with correlated shocks." Working paper. [3]

Barczyk, D. (2012), "Deficits, gifts, and bequests." Working paper. [20]

Barczyk, D. and M. Kredler (2013), "Macroeconomic implications of long-term-care policies." Working paper. [18, 21]

Barczyk, D. and M. Kredler (2014), "A dynamic model of altruistically-motivated transfers." *Review of Economic Dynamics*, 17, 303–328. [19]

Kushner , H. J. and P. Dupuis (2001), *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York. [6]

Tauchen, G. (1986), "Finite state Markov-chain approximations to univariate and vector autoregressions." *Economics Letters*, 20 (2), 177–181. [6, 8]

Submitted March, 2013. Final version accepted November, 2013.