

Federal Reserve Bank of Cleveland Working Paper Series

A Certified Howard Algorithm for Epstein–Zin Preferences

Anmol Bhandari, Jaroslav Borovička, and
Thomas Phelan

Working Paper No. 26-37

November 2026

Suggested citation: Bhandari, Anmol, Jaroslav Borovička, and Thomas Phelan. 2026. "A Certified Howard Algorithm for Epstein–Zin Preferences." Working Paper No. 26-37. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202637.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

See more working papers at: www.clevelandfed.org/research. Subscribe to email alerts to be notified when a new working paper is posted at: https://www.clevelandfed.org/subscriptions.

This work is licensed under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International. To view a copy of this license, visit https://creativecommons.org/licenses/by-nc-nd/4.0/

A Certified Howard Algorithm for Epstein–Zin Preferences

Anmol Bhandari* Jaroslav Borovička† Thomas Phelan‡

* University of Minnesota

† New York University

‡ Federal Reserve Bank of Cleveland

October 14, 2026

Abstract

This paper combines the variational utility representation of Geoffard (1996) and Dumas et al. (2000) with a variation of the Howard algorithm given in Bokanowski et al. (2009) to obtain a globally convergent algorithm for finite-dimensional Bellman equations arising from continuous-time problems with Epstein–Zin preferences. The variational representation of recursive utility introduces an auxiliary discount rate that makes policy evaluation linear. However, feasible discount policies can generate evaluation matrices that fail to be monotone. Our certification argument uses current value iterates or barrier functions to establish monotonicity of the matrices generated by the algorithm and provides conditions under which the value iterates converge monotonically to the unique admissible solution. We then illustrate the method in a consumption-savings problem with undiversifiable income risk and derive an explicit coefficient for the square-root behavior of savings near the borrowing constraint for agents with the lowest income. An extension to permanent income risk shows that the choice of initialization matters: the certified initial discount rates form an explicit interval that excludes the CRRA discount rate once permanent risk is large.

JEL classification: C63, D15, E21

Keywords: Epstein–Zin preferences, Howard's algorithm, consumption-savings problem, M-matrices

The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System. We thank Sergio Ocampo and Eric Young for helpful discussions and comments. Generative AI tools, including ChatGPT Astra, were used during the preparation of this manuscript to assist with mathematical arguments and exposition. The authors take full responsibility for its contents. Emails: bhandari@umn.edu, jb4457@nyu.edu, tom.phelan@clev.frb.org.

1 Introduction

Epstein–Zin preferences are commonly used in heterogeneous-agent macroeconomic models to allow risk aversion and intertemporal substitution to play distinct roles in household decisions. However, they also complicate computation: even for a fixed consumption function, the resulting Bellman equation is generally nonlinear in the value function. Consequently, fixing consumption no longer yields the linear policy-evaluation step used in conventional Howard iteration.

Our algorithm combines the variational representation of recursive utility (Geoffard 1996, Dumas et al. 2000) with the extension of Howard's algorithm to sup-inf problems given in Bokanowski et al. (2009). The variational representation introduces an auxiliary control, interpretable as an endogenous rate of discounting, and rewrites the recursive aggregator as either a supremum or an infimum over this control, depending on whether risk aversion exceeds the inverse elasticity of intertemporal substitution. For fixed values of consumption and the auxiliary discount rate, policy evaluation is once again a linear system. When the aggregator is a supremum, the problem is an ordinary control problem in the pair (consumption, discount rate), and Howard's algorithm applies as usual. When it is an infimum, the algorithm becomes nested: an *inner* loop of Howard iterations over the discount rate evaluates a fixed consumption policy, and an *outer* step then updates consumption. In either case, Howard's algorithm requires that each policy-evaluation matrix be monotone—invertible with a nonnegative inverse—and Bokanowski et al. (2009) assume this for *every* feasible policy. Because nothing ensures that the auxiliary discount rate is everywhere positive, this assumption is no longer automatic.

Our main contribution is to show that the monotonicity Howard's algorithm requires can be verified as the algorithm runs, rather than assumed for every feasible policy. Each step of the algorithm inverts a policy-evaluation matrix, and the step is well-behaved when that matrix has a nonnegative inverse. By a classical result on M-matrices, for the matrices arising here (whose off-diagonal entries are nonpositive), this is the case exactly when the matrix maps some strictly positive vector to a strictly positive vector; we call such a vector a *certificate*. The key observation is that the algorithm generates its own certificates. The policy-improvement step fixes the sign of the Bellman residual at the current value, and this sign is precisely what is needed for the current value itself (up to the sign of $1-\gamma$) to certify the matrix inverted in the next step. In two of the four parameter regions, delineated by whether γ exceeds one and whether γ exceeds ρ, this self-certification establishes the monotonicity of the next policy-evaluation matrix. In the other two regions, every feasible discount rate is positive, and a single fixed *barrier*—which in our applications is simply the utility of consuming at the highest (or lowest) feasible level forever—certifies every policy matrix at once. In either case, provided the algorithm is started from a policy whose

evaluation matrix is monotone, every subsequent matrix is monotone, and so under the stated barrier conditions, the value iterates increase monotonically and converge to the unique solution of the Bellman equation in the orthant on which the aggregator is defined. Our general framework applies to control problems with finitely many states and in which all controls (apart from the auxiliary discount rate) are restricted to compact sets. This is a setting that arises naturally from finite-difference discretizations of continuous-state HJB equations for models with Markovian uncertainty.

We apply our algorithm to a consumption-savings problem in which an agent with Epstein–Zin preferences faces fluctuating income and can save in a risk-free bond subject to borrowing constraints. For this application, a valid initialization is obtained by evaluating any feasible consumption policy with the auxiliary discount rate set equal to the subjective discount rate. This is simply the policy-evaluation step of the corresponding CRRA problem, and does not imply (or require) that subsequent auxiliary discount rates remain positive.

We then use the algorithm to study the savings behavior of agents near the borrowing constraint, a topic of considerable interest in the heterogeneous-agent literature. Building on the boundary analysis of Achdou et al. (2022) for CRRA preferences and Achdou and Tang (2026a) for Epstein–Zin preferences with late resolution of uncertainty, we characterize the savings of agents with the lowest income near the borrowing constraint. Under certain regularity assumptions, we show that such savings satisfy a "square-root" law and derive an expression for the associated coefficient that makes explicit the roles of intertemporal substitution, endogenous discounting, and marginal continuation values across income states. Quantitatively, the coefficient is more sensitive to changes in ρ than to changes in γ, and it appears to vary smoothly through the parameter boundaries $\gamma = 1$ and $\gamma = \rho$ at which the algorithm switches between its four cases.

We then document the practical performance of the algorithm in this income-fluctuation problem. In our experiments, each outer iteration adds roughly one and a half to two decimal digits of accuracy in consumption-equivalent units, so that high accuracy obtains within roughly ten iterations, and the discretization appears to converge at the first-order rate expected of upwind schemes.

Finally, we show that the requirement of a certified starting policy has bite once income has a permanent component, including large and persistent earnings losses such as those following job displacement. With substantial permanent risk, the CRRA initialization is no longer certified, and from natural starting values Newton's method fails while implicit time stepping requires small steps and thousands of iterations. Our theory characterizes exactly which constant initial discount rates are certified and suggests a natural one, the discount rate of an agent who consumes her permanent income. From it, the algorithm

converges monotonically in at most ten iterations, up to the point at which utility becomes infinite. The difficulty lies in the starting value rather than in the method: from a certified start, time stepping converges as well, while simply choosing a large initial discount rate is no remedy, because the certified range is bounded above.

Related literature. The continuous-time theory of recursive utility begins with the stochastic differential utility of Duffie and Epstein (1992). The variational (or Fenchel–Legendre) representation of recursive preferences was introduced by Geoffard (1996) in a deterministic setting and extended to uncertainty by Dumas et al. (2000).[1] These representations turn the original recursive problem into a sup-inf (or sup-sup) control problem in an auxiliary discount-rate process. On the existence and uniqueness of recursive utility, Epstein and Zin (1989) and Marinacci and Montrucchio (2010) give sufficient conditions under bounded consumption growth, Hansen and Scheinkman (2012) connect existence to a Perron–Frobenius eigenvalue problem, Borovička and Stachurski (2020) obtain necessary and sufficient spectral conditions together with a globally convergent computational method, Christensen (2022) extends existence and uniqueness to unbounded state spaces and unbounded flow utilities under a thin-tail condition on utility growth, and Pohl et al. (2019) and Pohl et al. (2024) establish existence by comparison with a CRRA benchmark.

On the numerical side, upwind finite-difference schemes and policy iteration for continuous-time heterogeneous-agent models with additive utility are by now standard (Achdou et al. 2022). Howard's algorithm is Newton's method applied to the Bellman equation (Puterman and Brumelle 1979), and its monotone convergence rests on the policy-evaluation matrices having nonnegative inverses; Bokanowski et al. (2009) establish convergence in this setting and introduce the nested iteration for sup-inf problems that we build on, under the assumption that *every* feasible policy generates such a matrix.

The closest work is Achdou and Tang (2026a) and Achdou and Tang (2026b). The first paper develops the constrained viscosity-solution theory for the consumption-savings problem with two income states under a preference for late resolution ($1 < \gamma < \rho$ in our notation)[2] and characterizes savings at the borrowing constraint. The second proposes upwind schemes for $\gamma > 1$ and $\rho > 1$ and proves convergence of two algorithms: for $\gamma \leq \rho$, a Howard–Newton iteration whose policy-evaluation step is an inner loop of Newton steps—the same iteration as ours in that case—and for $\gamma > \rho$, a Howard–Tarski–Kantorovich iteration in which a monotone fixed-point map is applied between full Howard solves; in the former case they also bound the distance between the discrete and continuous solutions.

Our contribution is complementary. We work with a general finite-state residual rather

[1] See also El Karoui et al. (1997).

[2] A preference for late resolution corresponds to $\gamma < \rho$ in our notation: relative risk aversion is smaller than the inverse elasticity of intertemporal substitution.

than a particular scheme, allow $\gamma < 1$ and $\rho < 1$, and obtain global monotone convergence from a certified initialization without requiring every policy matrix to be monotone. For late resolution this supplies the convergence proof for the inner Newton loop; for early resolution our algorithm is ordinary Howard iteration in the pair (c, v) rather than a fixed-point iteration, and the comparison argument delivers uniqueness in the admissible orthant, which Achdou and Tang (2026b) leave open. Our convergence results concern a fixed discretization; Appendix C establishes convergence under grid refinement for the bounded-asset, bounded-consumption problem by viscosity-solution methods, without a rate. In discrete time, Ren and Stachurski (2018) prove global convergence of value function iteration for Epstein–Zin dynamic programs in the early-resolution case, using monotone convex-operator methods and value functions bracketed between a lower and a strict upper solution—the discrete-time counterpart of the barriers employed here.

The remainder of the paper is organized as follows. Section 2 considers a two-state example in which both the difficulties alluded to above and their resolution can be verified with pen and paper. Section 3 develops the general theory and establishes global convergence. Section 4 considers a consumption-savings problem, derives an expression for the square-root savings law and illustrates the performance of the algorithm numerically, and extends the application to permanent income risk. Section 5 concludes the paper.

2 Motivating example

Prior to the general analysis, we consider a two-state endowment economy with recursive preferences that is simple enough for every step to be carried out with pen and paper. The example makes three points. First, despite the nonlinearity introduced by recursive preferences, an auxiliary discount rate turns the Bellman equation into a standard control problem to which Howard's algorithm (policy function iteration) applies. Second, a naive application of Howard's algorithm can fail: the policy-improvement step may be undefined, and altering the control set so that every update is well-defined can produce convergence to the wrong value. Third, these failures are avoidable: if policy evaluation at the initial policy is a monotone linear system (the matrix to be inverted has a nonnegative inverse), then the same is true at every subsequent update, and the values increase monotonically toward the solution.

Preferences. We use the continuous-time Epstein–Zin aggregator $f(c, w)$, where γ governs relative risk aversion and ρ the inverse elasticity of intertemporal substitution; we defer a description of its general form to Section 3. To make the example as transparent as possible, we set $\gamma = 3/2$ and $\rho = 1/2$, for which the aggregator reduces to the quadratic

form

$$f(c,w) = \delta\left(\sqrt{c}\frac{w^2}{2} + w\right). \tag{1}$$

Technology. The exogenous state $x_t \in \{1,2\}$ follows a continuous-time Markov chain in which state 1 switches to state 2 with intensity $\phi > 0$ and state 2 is absorbing. The endowment is $C_t = e^{y_t}$ and

$$dy_t = \mu_y(x_t)dt + \sigma_y(x_t)dB_t,$$

for some state-dependent drift μ_y and diffusion σ_y of log output.

The HJB equation for the value function is $0 = f(e^y, W) + \mathcal{L}W$, where f is the quadratic aggregator in (1) and $\mathcal{L}$ is the generator of (y_t, x_t). We guess-and-verify that the value function W satisfies $W(y,x) = w(x)e^{-y/2}$ for some w and substitute into the HJB equation to obtain

$$\begin{aligned} 0 &= \frac{\delta}{2}w_1^2 + (\delta + h_1 - \phi)w_1 + \phi w_2 \\ 0 &= \frac{\delta}{2}w_2^2 + (\delta + h_2)w_2 \end{aligned} \tag{2}$$

where $h_i = (1-\gamma)\mu_y(x_i) + (1-\gamma)^2\sigma_y(x_i)^2/2$ for $i = 1,2$. The system admits a solution with both components strictly negative if and only if

$$\delta + h_2 > 0 \tag{3}$$

in which case $w_2 = -2(1 + h_2/\delta)$ and w_1 is the negative solution to

$$0 = \frac{\delta}{2}w_1^2 + (\delta + h_1 - \phi)w_1 - 2\phi(1 + h_2/\delta).$$

Application of Howard. Now suppose that we are unaware of the closed-form solution and wish to solve the system of equations (2) iteratively, under the maintained assumption (3). In contrast with the case of CRRA utility, the right-hand side is nonlinear in w. However, elementary algebra shows that for $w < 0$, the quadratic function can be written as a maximum over affine functions,

$$w^2 = \max_{q>0} -q^2 - 2qw \tag{4}$$

which implies that for $D_\nu := (-\delta, \infty)$, we have

$$\frac{\delta}{2}w^2 + \delta w = \max_{\nu \in D_\nu} -\frac{(\nu+\delta)^2}{2\delta} - \nu w. \tag{5}$$

The value function $w = (w_1, w_2)$ then solves the Bellman equation

$$0 = \max_{\nu \in D_\nu^2} F(\nu) + T(\nu)w, \tag{6}$$

where $F(\nu) := -(\nu + \delta)^2/(2\delta)$ and

$$-T(\nu) = \begin{pmatrix} \nu_1 - h_1 + \phi & -\phi \\ 0 & \nu_2 - h_2 \end{pmatrix}.$$

A natural approach to solving an equation of the form (6) is the familiar policy function iteration (PFI) algorithm of Howard (1960), which we now recall for ease of reference.

Algorithm 2.1 (PFI; Howard (1960))**.** *Given a tolerance $\epsilon > 0$:*

1. *Choose an initial $\nu^0 \in D_\nu^2$.*
2. *Find w^0 satisfying $0 = F(\nu^0) + T(\nu^0)w^0$.*
3. *Update ν^1 to maximize $\nu \mapsto F(\nu) + T(\nu)w^0$, return to Step 2 with ν^0 replaced by ν^1, and repeat until convergence of successive w terms within tolerance ϵ.*

We now apply Algorithm 2.1 to the Bellman equation (6) for a given initial control ν^0. First note that the value of w_2^0 as a function of ν_2 is

$$w_2^0(\nu_2) = -\frac{(\nu_2 + \delta)^2}{2\delta(\nu_2 - h_2)}$$

which, upon substitution, gives w_1^0 as a function of ν,

$$w_1^0(\nu_1, \nu_2) = -\frac{(\nu_1 + \delta)^2 + \phi(\nu_2 + \delta)^2/(\nu_2 - h_2)}{2\delta(\nu_1 - h_1 + \phi)}.$$

Given (w_1^0, w_2^0), the optimal ν^1 from (5) will fail to be well-defined if $w_2^0 > 0$, which occurs when $\nu_2^0 < h_2$. For such initial conditions, Algorithm 2.1 breaks down entirely.

Now consider a candidate "fix" for this problem, in which we replace the set D_ν in (6) with its closure $\overline{D}_\nu = [-\delta, \infty)$ and apply Algorithm 2.1. Then $\nu_i^1 = -\delta + \max\{-\delta w_i^0, 0\}$ and the updated ν^1 can be written as a function of the previous ν,

$$\begin{aligned} \nu_2^1(\nu_2) + \delta &= \max\left\{\frac{(\nu_2 + \delta)^2}{2(\nu_2 - h_2)}, 0\right\} \\ \nu_1^1(\nu) + \delta &= \max\left\{\frac{(\nu_1 + \delta)^2 + \phi(\nu_2 + \delta)^2/(\nu_2 - h_2)}{2(\nu_1 - h_1 + \phi)}, 0\right\}. \end{aligned} \tag{7}$$

We now examine the behavior of the map (7). First note that if $\nu_2^0 < h_2$ then $\nu_2^n = -\delta$ for all $n \geq 1$, while if $\nu_2^0 > h_2$ then

$$\nu_2^1(\nu_2) - \nu_2 = \frac{(h_2+\delta)^2 - (\nu_2 - h_2)^2}{2(\nu_2 - h_2)}.$$

It follows that $\nu_2^1(\nu_2) > \nu_2$ if $2h_2 + \delta > \nu_2$ and $\nu_2^1(\nu_2) < \nu_2$ if $2h_2 + \delta < \nu_2$, which implies that $(\nu_2^n)_{n\geq 0}$ converges to $-\delta$ if $\nu_2^0 < h_2$ and converges to $\delta + 2h_2$ if $\nu_2^0 > h_2$. Similarly, if $\nu_2^0 > h_2$ and $\nu_1^0 > h_1 - \phi$, then $(\nu_1^n)_{n\geq 0}$ converges to the larger of the two roots of

$$0 = \frac{(\nu_1+\delta)^2}{2} - (\nu_1+\delta)(\nu_1 - h_1 + \phi) + 2\phi(\delta + h_2)$$

while if $\nu_1^0 < h_1 - \phi$ then $\nu_1^n = -\delta$ for $n \geq 1$. It follows that even if we change the constraint set D_ν in (6) so that policy updating is well-defined, the associated sequence may converge to the wrong value for some initial controls.

However, despite illustrating some problems with the naive application of Algorithm 2.1 from any initial condition, the above also hints at a resolution. To see this, note that by the definitions of w^n, ν^{n+1} and w^{n+1}, we have

$$F(\nu^{n+1}) + T(\nu^{n+1})w^{n+1} = 0 \leq B(w^n) = F(\nu^{n+1}) + T(\nu^{n+1})w^n$$

and hence

$$w^{n+1} - w^n = -T(\nu^{n+1})^{-1}B(w^n). \tag{8}$$

Because $0 \leq B(w^n)$ for all n, equation (8) shows that a sufficient condition for $w^{n+1} \geq w^n$ is if $-T(\nu^{n+1})^{-1} \geq 0$ componentwise. Further, the above shows that $-T(\nu^1)^{-1} \geq 0$ if $-T(\nu^0)^{-1} \geq 0$, which, by (8), ensures that the sequence $(w^n)_{n\geq 0}$ is increasing. The analysis that follows shows that this is not a coincidence, and that under certain conditions on the initial policy guesses, we can ensure monotonic convergence of our updates.

Doing it right. In this example, the appropriate initial conditions are transparent. Because $-T(\nu)$ is upper triangular, $-T(\nu)^{-1} \geq 0$ if and only if $\nu_2 > h_2$ and $\nu_1 > h_1 - \phi$: the thresholds are precisely the growth-adjusted rates at which the failures above originate. Suppose we initialize anywhere in this region, leaving the constraint set D_ν untouched. Writing $t := \nu_2^0 - h_2 > 0$ and $s := h_2 + \delta > 0$, the update (7) gives

$$\nu_2^1 + \delta = \frac{(t+s)^2}{2t} = \frac{t}{2} + s + \frac{s^2}{2t} \geq 2s,$$

with equality if and only if $t = s$. A single update therefore lands at or (weakly) above the fixed point $\delta + 2h_2$, every subsequent iterate remains in the monotone region, and $(\nu_2^n)_{n\geq 1}$

decreases monotonically to $\delta + 2h_2$, the optimal discount rate associated with the correct root $w_2 = -2(1 + h_2/\delta)$. Combined with the certification step noted above ($-T(\nu^{n+1})^{-1} \geq 0$ whenever $-T(\nu^n)^{-1} \geq 0$), equation (8) then implies that the value sequence $(w^n)_{n\geq 0}$ increases monotonically to the solution of (2) with the correct signs. Two lessons carry over to the general theory. Monotonicity of the initial policy matrix suffices (no modification of the control set is required) and the natural default $\nu^0 = \delta$ qualifies only when δ exceeds the growth-adjusted rates h_2 and $h_1 - \phi$; with growth in the environment, safe initial discount rates must be chosen accordingly, a point we revisit in Section 3.

3 General theory

We now outline a more general theory and discuss conditions under which convergence can be guaranteed. Throughout the section we assume $\delta, \gamma, \rho > 0$, $\gamma \neq 1$, and $\gamma \neq \rho$. The case $\rho = 1$ is included using the logarithmic aggregator in (10). The case $\gamma = \rho$ reduces to the standard CRRA problem and is covered in Achdou et al. (2022). All vector inequalities below are componentwise, and $\gg$ denotes strict inequality in every component.

3.1 Environment

In this section we define our dynamic programming equation formally and fix notation. Let $X := \{x_1, \ldots, x_M\}$ be the finite state space, where $M \geq 1$ is the number of states, and define the space of values

$$\mathcal{W} := \Big\{ w \in \mathbb{R}^M \Big| (1-\gamma)w \gg 0 \Big\}.$$

For scalar arguments $c > 0$ and $(1-\gamma)w > 0$, the Epstein–Zin aggregator is

$$f(c, w) = \delta \frac{(1-\gamma)}{1-\rho} w \left(c^{1-\rho} [(1-\gamma)w]^{-\frac{1-\rho}{1-\gamma}} - 1 \right) \tag{9}$$

for $\gamma, \rho \neq 1$, and for $\rho = 1$,

$$f(c, w) = \delta(1-\gamma)w \left(\ln c - \frac{1}{1-\gamma} \ln((1-\gamma)w) \right). \tag{10}$$

For $\rho \neq 1$ we define

$$\theta := \frac{1-\gamma}{1-\rho},$$

and note that $\theta = 1$ if and only if $\gamma = \rho$, the CRRA case excluded above. The quadratic aggregator (1) of Section 2 corresponds to $\gamma = 3/2$ and $\rho = 1/2$. Throughout this paper

$$u(c) := \frac{c^{1-\gamma}}{1-\gamma}$$

denotes the CRRA felicity function, and we record the identity $f(c, u(c)) = 0$ for every $c > 0$: the utility of a constant consumption stream c is $u(c)$. For $c \in (0,\infty)^M$ and $w \in \mathcal{W}$, we write $f(c, w)$ for the vector obtained by applying the aggregator state by state. The HJB equation is given by $0 = \max_c f(c, W) + \mathcal{L}^c W$ for some controlled generator $\mathcal{L}^c$, and so we will consider discrete analogues of the form

$$0 = B(w) := \max_{c \in \mathcal{C}} G_c(w) = \max_{c \in \mathcal{C}} f(c, w) + A(c)w \tag{11}$$

for some constraint set $\mathcal{C}$ and function A, where the maximum in (11) is interpreted componentwise. We will later consider explicit functional forms for the generator $\mathcal{L}^c$ and discrete analogue A, but for now we simply assume the following.

Assumption 3.1. *The policy set is $\mathcal{C} = \prod_{x \in X} C_x$, where each $C_x \subset (0,\infty)$ is nonempty and compact. The map $c \mapsto A(c)$ is continuous, the off-diagonal entries of $A(c)$ are nonnegative, and row x of $A(c)$ depends on c only through c_x.*

We will later see that the "upwinding" procedure from the numerical PDE literature and Achdou et al. (2022) produces discrete Bellman equations satisfying Assumption 3.1.[3]

The arguments below require *barriers*: elements of $\mathcal{W}$ at which the residual of every feasible policy has a known sign. The upper barrier $\overline{w}$ caps the increasing sequence of outer values in both cases, while the lower barrier $\underline{W}$ floors the decreasing inner values when the optimization over ν is a minimization; together they keep every iterate in $\mathcal{W}$, so no approximate value ever changes sign.

Assumption 3.2 (supersolution)**.** *There exists $\overline{w} \in \mathcal{W}$ such that $G_c(\overline{w}) \leq 0$ for every $c \in \mathcal{C}$.*

Assumption 3.3 (subsolution)**.** *There exists $\underline{W} \in \mathcal{W}$ such that $G_c(\underline{W}) \geq 0$ for every $c \in \mathcal{C}$.*

Because B is a componentwise maximum, Assumption 3.2 is equivalent to $B(\overline{w}) \leq 0$, so $\overline{w}$ is a supersolution of the Bellman equation (11). Barriers are also easy to exhibit in practice. We call A *conservative* if $A(c)\mathbf{1} = 0$ for every $c \in \mathcal{C}$, the discrete analogue of the generator of a Markov process without killing. In the conservative case, constant barriers can be constructed from bounds on feasible consumption.

[3] In particular, the requirement that off-diagonal components of $A(c)$ are nonnegative amounts to the assumption that the choice of finite differences is determined by the direction of the drift.

Lemma 3.1 (constant barriers)**.** *Suppose that Assumption 3.1 holds and that A is conservative. If $0 < c_- \leq c_x \leq c_+ < \infty$ for every $c \in \mathcal{C}$ and state x, then $\underline{w} := u(c_-)\mathbf{1}$ and $\overline{w} := u(c_+)\mathbf{1}$ satisfy Assumptions 3.3 and 3.2, respectively.*

Proof. Both vectors lie in $\mathcal{W}$ by inspection, and because $A(c)\mathbf{1} = 0$ for any $c \in \mathcal{C}$, we have $G_c(\kappa\mathbf{1})_x = f(c_x, \kappa)$ for every scalar κ and state x. Because f is strictly increasing in consumption and $f(z, u(z)) = 0$, we have $G_c(\overline{w})_x = f(c_x, u(c_+)) \leq f(c_+, u(c_+)) = 0$ and $G_c(\underline{W})_x = f(c_x, u(c_-)) \geq f(c_-, u(c_-)) = 0$ for every $c \in \mathcal{C}$ and state x, which are Assumptions 3.2 and 3.3. ∎

There are three parts to how the general theory builds on the motivating example. First, Lemma 3.2 introduces an auxiliary control and shows that the HJB equation can be written as the pointwise optimization over linear operators. Second, using a well-known result from linear algebra (Lemma 3.3), Lemma 3.4 shows how the monotonicity of each updating matrix can be established by choosing certificates judiciously. Third, Lemma 3.5 is a "comparison" result which will allow us to pass from orderings of the values of G_c and B to ordering of their arguments (the approximate value functions). Section 3.4 then defines our algorithm formally and proves monotonicity and convergence under conditions stated below.

3.2 Variational utility

In the motivating example in Section 2, we exploited the identity (4), which allowed us to write the quadratic HJB equation as the pointwise maximization of linear operators over an auxiliary control ν. This is a special case of the following more general Fenchel–Legendre duality, due to Geoffard (1996) and Dumas et al. (2000).

Lemma 3.2 (variational utility duality)**.** *For every w satisfying $(1-\gamma)w > 0$,*

$$f(c, w) = \begin{cases} \max_{\nu \in D_\nu} F(c, \nu) - \nu w, & \gamma > \rho, \\ \min_{\nu \in D_\nu} F(c, \nu) - \nu w, & \gamma < \rho \end{cases} \tag{12}$$

where $F(c, \nu) := \delta u(c) g(\nu)$ for u as defined in Section 3.1 and g given by

$$g(\nu) := \begin{cases} \left(\dfrac{\nu/\delta - \theta}{1-\theta}\right)^{1-\theta}, & \rho \neq 1, \\ e^{\nu/\delta - 1}, & \rho = 1, \end{cases}$$

and the extremum is attained at a unique policy $v^*(c,w) \in D_v$, where

$$D_v := \begin{cases} (\delta\theta, \infty), & \rho \neq 1 \text{ and } \theta < 1, \\ (-\infty, \delta\theta), & \rho \neq 1 \text{ and } \theta > 1, \\ \mathbb{R}, & \rho = 1. \end{cases}$$

Proof. See Appendix B. ■

Applying Lemma 3.2 state-by-state then gives an alternative (and more convenient) representation of the function $B(w)$ given in (11).

Definition 1 (nested Bellman formulation)**.** *For $c \in \mathcal{C}$, $v \in D_v^M$, and $w \in \mathcal{W}$, we define $Q(c,v) := \operatorname{diag}(v) - A(c)$, $T(c,v) := -Q(c,v)$, and the residual of policy (c,v) at w by*

$$R(c,v;w) := F(c,v) + T(c,v)w.$$

In this notation the Bellman equation can be written $B(w^*) \equiv 0$, where

$$B(w) = \begin{cases} \max\limits_{c \in \mathcal{C}} \max\limits_{v \in D_v^M} R(c,v;w), & \text{if } \gamma > \rho, \\ \max\limits_{c \in \mathcal{C}} \min\limits_{v \in D_v^M} R(c,v;w), & \text{if } \gamma < \rho. \end{cases} \tag{13}$$

Our analysis will use the formulation of the Bellman equation in (13) and proceed case by case. Two binary distinctions organize the argument. The first is the sign of $\gamma - \rho$, which by Lemma 3.2 determines whether the auxiliary discount rate is a maximizing or a minimizing control:

1. if $\gamma > \rho$, the optimization over v in (13) is a maximization, and the algorithm below is ordinary Howard iteration in the pair (c,v);
2. if $\gamma < \rho$, it is a minimization, and the algorithm acquires an inner loop over v for each fixed consumption policy.

The second is the sign of $1 - \gamma$, which fixes the sign of the flow payoff and hence the orthant in which values lie: under high risk aversion ($\gamma > 1$) we have $F \ll 0$ and $w \ll 0$ for $w \in \mathcal{W}$, while under low risk aversion ($\gamma < 1$) both are positive. The four combinations call for different arguments, and are summarized in Section 3.5.

3.3 Certificates and comparison

Recall that a matrix is a Z-matrix if its off-diagonal entries are nonpositive. A matrix Q is *monotone* if $Qx \leq Qy$ implies $x \leq y$. For $c \in \mathcal{C}$, let

$$\mathcal{N}(c) := \left\{ \nu \in D_\nu^M \middle| Q(c,\nu) \text{ is monotone} \right\}.$$

The following is the standard positive-vector criterion for nonsingular M-matrices. See Plemmons (1977) or Theorem 6.2.3 of Berman and Plemmons (1994).

Lemma 3.3 (certificate criterion)**.** *For any Z-matrix Q, the following are equivalent: (i) Q is monotone; (ii) Q is nonsingular and $Q^{-1} \geq 0$; and (iii) there exists $z \gg 0$ such that $Qz \gg 0$. When these conditions hold, the diagonal entries of Q^{-1} are strictly positive.*

Given a Z-matrix Q, a vector $z \gg 0$ satisfying $Qz \gg 0$ is referred to as a *certificate* for Q. We now characterize the certificates that will be used to establish monotonic convergence. These are of three types: initial certificates, self-certificates generated by previous value iterates, and barrier-generated certificates.

Lemma 3.4 (certificates)**.** *Under Assumption 3.1, we have the following certificates.*

1. *Initial certificates: if $\nu \gg A(c)\mathbf{1}$ then $Q(c,\nu)$ is monotone with certificate $\mathbf{1}$.*
2. *Self-certificates: for every $c \in \mathcal{C}$, $\nu \in D_\nu^M$, and $w \in \mathcal{W}$, if $(1-\gamma)R(c,\nu;w) \leq 0$, then $Q(c,\nu)$ is monotone with certificate $(1-\gamma)w$.*
3. *Barrier-generated certificates:*
   (a) *In Case 1 with $\gamma < 1$, if Assumption 3.2 holds then $Q(c,\nu)$ is monotone for all $(c,\nu) \in \mathcal{C} \times D_\nu^M$ with certificate $(1-\gamma)\overline{w}$.*
   (b) *In Case 2 with $\gamma > 1$, if Assumption 3.3 holds then $Q(c,\nu)$ is monotone for all $(c,\nu) \in \mathcal{C} \times D_\nu^M$ with certificate $(1-\gamma)\underline{w}$.*

Proof. By the definition of Q we have $Q(c,\nu)\mathbf{1} = \nu - A(c)\mathbf{1}$, and so part 1 follows from Lemma 3.3. For part 2, multiplying the definition of R by $1-\gamma$ and rearranging gives

$$Q(c,\nu)[(1-\gamma)w] = (1-\gamma)F(c,\nu) - (1-\gamma)R(c,\nu;w).$$

The first term on the right is strictly positive by Lemma 3.2. The desired residual inequality therefore implies $Q(c,\nu)[(1-\gamma)w] \gg 0$ and so $(1-\gamma)w$ is a certificate for Q. For part 3a, we have $R(c,\nu;\overline{w}) \leq G_c(\overline{w}) \leq 0$ and $\overline{w}$ is a certificate for Q, while in part 3b, we have $R(c,\nu;\underline{W}) \geq G_c(\underline{W}) \geq 0$ and $(1-\gamma)\underline{W}$ is a certificate for Q. ■

The certificates in Lemma 3.4 concern one policy matrix at a time. To turn the monotonicity of individual matrices into an ordering of candidate value functions—and, ultimately, of the iterates generated by the algorithm—we need a comparison principle: whenever one candidate value receives a nonnegative residual and another a nonpositive one, the two are ordered. The key step in the proof is to evaluate the auxiliary discount rate at whichever of the two candidates makes a self-certification argument available, so that comparison holds even though some feasible policy matrices fail to be monotone.

Lemma 3.5 (discrete comparison)**.** *Let Assumption 3.1 hold and fix $c \in \mathcal{C}$ and $u, w \in \mathcal{W}$. Then $G_c(u) \geq 0 \geq G_c(w) \implies u \leq w$ and $B(u) \geq 0 \geq B(w) \implies u \leq w$, and so (11) has at most one solution in $\mathcal{W}$.*

Proof. In Case 1, if $\nu = \nu^*(c, u)$ then $R(c, \nu; u) = G_c(u) \geq 0$ and $R(c, \nu; w) \leq G_c(w) \leq 0$, while in Case 2, if $\nu = \nu^*(c, w)$ then $R(c, \nu; w) = G_c(w) \leq 0$ and $R(c, \nu; u) \geq G_c(u) \geq 0$. In either case, we have

$$Q(c, \nu)u \leq F(c, \nu) \leq Q(c, \nu)w. \tag{14}$$

If $\gamma > 1$, the residual at u satisfies the self-certification part of Lemma 3.4; if $\gamma < 1$, the residual at w does. Thus $Q(c, \nu)$ is monotone. Applying its nonnegative inverse to (14) gives $u \leq w$. For the second claim, choose $c_u \in \mathcal{C}$ attaining $B(u)$ and note that $G_{c_u}(u) = B(u) \geq 0$ and $G_{c_u}(w) \leq B(w) \leq 0$. The first part of the lemma then gives the result. ∎

3.4 Certified Howard iteration

We now state the algorithm analyzed below. In Case 1, it is ordinary Howard iteration in the control (c, ν), and in Case 2, it is the Ho–3 Algorithm of Bokanowski et al. (2009): an inner Howard loop solves the minimization over ν for a fixed consumption policy, and the outer step improves consumption.

Algorithm 3.6 (certified Howard iteration)**.** *Choose an initial consumption policy $c^0 \in \mathcal{C}$.*

Case 1. *Choose $\nu^0 \in \mathcal{N}(c^0)$. For $n = 0, 1, \ldots$:*

1. *Evaluate the current policy by solving $Q(c^n, \nu^n)w^n = F(c^n, \nu^n)$.*
2. *Choose $c^{n+1} \in \arg\max_{c \in \mathcal{C}} G_c(w^n)$, $\nu^{n+1} := \nu^*(c^{n+1}, w^n)$, and return to Step 1.*

Case 2. *For $n = 0, 1, \ldots$:*

1. *Choose $\nu^{n,0} \in \mathcal{N}(c^n)$, and for $m = 0, 1, \ldots$, solve $Q(c^n, \nu^{n,m})w^{n,m} = F(c^n, \nu^{n,m})$ and set $\nu^{n,m+1} := \nu^*(c^n, w^{n,m})$. Denote the limit of this inner loop by $w^n := \lim_m w^{n,m}$.*

2. *Choose* $c^{n+1} \in \arg\max_{c \in \mathcal{C}} G_c(w^n)$ *and return to Step 1.*

Remark 1. *We emphasize that in Step 2 of Case 2, we are maximizing the objective* $G_c(w^n)$ *over* c*, and not the objective* $R(c, \nu^*(c^n, w^n); w^n)$*, and hence are varying* ν *to satisfy the interior optimum. This is an important distinction, because the proof below only uses the fact that* $Q(c^n, \nu^*(c^n, w^n))$ *is monotone, and makes no claims on* $Q(c^{n+1}, \nu^*(c^{n+1}, w^n))$*.*

Remark 2. *As noted in Puterman and Brumelle (1979), Howard's algorithm can be viewed as Newton's method applied to the function* B*.*[4] *As emphasized by Remark 5.3 of Bokanowski et al. (2009), this is no longer true in the nested case.*

3.5 Monotonicity and convergence

In this section we state and prove the main results of the paper: under conditions specified below, Algorithm 3.6 converges to the unique solution to the discrete Bellman equation. For clarity, we separate the monotonicity claim from the convergence claim.

The following table summarizes the case-by-case analysis. The two cases that use the current value are the parameter regions in which D_ν contains negative discount rates. In the other two cases, $\theta \in (0, 1)$, every feasible discount rate is positive, and a fixed barrier certifies every policy matrix.

	$\gamma > 1$	$\gamma < 1$
$\gamma > \rho$ (sup)	Current w certifies each update.	Supersolution certifies each update.
$\gamma < \rho$ (inf)	Subsolution certifies each update.	Current w certifies each update.

Proposition 3.7 (monotonicity of iterates)**.** *Let Assumptions 3.1 and 3.2 hold. In Case 1, suppose* $\nu^0 \in \mathcal{N}(c^0)$*. In Case 2, also let Assumption 3.3 hold, suppose* $\mathcal{N}(c)$ *is nonempty for every* $c \in \mathcal{C}$*, and that* $\nu^{n,0} \in \mathcal{N}(c^n)$ *for all* n*. Then:*

(i) in Case 2, for every n*,* $\underline{W} \leq w^{n,m+1} \leq w^{n,m}$*, the inner loop converges to the limit* w^n *that is the unique solution in* $\mathcal{W}$ *of* $G_{c^n}(w) = 0$*;*

(ii) in both cases, the outer values satisfy $w^n \leq w^{n+1} \leq \overline{w}$*.*

Proof. We first note that if $Q(c, \nu)w = F(c, \nu)$ and $Q(c, \nu)^{-1} \geq 0$, then

$$(1 - \gamma)w = Q(c, \nu)^{-1}[(1 - \gamma)F(c, \nu)] \gg 0. \tag{15}$$

[4] To see why, note that when the problem is of "sup-sup" type, given $w_t \in \mathcal{W}$, the envelope theorem implies that if (c_{t+1}, ν_{t+1}) uniquely attains the maximum in B then the Jacobian is $\nabla B(w_t) = T(c_{t+1}, \nu_{t+1})$, and so the updating step can be written $w_{t+1} = -\nabla B(w_t)^{-1}(B(w_t) - T(c_{t+1}, \nu_{t+1})w_t) = w_t - \nabla B(w_t)^{-1}B(w_t)$.

The strict inequality follows from Lemma 3.2 and the strictly positive diagonal of $Q(c,\nu)^{-1}$, and so every monotone evaluation produces a value in $\mathcal{W}$.

Consider Case 1. The initial Q matrix is monotone by construction. If the nth Q matrix is monotone then equation (15) gives $w^n \in \mathcal{W}$, and policy improvement gives

$$\begin{aligned} R(c^{n+1},\nu^{n+1};w^n) &= G_{c^{n+1}}(w^n) = B(w^n) \\ &\geq G_{c^n}(w^n) \geq R(c^n,\nu^n;w^n) = 0. \end{aligned} \tag{16}$$

If $\gamma > 1$, part 2 of Lemma 3.4 shows that $(1-\gamma)w^n$ is a certificate for $Q(c^{n+1},\nu^{n+1})$, while if $\gamma < 1$, the same matrix is monotone by part 3 of Lemma 3.4. Thus every policy-evaluation matrix in Case 1 is monotone, by induction. Further,

$$w^{n+1} - w^n = Q(c^{n+1},\nu^{n+1})^{-1} R(c^{n+1},\nu^{n+1};w^n) \geq 0.$$

Because the current policy is feasible, we have $B(w^{n+1}) \geq R(c^{n+1},\nu^{n+1};w^{n+1}) = 0$ and so Lemma 3.5 gives $w^{n+1} \leq \overline{w}$.

Now consider Case 2 and fix an integer n. The first matrix of the inner loop is monotone by construction, and if the mth matrix is monotone, then $w^{n,m} \in \mathcal{W}$ by (15), and

$$R(c^n,\nu^{n,m+1};w^{n,m}) = G_{c^n}(w^{n,m}) \leq R(c^n,\nu^{n,m};w^{n,m}) = 0. \tag{17}$$

If $\gamma < 1$, $(1-\gamma)w^{n,m}$ is a certificate by Lemma 3.4, while if $\gamma > 1$, Lemma 3.4 certifies it. Thus every inner matrix is monotone and

$$w^{n,m+1} - w^{n,m} = Q(c^n,\nu^{n,m+1})^{-1} R(c^n,\nu^{n,m+1};w^{n,m}) \leq 0.$$

Also, (17) and $G_{c^n}(\underline{W}) \geq 0$ imply $\underline{W} \leq w^{n,m}$ by Lemma 3.5. The inner sequence therefore has a limit $w^n \in \mathcal{W}$. To identify this limit, note that $w^{n,m}$ lie in the compact order interval $[\underline{W}, w^{n,0}] \subset \mathcal{W}$. We have $\nu^{n,m+1} = \nu^*(c^n, w^{n,m})$, and continuity of ν^* implies that the associated matrices $Q(c^n,\nu^{n,m+1})$ are bounded. Since

$$G_{c^n}(w^{n,m}) = Q(c^n,\nu^{n,m+1})\left(w^{n,m+1} - w^{n,m}\right),$$

the right-hand side converges to zero. Continuity gives $G_{c^n}(w^n) = 0$, and Lemma 3.5 gives uniqueness in $\mathcal{W}$. Finally,

$$G_{c^{n+1}}(w^n) = B(w^n) \geq G_{c^n}(w^n) = 0 = G_{c^{n+1}}(w^{n+1}).$$

Fixed-policy comparison gives $w^n \leq w^{n+1}$, while Bellman comparison with $\overline{w}$ gives $w^{n+1} \leq \overline{w}$, which completes the proof. ■

We can now state the convergence result.

Theorem 3.8 (global convergence and uniqueness)**.** *Under the assumptions and initialization of Proposition 3.7, the exact iterates of Algorithm 3.6 converge componentwise to a limit $w^* \in \mathcal{W}$. The limit is the unique solution in $\mathcal{W}$ of (11).*

Proof. Proposition 3.7 gives $w^n \leq w^{n+1} \leq \overline{w}$, and so w^n converges to some w^* in the order interval $K = [w^0, \overline{w}] \subset \mathcal{W}$. It remains to show that $B(w^*) = 0$. In Case 1, (16) and policy evaluation imply

$$B(w^n) = Q(c^{n+1}, v^{n+1})(w^{n+1} - w^n).$$

The sequence of matrices on the right is bounded: c^{n+1} lies in the compact set $\mathcal{C}$, w^n lies in K, and $v^{n+1} = v^*(c^{n+1}, w^n)$ is continuous. Since $w^{n+1} - w^n \to 0$, it follows that $B(w^n) \to 0$.

In Case 2 the definitions of c^{n+1} and w^n imply $B(w^n) = G_{c^{n+1}}(w^n)$ and $G_{c^{n+1}}(w^{n+1}) = 0$. The family $\{G_c : c \in \mathcal{C}\}$ is uniformly Lipschitz on K: the aggregator is continuously differentiable on the compact set $\mathcal{C} \times K$, and $A(c)$ is bounded. Therefore, for some $L < \infty$,

$$\|B(w^n)\| \leq L\|w^{n+1} - w^n\| \longrightarrow 0.$$

The map B is continuous because $\mathcal{C}$ is compact, so $B(w^*) = 0$, and uniqueness follows from Lemma 3.5. ■

When A is conservative, Lemma 3.1 supplies the barriers, and the algorithm can moreover be initialized at the constant discount rate δ.

Corollary 3.9 (canonical initialization)**.** *Let Assumption 3.1 hold and suppose $A(c)\mathbf{1} = 0$ for every $c \in \mathcal{C}$. In Case 1, for any initial consumption policy c^0, Algorithm 3.6 can be initialized at $v^0 = \delta\mathbf{1}$. In Case 2, every inner loop can be initialized at $v^{n,0} = \delta\mathbf{1}$.*

Proof. Assumption 3.1 and the finiteness of M give constants $0 < c_- \leq c_+ < \infty$ that bound every feasible consumption component. Lemma 3.1 therefore supplies the supersolution and uniform policy subsolution. Further, $Q(c, \delta\mathbf{1})\mathbf{1} = \delta\mathbf{1} \gg 0$, and so Lemma 3.4 gives $\delta\mathbf{1} \in \mathcal{N}(c)$ for every c. ■

4 Consumption-savings problem

We now consider a consumption-savings problem of the kind considered in Achdou and Tang (2026b). We first describe the continuous-state HJB equation before specifying how to construct a finite system of equations using finite differences.

4.1 Environment

Income risk and market structure. We assume that log income y_t evolves according to a finite-state income Markov process supported on $\Sigma_y := \{y_0, y_1, \ldots, y_{N_y}\}$ where $N_y \geq 1$ and $y_0 < \cdots < y_{N_y}$, and for $0 \leq j, k \leq N_y$ with $j \neq k$, the rate of transition from y_j to y_k is denoted by $\lambda_{jk} \geq 0$. The agent may save in a risk-free bond with return $r \in (0, \delta)$ and so the law of motion of wealth is given by $da_t = s_t dt$ where $s_t = (ra_t + e^{y_t} - c_t)$. Finally, the agent faces a borrowing constraint of the form $a_t \geq \underline{a}$ that is strictly tighter than the natural borrowing limit,

$$\underline{a} > -e^{y_0}/r.$$

Definition 2. *For $j = 0, \ldots, N_y$ we define the operator*

$$\mathcal{L}(c, v)W(a, y_j) = -vW(a, y_j) + (ra + e^{y_j} - c)\partial_a W(a, y_j) + \sum_{k=0}^{N_y} \lambda_{jk}(W(a, y_k) - W(a, y_j)).$$

The HJB equation for the value function in Case 1 is

$$0 = \max_{c \in \mathcal{C}} \max_{v \in D_v} F(c, v) + \mathcal{L}(c, v)W,$$

with the max over v replaced by a min in Case 2.

4.2 Discretization

We now wish to define a discrete Bellman equation of the form given in (11) for some control set $\mathcal{C}$ and matrix-valued function A. Appendix C relates the solution of such discrete problems to an appropriately defined (constrained viscosity) solution to the HJB equation in Definition 2.

The discretization of the asset space follows the standard "upwinding" procedure widely used in the finite-difference literature. Given $N_a \geq 1$ and asset bounds $\underline{a} < \overline{a}$, we define $\Sigma_a := \{\underline{a}, \underline{a} + \Delta_a, \ldots, \overline{a}\}$ where $\Delta_a := (\overline{a} - \underline{a})/N_a$. We use i to enumerate points in the asset grid and j and k to enumerate points in the income grid, and we write $X := \Sigma_a \times \Sigma_y$ for the state space and $M := (N_y + 1)(N_a + 1)$. We now define our sets of controls $\{C_x\}_{x \in X}$ and the function A.

Definition 3. *Given $c_-, c_+ > 0$ satisfying $c_- < r\underline{a} + e^{y_0}$ and $c_+ > r\overline{a} + e^{y_{N_y}}$, for any state $x = (a, y) \in X$, we define our control sets as follows:*

1. *if $a \neq \underline{a}, \overline{a}$, then $C_x = [c_-, c_+]$;*
2. *if $a = \underline{a}$, $C_x = [c_-, \min\{c_+, ra + e^y\}]$;*

3. *if* $a = \bar{a}$, $C_x = [\max\{c_-, ra + e^y\}, c_+]$.

For any $c \in \mathcal{C} := \prod_{x \in X} C_x$, *define*

$$[A(c)W]_{ij} = (ra_i + e^{y_j} - c_{ij})^+ W_{ij}^F - (ra_i + e^{y_j} - c_{ij})^- W_{ij}^B + \sum_{k=0}^{N_y} \lambda_{jk}(W_{ik} - W_{ij})$$

where the finite differences are given by

$$W_{ij}^F = \frac{1}{\Delta_a}(W_{i+1,j} - W_{ij}) \qquad W_{ij}^B = \frac{1}{\Delta_a}(W_{ij} - W_{i-1,j}).$$

The following verifies that the framework of Section 3 applies to this problem.

Lemma 4.1. *The control set* $\mathcal{C}$ *and function* A *given in Definition 3 satisfy Assumption 3.1, and the vectors* $\overline{w}$ *and* $\underline{w}$ *in Lemma 3.1 satisfy Assumptions 3.2 and 3.3, respectively.*

Proof. The function A satisfies Assumption 3.1 by inspection, because the definitions of C_x ensure that $A(c)$ never places a nonzero coefficient on points outside of the grid. This also implies that $A(c)$ is conservative, and so the remaining claims follow from Lemma 3.1. ■

Policy update. For $W \in \mathcal{W}$ and $c \in \mathcal{C}$, the optimal ν is

$$\nu^*(c, W) = \delta\Big(\theta + (1-\theta)c^{1-\rho}[(1-\gamma)W]^{-1/\theta}\Big)$$

regardless of whether or not the optimization over ν is a maximum or a minimum. The consumption update must incorporate the dependence of the optimal ν on c and W. Omitting terms that do not depend on c gives

$$\mathcal{O}(c) = \frac{\delta c^{1-\rho}}{1-\rho}\big((1-\gamma)W_{ij}\big)^{1-1/\theta} + (ra_i + e^{y_j} - c)^+ W_{ij}^F - (ra_i + e^{y_j} - c)^- W_{ij}^B.$$

To maximize this objective over the set $\mathcal{C}$ given in Definition 3, we consider five candidate maxima: the three points $c_-, ra + e^y$ and c_+, and the solutions to the first-order conditions

$$c_{Id} = \Big([(1-\gamma)W_{ij}]^{1/\theta - 1} W_{ij}^d / \delta\Big)^{-1/\rho}$$

for each direction $d \in \{F, B\}$.

Tolerance. Given $\gamma > 0$ and $W \in \mathcal{W}$, we define consumption-equivalent utility by

$$CE(W) = u^{-1}(W) = [(1-\gamma)W]^{\frac{1}{1-\gamma}}$$

and in our implementation of Algorithm 3.6 we will use the metric

$$d(W, V) := ||u^{-1}(W) - u^{-1}(V)||_\infty.$$

The consumption-equivalent transformation serves three purposes. First, it expresses values in units of consumption: by the identity $f(z, u(z)) = 0$, we have $u^{-1}(W) = z$ exactly when W is the utility of the constant stream z, so tolerances on CE bound errors in units of permanent consumption. Second, CE is a strictly increasing transformation on $\mathcal{W}$, so the monotonicity of the iterates established in Proposition 3.7 is preserved. Third, utility is neither comparable across preference parameters nor even of common sign on the two sides of $\gamma = 1$, whereas consumption equivalents are directly comparable, permitting common tolerances and error metrics across all of the experiments below.

4.3 Numerical illustration

We now illustrate the results above numerically. For the following exercises, we fix parameters governing income and the interest rate and vary preferences.

Income process. We assume that y follows a finite-state approximation to an Ornstein-Uhlenbeck process. Given lower and upper bounds $\underline{y} < 0 < \overline{y}$, we define an increment $\Delta_y := (\overline{y} - \underline{y})/N_y$ for log income and consider an equispaced grid $\Sigma_y := \{\underline{y}, \underline{y} + \Delta_y, \ldots, \overline{y}\}$. For constants $\overline{\mu}, \overline{\sigma} > 0$, we then define the income transition rates as follows. For non-boundary points $j \neq 0, N_y$,

$$\lambda_{j,j\pm1} = \frac{1}{\Delta_y^2}\left(\frac{\overline{\sigma}^2}{2} + \Delta_y[-\overline{\mu} y_j]^\pm\right)$$

while for the boundary points we keep only inward transitions (which amounts to setting volatility to zero on the boundary)

$$\lambda_{01} = \frac{1}{\Delta_y}[-\overline{\mu} y_0]^+ \qquad\qquad \lambda_{N_y,N_y-1} = \frac{1}{\Delta_y}[-\overline{\mu} y_{N_y}]^-$$

with all other transitions zero. We follow the appendix of Achdou et al. (2022) and choose $\overline{\mu} = -\ln \rho_y, \rho_y = 0.95$ and $\overline{\sigma} = \overline{\sigma}_\kappa \sqrt{2\overline{\mu}}, \overline{\sigma}_\kappa = 0.2$, where ρ_y corresponds to the annual autocorrelation and $\overline{\sigma}_\kappa$ is the standard deviation of the resulting stationary distribution.

Initial policies. In the following numerical experiments, the initial consumption policy is *zero net saving*, $c^0 := ra + e^y$, and the initial discount policy is the CRRA value, $\nu^0 = \delta$.

Example. We first plot several examples before turning to a more systematic exploration. Because the arguments employed in this paper differ depending on the signs of $\gamma - \rho$ and $1 - \gamma$, we will consider four examples for illustration. We fix the grid at the

moderate size $N = (200, 20)$ and focus on the savings functions. Elementary calculations reviewed in Lemma B.1 show that in the case of deterministic income and no borrowing constraints, for all y we have

$$\lim_{a \to \infty} \frac{s(a, y)}{a} = \frac{r - \delta}{\rho}.$$

Motivated by this, we choose $r = 0.036 < \delta = 0.04$ and check ex post that dissaving occurs at sufficiently large assets. For our examples we choose $\gamma = 0.75, 1.5$ and $\rho = 0.5, 2.0$, so that all four of the cases considered in Section 3.2 are represented. Figure 1 depicts the associated savings functions.

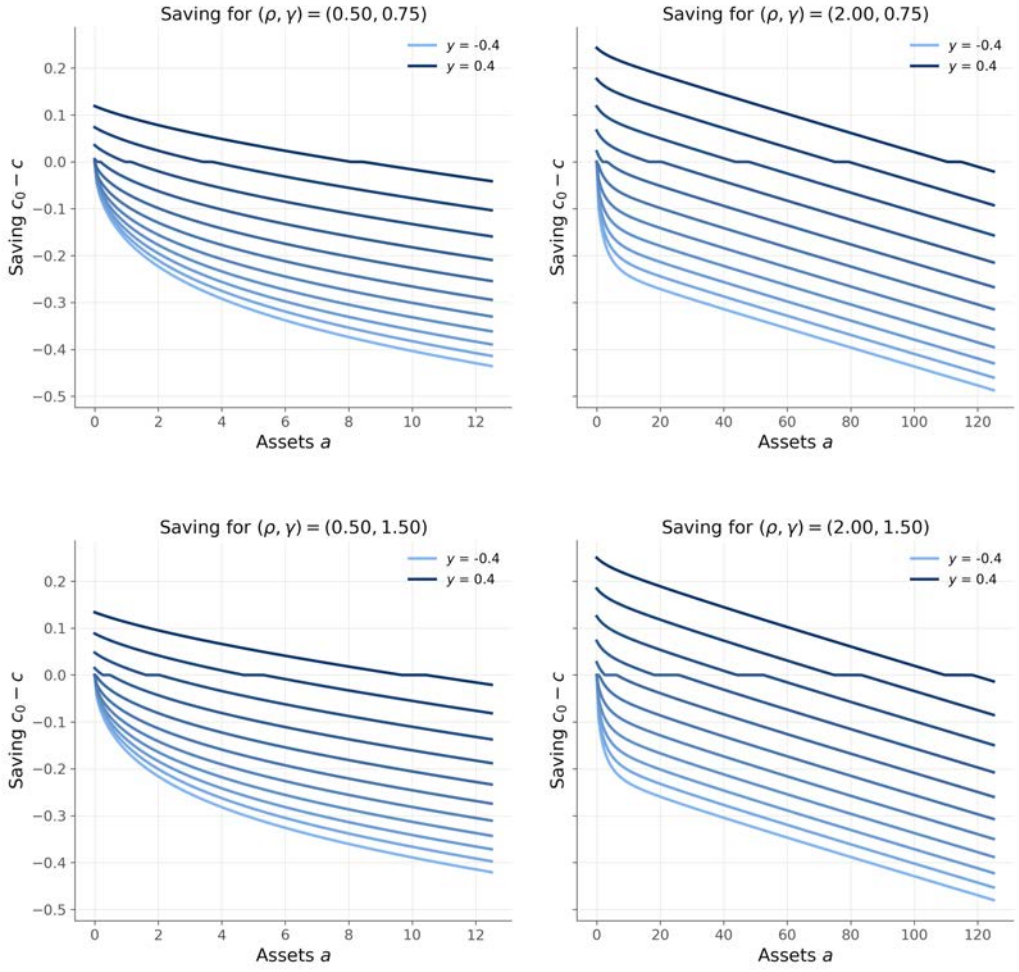


Figure 1: Savings policies in the four preference configurations. *Notes:* Each panel plots the savings function $s(a, y) = ra + e^y - c(a, y)$ against assets for one (ρ, γ) pair, with one line per depicted income state (light to dark from $y = -0.4$ to $y = 0.4$). Parameters: $r = 0.036$, $\delta = 0.04$, $\underline{a} = 0$, income process as described in the text, grid $N = (200, 20)$; the upper bound $\bar{a}$ is chosen so that dissaving occurs there, and is an order of magnitude larger for $\rho = 2$ than for $\rho = 0.5$. Flat segments occur where savings change sign and the scheme switches between forward and backward differences.

Note that in Figure 1 we observe a "kink" at the points in which savings are zero. This feature is not specific to Epstein–Zin preferences, but instead occurs because these are the asset levels where we switch from forward to backward differences. The size of this kink appears to be increasing in γ, which is intuitive because γ governs the curvature of the value function and hence the discrepancy between candidate first derivatives. Further, as expected, when ρ is higher (so that EIS is lower), the savings lines are flatter and we must increase $\bar{a}$ in order to ensure dissaving. Finally, for all four examples in Figure 1, it appears that for the lower income level, the slope in the savings function diverges to $-\infty$ as we approach the lowest asset level; we return to this behavior in Section 4.4.

Monotonicity and speed. Proposition 3.7 shows that any sequence $(w^n)_{n\geq 0}$ produced by Algorithm 3.6 is monotonically increasing. Figure 2 illustrates both this property and the speed of convergence: for each outer iteration n it plots $-\log_{10}\left(1 - u^{-1}(w^n)/u^{-1}(w^*)\right)$ against assets at a fixed income state, where w^* denotes the converged solution. Because CE is increasing and $w^n \uparrow w^*$, the argument of the logarithm is a positive and shrinking relative shortfall, and the plotted quantity can be read as the number of accurate decimal digits of the consumption-equivalent value at each state. Monotonicity appears as the strict vertical ordering of the curves: successive iterates shift upward and never cross. Convergence is also fast. The first iterate—the CRRA evaluation of the initial policy—is already accurate to roughly two digits; each subsequent outer iteration adds roughly one and a half to two digits, nearly uniformly across the state space; and fewer than ten outer iterations drive the relative error below 10^{-10} in all four preference configurations (the algorithm terminates once successive consumption-equivalent values agree to within the solver tolerance of 10^{-8}). Accuracy is lowest near the borrowing constraint, where policies vary most rapidly.

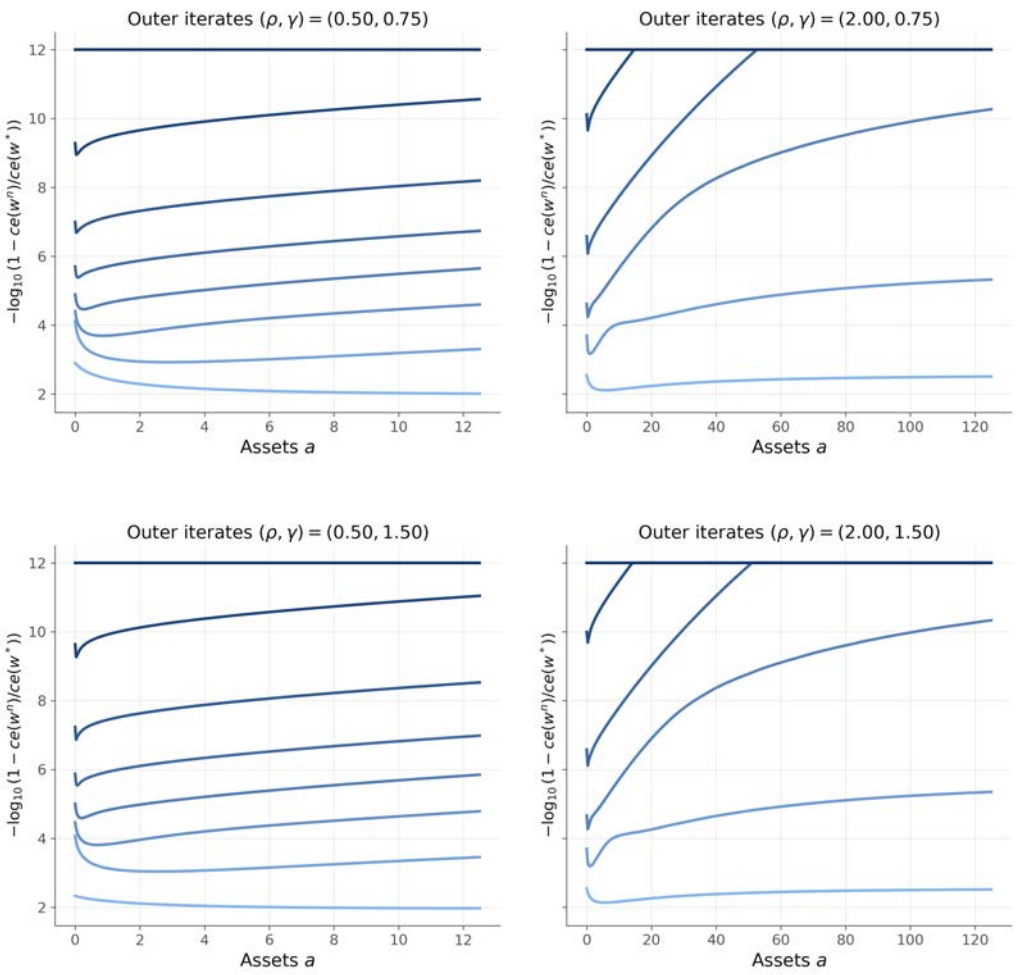


Figure 2: Monotone convergence of the outer iterates. *Notes:* Each panel plots $-\log_{10}(1 - u^{-1}(w^n)/u^{-1}(w^*))$ against assets for successive outer iterations n (light to dark) at a fixed income state, where w^* is the converged solution; the plotted value can be read as the number of accurate decimal digits of the consumption-equivalent value. Curves shift upward and never cross, illustrating Proposition 3.7; the topmost flat line is the converged solution itself, shown at the 10^{-12} display floor. Curves correspond to the middle income state $y = 0$; grid $N = (400, 20)$, remaining parameters as in Figure 1.

Accuracy. We now explore the accuracy of our algorithm. As discussed in Achdou et al. (2022), for continuous-time problems the "Euler equation errors" familiar from the discrete-time literature are uninformative, because the first-order condition for consumption holds by construction up to solver tolerance. We therefore fix the income grid, compute policy and value functions on a fine asset grid with $N_{\text{fine}} = 20{,}480$ nodes, and compare solutions computed on coarser grids $N_a = 40, 80, \ldots, 5{,}120$ against this reference. Figure 3 plots the sup-norm gaps of the consumption policy (left panel) and of the value

function in consumption-equivalent units (right panel) on a log-log scale. Three points stand out. First, the value error declines linearly in the logarithm of the grid size with slope close to -1 for all four preference configurations: halving the grid spacing halves the error, the first-order rate expected of an upwind scheme. Second, values converge more quickly and far more smoothly than policies: the policy error decays at a slower and noisier rate, reflecting the kinks at zero-savings points visible in Figure 1. Third, accuracy is respectable already on coarse grids: at $N_a \approx 200$, the grid used in Figure 1, the consumption-equivalent value error lies between roughly 3×10^{-4} and 4×10^{-3}.

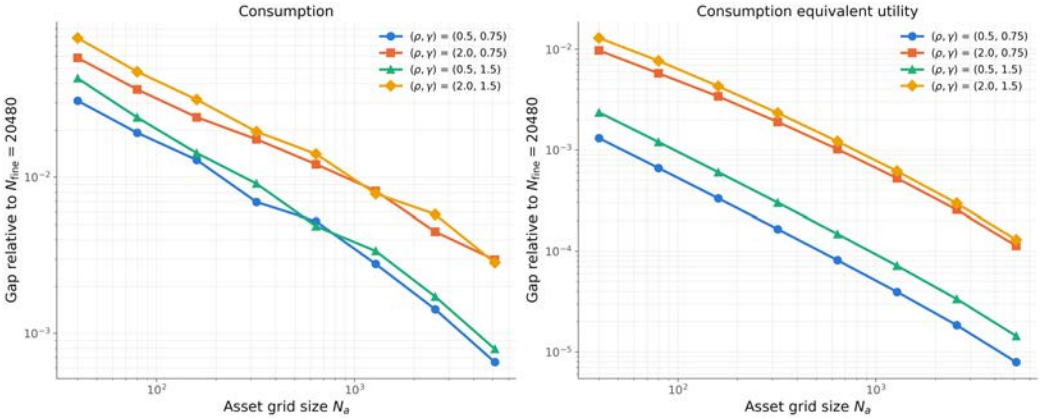


Figure 3: Grid refinement. *Notes:* Sup-norm distance between solutions computed on an asset grid with N_a nodes and a reference solution with $N_{\text{fine}} = 20{,}480$ nodes, for the four preference configurations of Figure 1. Left panel: consumption policy. Right panel: value function in consumption-equivalent units. Log-log scale; a slope of -1 corresponds to first-order convergence.

4.4 Savings near the borrowing constraint

We next study the savings behavior of low-income agents near the borrowing constraint. For the CRRA case with two income types, Achdou et al. (2022) show that the savings function obeys a *square-root law*: savings vanish like $\sqrt{a - \underline{a}}$ as assets approach the constraint, with a coefficient available in closed form. We now show that this law extends to Epstein–Zin preferences, derive the coefficient, and verify the prediction numerically.

It is convenient to work with one income state at a time. For $j = 0, \ldots, N_y$ we write $W_j(a) := W(a, y_j)$ and $c_j(a) := c(a, y_j)$ for the value and consumption functions in state j, and $s_j(a) := ra + e^{y_j} - c_j(a)$ for the savings function. Using Lemma 3.2, the HJB equation in Definition 2 simplifies to

$$0 = \sup_{c>0} f(c, W_j(a)) + (ra + e^{y_j} - c)W_j'(a) + \sum_{k=0}^{N_y} \lambda_{jk}\big(W_k(a) - W_j(a)\big), \tag{18}$$

for $j = 0, \dots, N_y$, and using (9), the first-order condition for consumption becomes

$$W_j'(a) = \delta c_j(a)^{-\rho} \big[(1-\gamma) W_j(a)\big]^{\frac{\rho-\gamma}{1-\gamma}}. \tag{19}$$

We now state the following regularity assumption that will allow us to sharply characterize the savings behavior near the borrowing constraint.

Assumption 4.1. *On $(\underline{a}, \infty)$, the value functions are C^2 and the policies c_j are C^1, and the limits $\underline{c}_j := \lim_{a \downarrow \underline{a}} c_j(a)$ and $\underline{W}_j := \lim_{a \downarrow \underline{a}} W_j(a)$ exist and satisfy $\underline{c}_j, (1-\gamma)\underline{W}_j \in (0, \infty)$.*

Proposition 4.2 (square-root law)**.** *Under Assumption 4.1, $\lim_{a \downarrow \underline{a}} -s_0(a)/\sqrt{a - \underline{a}} = \widehat{\kappa}$, where*

$$\widehat{\kappa}^2 = \frac{2\underline{c}_0}{\rho} \left[\underline{\nu}_0^* - r - \sum_{k=0}^{N_y} \lambda_{0k} \left(\frac{W_k'(\underline{a})}{W_0'(\underline{a})} - 1 \right) \right]. \tag{20}$$

Proof. The idea of the proof is to use (18) to derive two expressions for $s_0(a) W_0''(a)/W_0'(a)$, equate these to one another, and use this to obtain an expression for $\lim_{a \downarrow \underline{a}} s_0(a) s_0'(a)$. Differentiating (18) for $j = 0$ with respect to a and applying the envelope theorem gives

$$0 = W_0'(a) f_w(c_0(a), W_0(a)) + r W_0'(a) + s_0(a) W_0''(a) + \sum_{k=0}^{N_y} \lambda_{0k} \big(W_k'(a) - W_0'(a) \big).$$

Writing $f_w(c_0(a), W_0(a)) = -\nu_0^*(a)$ and dividing by $W_0'(a)$ gives

$$\frac{s_0(a) W_0''(a)}{W_0'(a)} = \nu_0^*(a) - r - \sum_{k=0}^{N_y} \lambda_{0k} \left(\frac{W_k'(a)}{W_0'(a)} - 1 \right).$$

Taking the log of both sides of (19) and differentiating with respect to assets gives

$$\frac{s_0(a) W_0''(a)}{W_0'(a)} = -\rho \frac{s_0(a) c_0'(a)}{c_0(a)} + s_0(a) \left(\frac{\rho - \gamma}{1 - \gamma} \right) \frac{W_0'(a)}{W_0(a)}.$$

Combining the two expressions and using $c_0'(a) = r - s_0'(a)$ gives

$$\nu_0^*(a) - r - \sum_{k=0}^{N_y} \lambda_{0k} \left(\frac{W_k'(a)}{W_0'(a)} - 1 \right) = \frac{\rho s_0(a) s_0'(a)}{c_0(a)} + s_0(a) \left(-\frac{\rho r}{c_0(a)} + \left(\frac{\rho - \gamma}{1 - \gamma} \right) \frac{W_0'(a)}{W_0(a)} \right).$$

The term in parentheses on the right is bounded as $a \downarrow \underline{a}$ while $s_0(a) \to 0$, so taking limits and applying l'Hôpital's rule gives $\lim_{a \downarrow \underline{a}} s_0^2/(a - \underline{a}) = \lim_{a \downarrow \underline{a}} 2 s_0 s_0'$ and hence (20), where we take the negative root because the constrained agent dissaves toward $\underline{a}$. ■

Formula (20) has a transparent economic interpretation. The bracketed term is the effective impatience wedge of the constrained agent: the endogenous discount rate $\underline{\nu}_0^*$, net

of the return r on wealth and net of the effect of income transitions, which is captured by the marginal-value ratios $W_k'(\underline{a})/W_0'(\underline{a})$. Because wealth is less valuable in better income states, these ratios lie below one for upward transitions, so the possibility of an income improvement acts like additional impatience and speeds decumulation. The prefactor $2\underline{c}_0/\rho$ scales the wedge by the level of consumption and by the elasticity of intertemporal substitution: the more willing the agent is to substitute consumption over time (the lower is ρ), the faster assets are run down. Notably, risk aversion appears nowhere directly: γ matters only through the continuation objects $\underline{\nu}_0^*$ and W_k'/W_0'.

The square-root law has a mechanical interpretation that we find useful. Consider a car traveling at speed V that brakes at a constant rate C. Its average speed before stopping is $V/2$ and it takes V/C units of time to stop, so it travels a distance $D = V^2/2C$; equivalently, at distance D from its stopping point its speed is $\sqrt{2CD}$. Whenever deceleration is constant—or, as here, tends to a constant at the stopping point—speed is proportional to the square root of the remaining distance. Now let the distance be the agent's assets above the borrowing constraint, the speed the rate of dissaving, and the braking rate the rate at which dissaving slows as the constraint approaches. The proof of Proposition 4.2 establishes that this braking rate has a finite, strictly positive limit at the constraint,

$$C := \lim_{a \downarrow \underline{a}} s_0(a) s_0'(a) = \frac{\widehat{\kappa}^2}{2} = \frac{\underline{c}_0}{\rho}\left[\underline{\nu}_0^* - r - \sum_{k=0}^{N_y} \lambda_{0k}\left(\frac{W_k'(\underline{a})}{W_0'(\underline{a})} - 1\right)\right], \tag{21}$$

where $s_0 s_0'$ is, by the chain rule, the "deceleration" $\dot{s}_t$ of assets along the optimal path. Further, near the borrowing constraint we have $\dot{s}_t = r s_t - \dot{c}_t \approx -\dot{c}_t$, and so the fact that this rate tends to a finite nonzero limit is intuitively a consequence of the Euler equation governing $\dot{c}_t$. This rate is governed by the elasticity of intertemporal substitution, which determines how quickly the agent is willing to let consumption fall as assets run out, and by the impatience wedge, which in our calibration is dominated by the prospect of an income improvement—the force that makes a constrained agent willing to run assets down in the first place.[5] Because the braking rate is bounded below, the constraint is reached in finite time, $T = 2\widehat{\kappa}^{-1}\sqrt{a_0 - \underline{a}}$ from initial assets a_0, whereas linear decay would only approach it asymptotically. Appendix A strips the problem down to a deterministic control problem with a linear flow cost and a convex cost of adjustment, and shows that the same braking argument yields the square-root law there, with the impatience wedge in the role of the flow cost and $\underline{c}_0/\rho$ in the role of the inverse curvature of the adjustment cost.

Figure 4 asks how well the law describes the computed policy. For each of the four preference configurations it overlays the savings function in the lowest income state on

[5] In the calibration of Figure 5, the transition term in the bracket of (21) is larger than the term $\underline{\nu}_0^* - r$ in every example we computed—by an order of magnitude at low ρ—and the latter is negative at high risk aversion.

$-\hat{\kappa}\sqrt{a-\underline{a}}$, with $\hat{\kappa}$ taken from (20). The coefficient is built from consumption, the endogenous discount rate and marginal-value ratios at the constraint. The two agree closely: within 1–2 percent for assets up to 5 percent of average income, and within 8–13 percent out to half of average income, beyond which the higher-order terms neglected in the proposition begin to matter.

Figure 5 provides comparative statics, reporting the value of $\hat{\kappa}$ for $\gamma \in [0.25, 3.25]$ and $\rho \in \{0.4, 0.8, 1.2, 1.6, 2\}$. Three features stand out. First, the coefficients are positive and vary smoothly across the entire parameter range, confirming that the square-root law extends beyond the CRRA case. Second, and consistent with the interpretation of (20), the coefficient is governed almost entirely by intertemporal substitution: moving ρ from 0.4 to 2 lowers $\hat{\kappa}$ by about 14 percent, whereas moving γ across the whole plotted range lowers it by under 3 percent. Third, the estimates pass smoothly through $\gamma = 1$ and $\gamma = \rho$, the boundaries at which the algorithm switches among the four cases of Section 3.

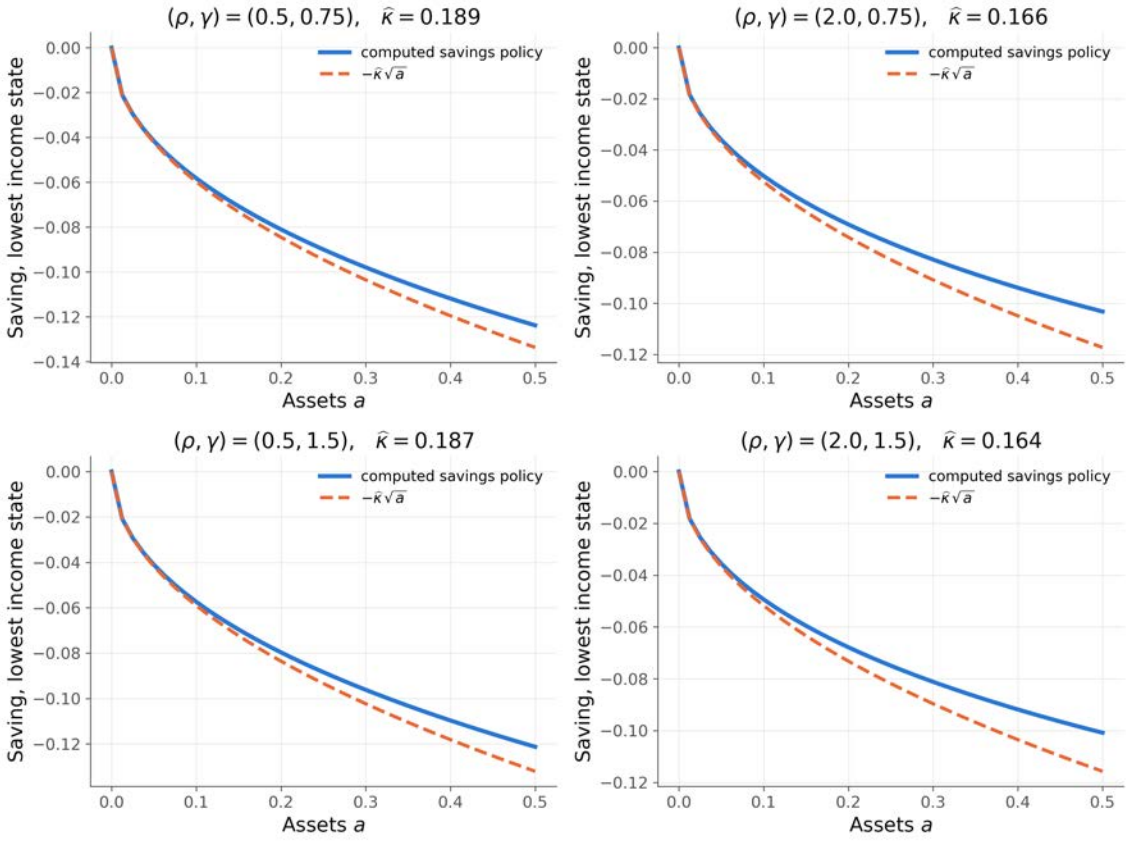


Figure 4: The square-root law near the borrowing constraint $\underline{a} = 0$. *Notes:* Solid lines are the computed savings function in the lowest income state; dashed lines are $-\hat{\kappa}\sqrt{a-\underline{a}}$, with $\hat{\kappa}$ given by the right-hand side of (20). The coefficient is computed independently of the interior savings policy. The horizontal range shown is roughly the lowest three quarters of one year's income. Asset grid $N_a = 16{,}000$ on $[0, 200]$; remaining parameters as in Figure 1.

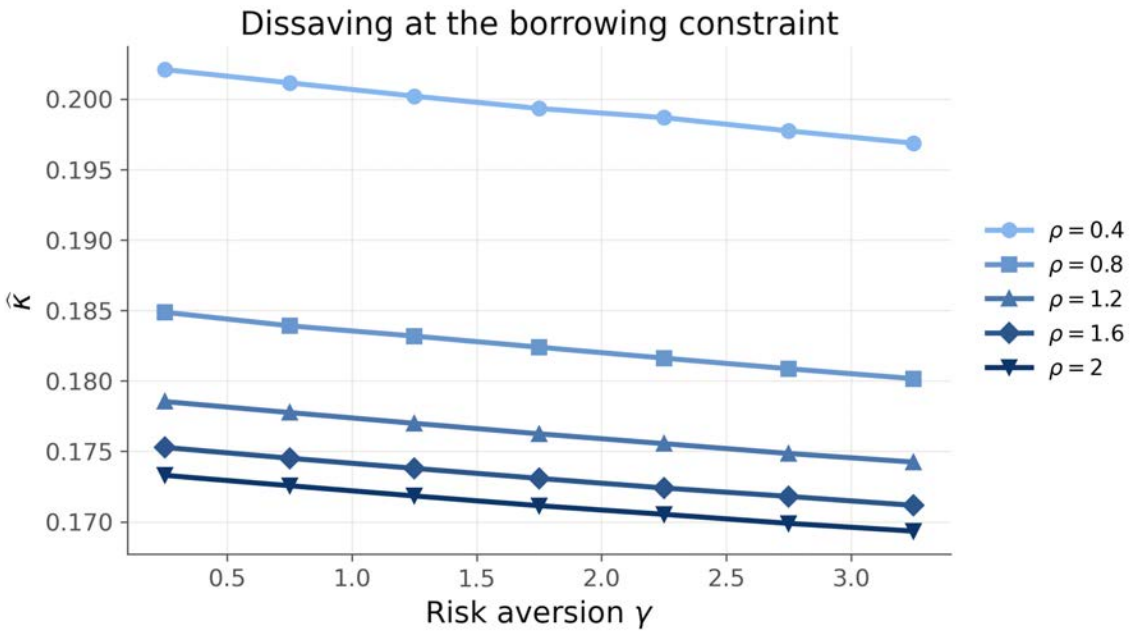


Figure 5: Dissaving at the borrowing constraint. *Notes:* The square-root coefficient $\widehat{\kappa}$ of Proposition 4.2 in the lowest income state, plotted against risk aversion γ for five values of the inverse elasticity of intertemporal substitution ρ. Because $\widehat{\kappa}$ is a limit at a point where the savings function has infinite slope, it is estimated on asset grids of $N_a = 4{,}000$, 8,000 and 16,000 nodes and Richardson-extrapolated; the two estimators of (20) extrapolate to the same value to within 10^{-3} at every point shown. Borrowing constraint $\underline{a} = 0$, asset domain $[0, 200]$; income process and (r, δ) as in Figure 1.

4.5 Permanent income risk

In the experiments of Section 4.3, the matrix $A(c)$ was conservative (zero row sums) and so Corollary 3.9 guaranteed that Algorithm 3.6 converged from the initial guess $\nu^0 = \delta$. To illustrate how the algorithm and certificate argument apply to non-conservative settings, we now add a *permanent* component to income. Despite the additional state variable, the homogeneity of our aggregator ensures that the resulting discretized system reduces to a two-dimensional problem of the form given in Section 3. However, in contrast with Section 4.3, now $A(c)\mathbf{1}$ is a *nonzero* constant that depends on the growth of permanent income, and the choice of initialization becomes more subtle: once that constant exceeds δ, the natural starting values are no longer certified and a naive application of Newton's method may fail. The theory of Section 3 identifies exactly which constant initializations are certified, and supplies a natural one.

Environment. Income is $Y_t = P_t e^{z_t}$, where z_t is the finite-state process of Section 4.3 and P_t is a permanent component, whose logarithm is a pure-jump process with drift μ_p, symmetric jumps of size $\pm\eta$ at rate $\lambda = \sigma_p^2/(2\eta^2)$, and negative jumps of size $-m\eta$ at

rate λ_d for some integer $m \geq 1$.[6] The latter represent idiosyncratic disasters: infrequent, large, negative shocks such as job displacement, disability, or health events that permanently lower earnings.[7] Wealth evolves as $da_t = (ra_t + Y_t - c_t)dt$ subject to $a_t \geq \underline{a}P_t$.[8] Because the aggregator is homogeneous of degree $1-\gamma$ in (c, W), the value function satisfies $W(a, z, P) = P^{1-\gamma} w(\hat{a}, z)$, where $\hat{a} = a/P$ denotes normalized wealth and $w(a, z) := W(a, z, 1)$. The HJB equation for w coincides with that in Definition 2 with three modifications: 1) the drift of normalized wealth is $(r - \mu_p)\hat{a} + e^z - \hat{c}$ where $\hat{c} = c/P$; 2) a jump of $\ln P$ by κ moves $\hat{a}$ to $\hat{a}e^{-\kappa}$ and multiplies continuation utility by $e^{(1-\gamma)\kappa}$; and 3) the drift of $\ln P$ contributes the additional term $(1-\gamma)\mu_p w$. It is instructive to write

$$h := (1-\gamma)\mu_p + \lambda\left(e^{(1-\gamma)\eta} - 1\right) + \lambda\left(e^{-(1-\gamma)\eta} - 1\right) + \lambda_d\left(e^{-(1-\gamma)m\eta} - 1\right)$$

for the expected growth rate of $P_t^{1-\gamma}$, that is, $\mathbb{E}_t[dP_t^{1-\gamma}] = hP_t^{1-\gamma}dt$. Since $1-\gamma < 0$ in our calibration, permanent risk and disasters raise h and growth lowers it.

Discretization. We use the upwind scheme of Definition 3 on the grid with ith element $\hat{a}_i = \underline{a}e^{i\eta}$, and we declare that any jumps that would leave the grid are absorbed at its ends. The resulting $A(c)$ satisfies Assumption 3.1 and has constant row sums, $A(c)\mathbf{1} = h\mathbf{1}$ for every $c \in \mathcal{C}$: the upwind drift terms and the transitions of z have zero row sums, and each jump κ of $\ln P$ contributes its rate times $e^{(1-\gamma)\kappa} - 1$. However, for $h \neq 0$, Lemma 3.1 and Corollary 3.9 do not apply. The following two results replace them.

Lemma 4.3 (barriers with constant row sums)**.** *Suppose that Assumption 3.1 holds, that for some $h \in \mathbb{R}$ we have $A(c)\mathbf{1} = h\mathbf{1}$ for every $c \in \mathcal{C}$, and that $0 < c_- \leq c_x \leq c_+ < \infty$ for every $c \in \mathcal{C}$ and state x. If $\rho \neq 1$ and $1 - h/(\delta\theta) > 0$, then*

$$\underline{w} := u(c_-)\left(1 - \frac{h}{\delta\theta}\right)^{-\theta}\mathbf{1} \qquad \text{and} \qquad \overline{w} := u(c_+)\left(1 - \frac{h}{\delta\theta}\right)^{-\theta}\mathbf{1}$$

satisfy Assumptions 3.3 and 3.2, respectively. If $\rho = 1$, the same holds with the factor $e^{h/\delta}$ in place of $(1 - h/(\delta\theta))^{-\theta}$.

Proof. Let $\rho \neq 1$. For any $m > 0$, the vector $u(m)\mathbf{1}$ lies in $\mathcal{W}$, $A(c)u(m)\mathbf{1} = hu(m)\mathbf{1}$, and

[6] The size of the jump is chosen so that in the below discretization, all transitions remain on the grid. Note that for the above specification, as $\eta \to 0$ the process $\ln P_t$ approaches a diffusion process with variance σ_p^2.

[7] Long-run earnings losses of 20–30 percent after job displacement are documented by Jacobson et al. (1993) and Davis and von Wachter (2011), with annual displacement rates of 2-4 percent for established workers; Guvenen et al. (2021) show that infrequent large shocks are a dominant feature of individual earnings risk; Low et al. (2010) treat disability. The range $\sigma_p \in [0.10, 0.15]$ is that of Storesletten et al. (2004).

[8] In our numerical examples we take $\underline{a} = 0.01$ so that the logarithmic grid below has a finite lower bound. This is a slight tightening of the constraint $a_t \geq 0$ and is immaterial for the comparison of algorithms.

substituting $(1-\gamma)u(m) = m^{1-\gamma}$ into (9) gives

$$G_c(u(m)\mathbf{1})_x = f(c_x, u(m)) + hu(m) = m^{1-\gamma}\varphi(c_x/m), \qquad \varphi(\xi) := \delta\,\frac{\xi^{1-\rho}-1}{1-\rho} + \frac{h}{1-\gamma}.$$

The function φ is strictly increasing on $(0,\infty)$ and vanishes at $\bar{\xi} := (1-h/(\delta\theta))^{1/(1-\rho)}$. Taking $m = c_+/\bar{\xi}$ gives $c_x/m \le \bar{\xi}$ and hence $G_c(\overline{w})_x \le 0$ for every $c \in \mathcal{C}$ and x, while taking $m = c_-/\bar{\xi}$ gives $G_c(\underline{w})_x \ge 0$. Finally, $u(c_\pm/\bar{\xi}) = u(c_\pm)\bar{\xi}^{-(1-\gamma)} = u(c_\pm)(1-h/(\delta\theta))^{-\theta}$. The case $\rho = 1$ is identical with $\varphi(\xi) = \delta\ln\xi + h/(1-\gamma)$ and $\bar{\xi} = e^{-h/(\delta(1-\gamma))}$. ■

The barriers are the values of consuming c_+P_t and c_-P_t forever, and $h = 0$ recovers Lemma 3.1. The condition $1 - h/(\delta\theta) > 0$ is the condition for these values to be finite.[9]

Lemma 4.4 (constant initializations)**.** *Suppose that $A(c)\mathbf{1} = h\mathbf{1}$ for every $c \in \mathcal{C}$. Then for every $c \in \mathcal{C}$ and $\nu \in D_\nu$, the matrix $Q(c, \nu\mathbf{1})$ is monotone if and only if $\nu > h$.*

Proof. If $\nu > h$, then $Q(c,\nu\mathbf{1})\mathbf{1} = (\nu - h)\mathbf{1} \gg 0$ and $\mathbf{1}$ is a certificate. Conversely, $A(c)$ has nonnegative off-diagonal entries and the strictly positive eigenvector $\mathbf{1}$ with eigenvalue h, so by the Perron–Frobenius theorem h is its largest real eigenvalue and admits a left eigenvector $y \ge 0$, $y \ne 0$, with $y^\top A(c) = h\,y^\top$. If $\nu \le h$ and $z \gg 0$ satisfied $Q(c,\nu\mathbf{1})z \gg 0$, then $0 < y^\top Q(c,\nu\mathbf{1})z = (\nu - h)\,y^\top z \le 0$, a contradiction. By Lemma 3.3, this implies that $Q(c,\nu\mathbf{1})$ is not monotone. ■

For constant initializations, part 1 of Lemma 3.4 is therefore not only sufficient but necessary, and whether a constant ν^0 is certified does not depend on the consumption policy. We set $\gamma = 5$ and $\rho = 2$, so that $\theta = 4$ and $D_\nu = (-\infty, \delta\theta)$. The certified constant initializations are then exactly the "window" $(h, \delta\theta)$, which is nonempty precisely when the barriers of Lemma 4.3 exist, and Theorem 3.8 applies from any of them. A natural choice is the discount rate associated with consuming permanent income forever, which we denote by ν_E (for "endowment"),

$$\nu_E := \nu^*\big(1, w_E\big) = h + \frac{\delta\theta - h}{\theta}, \tag{22}$$

where w_E is the value of consuming permanent income forever, given in the preceding footnote. Since $\theta > 1$, ν_E lies strictly inside the window whenever utility is finite, and it reduces to δ when $h = 0$. Two features of the window are worth stressing. First, the CRRA initialization $\nu^0 = \delta$ is certified if and only if $h < \delta$, and no choice of consumption policy changes this. Second, because D_ν is bounded above when $\theta > 1$, the window is bounded

[9] If consumption equals permanent income, $c_t = P_t$—an endowment economy in the terminology of asset pricing—the normalized Bellman equation $0 = f(1,w) + hw$ has the solution $(1-\gamma)w = (1-h/(\delta\theta))^{-\theta}$, which exists if and only if $1 - h/(\delta\theta) > 0$.

and shrinks to a point as $h \uparrow \delta\theta$, and so starting from a large discount rate is not a safe rule.

Starting values. All experiments use the initial consumption policy of zero net saving, $c^0 = (r - \mu_p)\hat{a} + e^z$, and one of three starting values:

(C) the *certified* value w_{C}^0, which solves $Q(c^0, \nu_E \mathbf{1})w = F(c^0, \nu_E \mathbf{1})$. By Lemma 4.4, $\nu_E \mathbf{1} \in \mathcal{N}(c^0)$, so $w_{\mathrm{C}}^0 \in \mathcal{W}$ and $B(w_{\mathrm{C}}^0) \geq 0$;

(R) the *CRRA* value w_{R}^0, which solves $Q(c^0, \delta \mathbf{1})w = F(c^0, \delta \mathbf{1}) = \delta u(c^0)$. By Lemma 4.4, the matrix is monotone if and only if $h < \delta$; otherwise w_{R}^0 need not lie in $\mathcal{W}$;

(U) the value $w_{\mathrm{U}}^0 = u(c^0)$, which lies in $\mathcal{W}$ but is not the value of any policy when $h \neq 0$, and $B(w_{\mathrm{U}}^0)$ has no definite sign.

Only (C) uses the theory of Section 3. Value (R) is what a user of the CRRA code would do, and (U) is an obvious choice for a user who starts from a value rather than a policy.

Iterations. From each starting value we run two iterations. Each step of either one takes the current value w^k, updates the policy to (c^{k+1}, ν^{k+1}) with $c^{k+1} \in \arg\max_c G_c(w^k)$ and $\nu^{k+1} = \nu^*(c^{k+1}, w^k)$, and then solves one sparse linear system:

$$\text{Newton on } B: \qquad Q(c^{k+1}, \nu^{k+1})\, w^{k+1} = F(c^{k+1}, \nu^{k+1}), \tag{23}$$

$$\text{implicit, step } \Delta: \qquad \left(\tfrac{1}{\Delta} I + Q(c^{k+1}, \nu^{k+1})\right) w^{k+1} = F(c^{k+1}, \nu^{k+1}) + \tfrac{1}{\Delta} w^k. \tag{24}$$

Iteration (23) is Algorithm 3.6 in Case 1; started from (C) it is exactly our algorithm, and started from (R) or (U) it is the same iteration without the certificate. Iteration (24) is an extension of the implicit scheme of Achdou et al. (2022) to the setting with non-CRRA utility.[10] Rearranged, it reads $w^{k+1} = w^k + (\tfrac{1}{\Delta} I + Q)^{-1} B(w^k)$, a damped Newton step that reduces to (23) as $\Delta \to \infty$.[11] We stop when successive values differ by less than 10^{-8} in consumption-equivalent units, and we record a failure when a starting value or an iterate lies outside $\mathcal{W}$, at which point the next policy update is undefined.

Results. Table 1 reports the number of iterations after the starting value as σ_p increases, with $\mu_p = 0.02$, $\eta = 0.06$, $\lambda_d = 0.02$ and $m = 6$—a 30 percent permanent earnings loss at a rate of 2 percent per year—and the remaining parameters as in Section 4.3.

[10] It is one backward-Euler step of length Δ of the time-dependent equation $\partial_s w = F + Tw$, with the policy frozen at its value at the start of the step.

[11] Near the solution the error contracts by a factor of approximately $1/(1 + \Delta\lambda_{\min})$ per step, where $\lambda_{\min}$ is the smallest eigenvalue of $Q(c^*, \nu^*)$, so convergence is linear for every finite Δ.

Panel A: Newton on B, iteration (23)

					starting value		
σ_p	h	ν_E	$\lambda_{\min}$	$\nu^* < h$	(C) certified	(R) CRRA	(U) $u(c^0)$
0.00	-0.016	0.028	0.045	0%	10	12	20*
0.05	0.005	0.044	0.042	0%	7	6	19*
0.08	0.036	0.067	0.034	80%	6	7	12*
0.10	0.065	0.089	0.026	90%	6	invalid	fails
0.12	0.100	0.115	0.016	92%	7	invalid	fails
0.14	0.142	0.147	0.0046	93%	7	invalid	fails
0.145	0.153	0.155	0.0017	93%	7	invalid	fails

Panel B: implicit time stepping, iteration (24)

	from (U) $u(c^0)$				from (C) certified		
σ_p	$\Delta = 1$	$\Delta = 10$	$\Delta = 100$	largest Δ	$\Delta = 1$	$\Delta = 10$	$\Delta = 100$
0.00	357*	50*	19*	20* (10^3)	327	46	12
0.05	389*	53*	16*	18* (10^3)	377	52	13
0.08	456*	61*	15*	13* (10^3)	468	62	15
0.10	545*	72*	fails	42* (20)	604	78	17
0.12	816*	101*	fails	101* (10)	940	118	22
0.14	2,601*	fails	fails	594* (5)	2,756	334	47
0.145	6,127*	fails	fails	1,422* (5)	6,445	791	102

Table 1: Permanent income risk: iterations to convergence by starting value. *Notes:* Entries are iterations after the starting value, each consisting of one policy update and one sparse linear solve. *: values do not increase monotonically. "Fails": the first iterate lies outside $\mathcal{W}$. "Invalid": the starting value itself lies outside $\mathcal{W}$. In Panel B, the largest step that converges is searched over $\Delta \in \{1, 2, 5, 10, 20, 50, 100, 200, 500, 1000\}$ and reported in parentheses. The certified window for constant initializations is $(h, \delta\theta) = (h, 0.16)$, and ν_E is given by (22). $\lambda_{\min}$ is the smallest eigenvalue of $Q(c^*, \nu^*)$ and the column $\nu^* < h$ is the share of grid points at which the discount rate at the solution lies below h. At $\sigma_p = 0.15$, $h > \delta\theta$ and utility is infinite. Parameters: $\gamma = 5$, $\rho = 2$, $\mu_p = 0.02$, $\eta = 0.06$, $\lambda_d = 0.02$, $m = 6$; grid $\hat{a} \in [0.01, 50]$ with 143 logarithmically spaced nodes and $N_y = 10$; remaining parameters as in Figure 1.

Five features stand out. First, our algorithm—Newton on B from the certified value—converges in six to ten iterations, with values increasing monotonically, at every σ_p up to the finiteness boundary, as Theorem 3.8 guarantees given Lemmas 4.3 and 4.4. Second, the CRRA starting value becomes invalid exactly when $h > \delta$, as Lemma 4.4 predicts: $Q(c^0, \delta\mathbf{1})$ is not monotone, and the first policy evaluation returns a value of the wrong sign. The cause is the constant discount rate, not the consumption policy: by Lemma 4.4, no consumption policy would make $\nu^0 = \delta$ admissible. Newton on B from $u(c^0)$ fails at its first step in the same region, and where it does converge its values are not monotone. Third, time stepping from $u(c^0)$ converges only for small steps: the largest step that works

falls from 10^3 to 5 as σ_p rises, the number of steps rises to 1,422 at that step and to 6,127 at a unit step, and the values are never monotone. Fourth, time stepping from the certified value converges at every step size, with monotone values, and approaches our algorithm as Δ grows; a finite step only slows it down. The difficulties in the other columns are therefore difficulties of the starting value, not of Newton's method or of time stepping as such, and the certificate is what supplies a starting value from which both work.

Fifth, there is no tension between Lemma 4.4, which says that a constant discount rate must exceed h, and the fact that at the solution ν^* lies below h in 80-93 percent of states once $\sigma_p \geq 0.08$. The solution's discount rates are not constant, and the certificate for $Q(c^*, \nu^*)$ is not $\mathbf{1}$ but instead $z = (1-\gamma)w^*$, because $Q(c^*, \nu^*)z = (1-\gamma)F(c^*, \nu^*) \gg 0$, that is, $\nu^*_x > (A(c^*)z)_x / z_x$ in every state. The relevant comparison is therefore with a value-weighted growth rate that varies across states, not with h; numerically the margin is positive everywhere, though as small as 0.0006 near the boundary. This is the self-certification of part 2 of Lemma 3.4 at work, and it is what carries the algorithm from a constant initialization above h to a solution whose discount rates are mostly below it. Finally, $\lambda_{\min}$ is within 10 percent of $(\delta\theta - h)/\theta$, its value when consumption equals permanent income, so the linear rate of time stepping, and its collapse near the boundary, can be read off from the primitives.

For practitioners, the lesson from the theory is simple: start from a constant discount rate in the certified window $(h, \delta\theta)$, for instance ν_E. The CRRA default $\nu^0 = \delta$ is certified only when $h < \delta$, which in our calibration fails once the volatility of permanent shocks exceeds about 8 percent (the finiteness boundary is at about 15 percent), within the range estimated in the literature. A large discount rate is not a safe alternative when $\theta > 1$, because the window is bounded above and narrows to zero as utility approaches infinity; at $\sigma_p = 0.145$ it has width 0.007. Our experiments cover one parameter cell (Case 1 with $\theta > 1$) and a specification with constant row sums; when $A(c)\mathbf{1}$ varies with the state, the same reasoning applies with h replaced by the largest eigenvalue of $A(c^0)$, and a constant ν^0 above the largest row sum of $A(c^0)$ is always certified.

5 Conclusion

This paper combines the variational representation of recursive utility of Geoffard (1996) with a "nested" extension of Howard iteration of Bokanowski et al. (2009) to solve finite-dimensional Bellman equations with continuous-time Epstein–Zin preferences. Our main contribution is a certification argument that establishes convergence without requiring every feasible policy to generate a monotone policy-evaluation matrix. Under our stated conditions, current values and barrier functions certify the matrices generated by the algo-

rithm, and the value iterates increase monotonically toward the unique admissible solution regardless of the sign of the subsequent auxiliary discount rates.

The numerical experiments illustrate the practical performance of the method across the four preference configurations considered in the analysis, with high accuracy relative to the discrete solution attained within roughly ten outer iterations. The consumption-savings application also characterizes the square-root behavior of dissaving near the borrowing constraint. The expression for its coefficient relates intertemporal substitution to endogenous discounting and marginal continuation values, and provides a check on the computed policies. Over the parameter ranges examined, the coefficient varies substantially more with intertemporal substitution than with risk aversion.

The extension to permanent income risk shows where the certificate argument matters in practice. When income has a permanent component, the policy-evaluation matrices are no longer conservative and the choice of starting policy is no longer innocuous. The theory delivers the interval of certified constant initial discount rates, barriers whenever utility is finite, and a natural starting value, and from that value the cost of the algorithm does not rise as the problem approaches the point at which utility becomes infinite. At the solution, the auxiliary discount rate lies below the lower end of that interval in most states, so that it is the current value (up to a sign change), rather than any uniform bound on discount rates, that certifies each step.

References

Achdou, Yves and Qing Tang. 2026a. Continuous-Time Heterogeneous Agent Models with Recursive Utility and Preference for Late Resolution. arXiv:2603.07782.

———. 2026b. Finite Difference Methods for a Continuous-Time Heterogeneous Agent Model with Recursive Utility. arXiv:2606.23408.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2022. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies* 89 (1):45–86.

Barles, G. and P. E. Souganidis. 1991. Convergence of Approximation Schemes for Fully Nonlinear Second Order Equations. *Asymptotic Analysis* 4 (3):271–283.

Berman, Abraham and Robert J. Plemmons. 1994. *Nonnegative Matrices in the Mathematical Sciences*, vol. 9 of *Classics in Applied Mathematics*. Philadelphia: Society for Industrial and Applied Mathematics.

Bokanowski, Olivier, Stefania Maroso, and Hasnaa Zidani. 2009. Some Convergence Results for Howard's Algorithm. *SIAM Journal on Numerical Analysis* 47 (4):3001–3026.

Borovička, Jaroslav and John Stachurski. 2020. Necessary and Sufficient Conditions for Existence and Uniqueness of Recursive Utilities. *Journal of Finance* 75 (3):1457–1493.

Briani, Ariela, Fabio Camilli, and Hasnaa Zidani. 2012. Approximation Schemes for Monotone Systems of Nonlinear Second-Order Partial Differential Equations: Convergence Result and Error Estimate. *Differential Equations and Applications* 4 (2):297–317.

Christensen, Timothy M. 2022. Existence and Uniqueness of Recursive Utilities without Boundedness. *Journal of Economic Theory* 200:105413.

Davis, Steven J. and Till von Wachter. 2011. Recessions and the Costs of Job Loss. *Brookings Papers on Economic Activity* 2011 (2):1–72.

Duffie, Darrell and Larry G. Epstein. 1992. Stochastic Differential Utility. *Econometrica* 60 (2):353–394.

Dumas, Bernard, Raman Uppal, and Tan Wang. 2000. Efficient Intertemporal Allocations with Recursive Utility. *Journal of Economic Theory* 93 (2):240–259.

El Karoui, N., S. Peng, and M. C. Quenez. 1997. Backward Stochastic Differential Equations in Finance. *Mathematical Finance* 7 (1):1–71.

Epstein, Larry G. and Stanley E. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57 (4):937–969.

Geoffard, Pierre-Yves. 1996. Discounting and Optimizing: Capital Accumulation Problems as Variational Minmax Problems. *Journal of Economic Theory* 69 (1):53–70.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics? *Econometrica* 89 (5):2303–2339.

Hansen, Lars Peter and José A. Scheinkman. 2012. Recursive Utility in a Markov Environment with Stochastic Growth. *Proceedings of the National Academy of Sciences* 109 (30):11967–11972.

Howard, Ronald A. 1960. *Dynamic Programming and Markov Processes*. Oxford, England: John Wiley.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. Earnings Losses of Displaced Workers. *American Economic Review* 83 (4):685–709.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. Wage Risk and Employment Risk over the Life Cycle. *American Economic Review* 100 (4):1432–1467.

Marinacci, Massimo and Luigi Montrucchio. 2010. Unique Solutions for Stochastic Recursive Utilities. *Journal of Economic Theory* 145 (5):1776–1804.

Plemmons, Robert J. 1977. M-matrix Characterizations. I—Nonsingular M-matrices. *Linear Algebra and its Applications* 18 (2):175–188.

Pohl, Walter, Karl Schmedders, and Ole Wilms. 2019. Relative Existence for Recursive Utility. Working paper, NHH Norwegian School of Economics.

———. 2024. Existence of the Wealth-Consumption Ratio in Asset Pricing Models with Recursive Preferences. *The Review of Financial Studies* 37 (3):989–1032.

Puterman, Martin L. and Shelby L. Brumelle. 1979. On the Convergence of Policy Iteration in Stationary Dynamic Programming. *Mathematics of Operations Research* 4 (1):60–69.

Ren, Guanlong and John Stachurski. 2018. Dynamic Programming with Recursive Preferences: Optimality and Applications. arXiv:1812.05748.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. Consumption and Risk Sharing over the Life Cycle. *Journal of Monetary Economics* 51 (3):609–633.

Appendix

A A toy model of the square-root law

The square-root law of Proposition 4.2 does not depend on recursive preferences, on income risk, or on discounting. This appendix isolates the mechanism in a deterministic control problem with two ingredients: a linear flow cost of remaining above a boundary, and a strictly convex cost of moving toward it. Let $a_t \geq 0$ denote the state and $s_t \leq 0$ its drift, $\dot{a}_t = s_t$, with $a(0) = a_0 \geq 0$, and consider

$$U(a_0) := \sup_s \int_0^\infty [-Ba_t - \eta(s_t)]dt, \qquad B > 0,$$

where the supremum is over *admissible controls*: measurable, locally integrable $s_t \leq 0$ such that $a_t = a_0 + \int_0^t s_\tau d\tau \geq 0$ for all t. The parameter B is the flow cost of being away from the boundary and η is the cost of adjustment.

Assumption A.1. *The cost function satisfies $\eta \in C^2((-\infty, 0])$, $\eta(0) = 0$, and $\eta''(s) \geq m > 0$ for every $s \leq 0$ and some constant m.*

Proposition A.1. *Under Assumption A.1, $U(a_0)$ is finite for every $a_0 \geq 0$. The optimal policy satisfies $s(0) = 0$ and, for $a > 0$, $s(a)$ is the unique negative solution of $G(s(a)) = Ba$, where $G(s) := s\eta'(s) - \eta(s)$. The value function is $U(a) = \int_0^a \eta'(s(z))dz$, which is $C^1([0,\infty)) \cap C^2((0,\infty))$ and strictly concave, and, with $\kappa := \sqrt{2B/\eta''(0)}$,*

$$\lim_{a \downarrow 0} \frac{s(a)}{\sqrt{a}} = -\kappa. \tag{25}$$

Proof. Because $G(0) = 0$ and $G'(s) = s\eta''(s) < 0$ for $s < 0$, with $G(s) = \int_0^s z\eta''(z)dz \geq ms^2/2$, the function G maps $(-\infty, 0]$ continuously and one-to-one onto $[0,\infty)$, so $G(s(a)) = Ba$ defines a negative control at every positive state. The implied $s(a)$ is continuous at zero and C^1 for $a > 0$. Implicit differentiation gives $s(a)\eta''(s(a))s'(a) = B$, so $s'(a) < 0$. Continuity of η'' at zero gives $\lim_{s \uparrow 0} 2G(s)/s^2 = \eta''(0)$, and combining this with $G(s(a)) = Ba$ proves (25). A standard verification argument then shows that U solves the original problem. ∎

The braking argument. The proposition can be read as a statement about acceleration. Along the optimal path $\dot{a}_t = s(a_t)$, so by the chain rule

$$\ddot{a}_t = \frac{d}{dt}s(a_t) = s(a_t)s'(a_t) = \frac{1}{2}\frac{d}{da}s(a)^2\Big|_{a=a_t}.$$

Differentiating the HJB equation $0 = \max_{s \leq 0}\{-Ba - \eta(s) + U'(a)s\}$ with respect to a and using the envelope theorem gives $0 = -B + sU''(a)$, while differentiating the first-order condition $\eta'(s) = U'(a)$ gives $U''(a) = \eta''(s)s'(a)$. Combining the two,

$$\ddot{a}_t = s(a_t)s'(a_t) = \frac{B}{\eta''(s(a_t))} \longrightarrow \frac{B}{\eta''(0)} \qquad \text{as } a_t \downarrow 0. \tag{26}$$

The state decelerates toward the boundary at a rate that stays bounded away from both zero and infinity: away from zero because $B > 0$ gives the agent a strict reason to keep moving, and away from infinity because moving is strictly convexly costly.

A bounded, nonzero deceleration forces the square root. If $s(a) \approx -\kappa a^p$ near the boundary for some $p > 0$, then $s(a)s'(a) = \kappa^2 p a^{2p-1}$, which has a finite nonzero limit only when $p = 1/2$, in which case the limit is $\kappa^2/2$. Matching to (26) gives $\kappa^2 = 2B/\eta''(0)$, which is (25). The policy function itself has infinite slope at the boundary, $s'(a) \to -\infty$, but $s(a) \to 0$ at exactly the rate that keeps the product ss' finite; the square root is the unique shape with this property. In the time domain this is the elementary braking problem: a state decelerating at constant rate C and coming to rest at zero at time T satisfies $\dot{a}(t) = C(t - T)$ and $a(t) = \frac{1}{2}C(t-T)^2$, so eliminating t gives $\dot{a} = -\sqrt{2Ca}$. In particular the boundary is reached in finite time, $T = 2\sqrt{a_0}/\kappa$, in contrast with linear decay $\dot{a} = -\lambda a$, under which it is only approached asymptotically.

The consumption-savings problem of Section 4.4 has the same structure. Equation (21), which is established in the course of proving Proposition 4.2, is (26) with the impatience wedge in the role of the flow cost B and $\underline{c}_0/\rho$ in the role of the inverse curvature $1/\eta''(0)$: the wedge is what makes the constrained agent willing to run assets down, and the elasticity of intertemporal substitution governs how costly it is to do so quickly. Formula (20) is (25) under these substitutions.

B Miscellaneous

Proof of Lemma 3.2. Suppose first that $\rho \neq 1$ and put $z(\nu) = (\nu/\delta - \theta)/(1-\theta)$. On D_ν, $z > 0$ and

$$g'(\nu) = \frac{z(\nu)^{-\theta}}{\delta}, \qquad g''(\nu) = -\frac{\theta}{\delta^2(1-\theta)} z(\nu)^{-\theta-1}.$$

The second derivative of $F(c,\nu) - \nu w$ has the sign of $-1/(\gamma - \rho)$. The objective is therefore strictly concave when $\gamma > \rho$ and strictly convex when $\gamma < \rho$. Its first-order condition is

$$z(\nu)^{-\theta} = \frac{(1-\gamma)w}{c^{1-\gamma}},$$

which gives the uniqueness of ν^*; the implied $z(\nu^*)$ is strictly positive, so $\nu^* \in D_\nu$. When $\rho = 1$, the second derivative has the sign of $1-\gamma$, and the first-order condition is $e^{\nu/\delta - 1} = (1-\gamma)w/c^{1-\gamma}$. Strict curvature makes these stationary points the unique global extrema and substitution gives the representation (12). ■

We now note that the solution to the income fluctuation problem may be obtained in closed form when there are no income shocks or ad hoc borrowing constraints (so that the assets are only restricted to lie above the negative of the present discounted value of income). If we write y for the constant value of log income and suppose that the borrowing constraint is set at the natural level $\underline{a} = -e^y/r$, then the HJB equation becomes

$$0 = \max_{c \geq 0} \ \delta\theta W(a)\Big(c^{1-\rho}((1-\gamma)W(a))^{-1/\theta} - 1\Big) + W'(a)(ra + e^y - c). \tag{27}$$

We denote the solution to (27) for a fixed y by W^y.

Lemma B.1. *If $\delta > r(1-\rho)$ then the solution W^y to (27) is*

$$W^y(a) = \frac{\overline{W}(a + e^y/r)^{1-\gamma}}{1-\gamma}$$

where $\overline{W} = \delta^\theta(\delta/\rho - r(1/\rho - 1))^{-\rho\theta}$ while the drift in total wealth is $r - \bar{c} = (r-\delta)/\rho$.

Proof. We guess-and-verify that this HJB has a solution of the above form for some scalar $\overline{W} > 0$ and that $c(a) = \bar{c}(a + e^y/r)$ for some constant $\bar{c} > 0$. Substituting into the HJB gives

$$\begin{aligned} 0 = \max_{\bar{c}>0} \ & \frac{\delta\overline{W}}{1-\rho}(a + e^y/r)^{1-\gamma}\left(\bar{c}^{1-\rho}(a + e^y/r)^{1-\rho}\left(\overline{W}(a + e^y/r)^{1-\gamma}\right)^{-1/\theta} - 1\right) \\ & + (r - \bar{c})\overline{W}(a + e^y/r)^{1-\gamma} \end{aligned}$$

which simplifies to

$$0 = \max_{\bar{c}>0} \ \frac{\delta}{1-\rho}\left(\bar{c}^{1-\rho}\overline{W}^{-1/\theta} - 1\right) + r - \bar{c}. \tag{28}$$

The first-order condition gives $\overline{c}^{\rho} = \delta \overline{W}^{-1/\theta}$ and substitution into (28) gives

$$0 = \frac{\delta}{1-\rho}\left(\delta^{1/\rho-1}\overline{W}^{-(1-\rho)\frac{1/\rho-1}{1-\gamma}}\overline{W}^{-\frac{1-\rho}{1-\gamma}} - 1\right) + r - \delta^{1/\rho}\overline{W}^{-\frac{1/\rho-1}{1-\gamma}}$$

$$(1-\rho)\delta^{1/\rho}\overline{W}^{-\frac{1/\rho-1}{1-\gamma}} = \delta^{1/\rho}\overline{W}^{-\frac{1/\rho-1}{1-\gamma}} - \delta + r(1-\rho)$$

and hence $(\delta - r(1-\rho))\overline{W}^{\frac{1/\rho-1}{1-\gamma}} = \delta^{1/\rho}\rho$, which rearranges to $\overline{W}$. The constant $\overline{c}$ is then given by

$$\overline{c} = \delta^{1/\rho}\overline{W}^{-\frac{1/\rho-1}{1-\gamma}} = \delta^{1/\rho}\left(\frac{\delta^{1/\rho}\rho}{\delta - r(1-\rho)}\right)^{-1} = \delta/\rho - r(1/\rho - 1)$$

which rearranges as claimed. ■

C Comparison principle and convergence

In this appendix we establish a comparison principle for the bounded-asset, bounded-consumption problem with finitely many income states and prove convergence of the upwind discretization in Section 4. The proof uses standard viscosity-solution techniques, which we present in a self-contained form tailored to our finite-state income process and state constraints.[12] Definition 4 introduces the solution notion. Proposition C.1 establishes the subsolution and supersolution properties of the upper and lower half-relaxed limits. Proposition C.2 establishes comparison, and Corollary C.3 combines these results to obtain uniform convergence.

Continuous-state problem. We first define the continuous-state problem formally and relate it to the problems considered in the main text. We write $I := [\underline{a}, \overline{a}]$ for the asset space and $I^\circ := (\underline{a}, \overline{a})$ for its interior. We again assume $r \in (0, \delta)$ and that income takes the values e^{y_j}, $j = 0, \dots, N_y$, where $y_0 < \cdots < y_{N_y}$ with transition intensities $\lambda_{jk} \geq 0$, and we choose c_- and c_+ such that

$$0 < c_- < r\underline{a} + e^{y_0}, \qquad r\overline{a} + e^{y_{N_y}} < c_+ < \infty, \tag{29}$$

and the control sets

$$\mathcal{C}_j(a) = \begin{cases} [c_-, c_+], & \underline{a} < a < \overline{a}, \\ [c_-, r\underline{a} + e^{y_j}], & a = \underline{a}, \\ [r\overline{a} + e^{y_j}, c_+], & a = \overline{a}. \end{cases}$$

The assumption (29) motivates the definition

$$\eta := \min\{r\underline{a} + e^{y_0} - c_-, c_+ - r\overline{a} - e^{y_{N_y}}\} > 0.$$

Given the Epstein–Zin aggregator f and the savings function $s_j(a, c) = ra + e^{y_j} - c$, for any vector

[12] Achdou and Tang (2026a) study an asset domain that is unbounded above, and their main analysis assumes, in our notation, $\gamma > 1$ and $\rho > \gamma$. We instead consider a bounded asset interval and a compact consumption set bounded away from zero. In this setting, the transformed comparison argument below applies to both $\gamma > \rho$ and $\rho > \gamma$, subject to the parameter assumptions maintained throughout the paper.

$w = (w_0, \ldots, w_{N_y})$ we define

$$\mathcal{B}_j(a, w, p) = \max_{c \in \mathcal{C}_j(a)} f(c, w_j) + s_j(a, c)p + \sum_k \lambda_{jk}(w_k - w_j).$$

The HJB equation is then

$$\mathcal{B}_j(a, w(a), w_j'(a)) = 0, \quad a \in [\underline{a}, \overline{a}],\ j \in \{0, \ldots, N_y\}. \tag{30}$$

Viscosity solutions. Before stating and proving any claims regarding the convergence of the discrete solutions characterized in the main text, we must define a precise solution notion for the HJB equation (30). Our setting differs from the environment in Barles and Souganidis (1991) because of the presence of discrete income values, which makes it an example of what Briani et al. (2012) call a *weakly coupled system*.[13] We then have the following notion, adapted from Achdou and Tang (2026a) and Briani et al. (2012).

Definition 4. *An upper semicontinuous (u.s.c.) function $U : I \to \mathbb{R}^{N_y+1}$ is a viscosity subsolution to (30) on $\Omega \subseteq I$ if for all j, if $U_j - \varphi$ has a local maximum (relative to I) at $a \in \Omega$ with $\varphi \in C^1$, then*

$$\mathcal{B}_j(a, U(a), \varphi'(a)) \geq 0.$$

A lower semicontinuous (l.s.c.) function V is a viscosity supersolution to (30) on $\Omega \subseteq I$ if for all j, if $V_j - \varphi$ has a local minimum (relative to I) at $a \in \Omega$ with $\varphi \in C^1$, then

$$\mathcal{B}_j(a, V(a), \varphi'(a)) \leq 0.$$

A continuous function on I is a constrained viscosity solution *if it is a viscosity subsolution on I and a viscosity supersolution on I°.*

We write W^h for the solution to the discrete Bellman equation in Definition 1 for the A operator and control sets $\mathcal{C}$ given in Definition 3, where $h := (\overline{a} - \underline{a})/N_a$ indicates the fineness of the grid, now denoted by X_h to illustrate the dependence on h. Similarly, in this appendix, we illustrate the dependence of the discrete Bellman function on h with a superscript, so that W^h is the solution to

$$B_{ij}^h(W^h) = 0, \quad i = 0, \ldots, N_a,\ j = 0, \ldots, N_y.$$

Further, we know that $W_{ij}^h \in [u(c_-), u(c_+)]$ for $i = 0, \ldots, N_a$ and $j = 0, \ldots, N_y$, which ensures that the following upper and lower half-relaxed limits are well-defined and $[u(c_-), u(c_+)]$-valued,

$$\overline{W}_j(a) = \limsup_{\substack{h \to 0 \\ a_i \in X_h,\ a_i \to a}} W_{ij}^h, \qquad \underline{W}_j(a) = \liminf_{\substack{h \to 0 \\ a_i \in X_h,\ a_i \to a}} W_{ij}^h. \tag{31}$$

The following applies the argument of Theorem 2.1 of Barles and Souganidis (1991) componentwise as in Proposition 3.3 of Briani et al. (2012).

Proposition C.1 (half-relaxed limits)**.** *The half-relaxed limits in (31) have the following properties:*

[13] See their Example 2.4. Also note the different convention: our B is their $-F$.

1. *$\overline{W}$ is an upper semicontinuous viscosity subsolution on I.*
2. *$\underline{W}$ is a lower semicontinuous viscosity supersolution on I°.*
3. *For every j and $a_0 \in \{\underline{a}, \overline{a}\}$,*

$$\underline{W}_j(a_0) = \liminf_{\substack{a \to a_0 \\ a \in I^\circ}} \underline{W}_j(a). \tag{32}$$

Proof. We first establish the subsolution and supersolution inequalities for interior points. Suppose that for some $\phi \in C^1(I)$ and index j, $\overline{W}_j - \phi$ attains a maximum at some $a \in I^\circ$. Because there is no loss in assuming that this maximum is strict in a neighborhood of a, we may choose sequences $h_n \downarrow 0$ and $a_{i_n} \in X_{h_n}$ satisfying $a_{i_n} \to a$, $W^{h_n}_{i_n j} \to \overline{W}_j(a)$, and $W^{h_n}_{l j} \le W^{h_n}_{i_n j} + \phi(a_l) - \phi(a_{i_n})$ for $n \ge 1$ and indices l adjacent to i_n on the grid X_{h_n}. We therefore define

$$\widetilde{W}^n_{lk} = \begin{cases} W^{h_n}_{i_n j} + \phi(a_l) - \phi(a_{i_n}) & k = j, \\ W^{h_n}_{i_n k} & k \ne j \end{cases}$$

and note that because $\widetilde{W}^n_{ij} \ge W^{h_n}_{ij}$ for i adjacent to i_n, with equality at i_n, we have $B^{h_n}_{i_n j}(\widetilde{W}^n) \ge B^{h_n}_{i_n j}(W^{h_n}) = 0$, because the $B^{h_n}_j$ operator places nonnegative weights on off-diagonal asset points and the above construction leaves the values at the other income points at i_n unchanged. By considering a subsequence if necessary, we may assume that $W^{h_n}_{i_n \cdot}$ converges to a vector w satisfying $w_j = \overline{W}_j(a)$ and $w_k \le \overline{W}_k(a)$, $k \ne j$, so that

$$\mathcal{B}_j(a, \overline{W}(a), \phi'(a)) \ge \mathcal{B}_j(a, w, \phi'(a)) = \lim_{n \to \infty} B^{h_n}_{i_n j}(\widetilde{W}^n) \ge 0$$

which proves the subsolution property for interior points. The supersolution property at interior points follows from a symmetric argument.

For the subsolution property on the boundary points, we fix ϕ as in the definition for $a_0 \in \{\underline{a}, \overline{a}\}$ and first note we have

$$\max_{c \in [c_-, c_+]} f(c, \overline{W}_j(a_0)) + s_j(a_0, c)\phi'(a_0) + \sum_k \lambda_{jk}(\overline{W}_k(a_0) - \overline{W}_j(a_0)) \ge 0. \tag{33}$$

After passing to a subsequence, the points at which the maximization occurs are either interior nodes, which use the full control interval, or endpoints, whose control set is a subset of that interval. The off-component passage is the same as above (i.e. $\overline{W}(a_0)$ weakly dominates any limit w in the components $k \ne j$). To establish (33) for the true control set $C_j(a_0)$ (instead of $[c_-, c_+]$), define $w = \overline{W}(a_0)$, $m = ra_0 + e^{y_j} \in (c_-, c_+)$, $p = \phi'(a_0)$, $p_c = f_c(m, w_j)$, and

$$\widehat{p} = \begin{cases} \max\{p, p_c\}, & a_0 = \underline{a}, \\ \min\{p, p_c\}, & a_0 = \overline{a}. \end{cases}$$

The Epstein–Zin aggregator is concave in consumption, and so

$$\max_{c \in C_j(a_0)} f(c, w_j) + (m - c)p = \max_{c \in [c_-, c_+]} f(c, w_j) + (m - c)\widehat{p}. \tag{34}$$

To see (34), note that at $a_0 = \underline{a}$ the maximizer over $[c_-, c_+]$ lies in $C_j(a_0)$ if $p \geq p_c$, while if $p < p_c$, (34) holds because both sides reduce to $f(m, w_j)$. The justification for (34) at $a_0 = \overline{a}$ is analogous. The function $\widehat{\phi}(a) = \phi(a) + (\widehat{p} - p)(a - a_0)$ remains an upper test relative to I and so (33) is also satisfied with $\widehat{\phi}$ in place of ϕ, which together with (34) gives $\mathcal{B}_j(a_0, \overline{W}(a_0), \phi'(a_0)) \geq 0$.

To prove (32), we use the fact that W^h solves $0 = \max_{c \in \mathcal{C}} f(c, W^h) + A(c)W^h$ (for $A(c)$ given in Definition 3) to place bounds on the finite differences of W^h. To make this precise, note that the compactness of the value and control sets implies the existence of a constant C_0 bounding $|f(c, W^h_{ij}) + \sum_k \lambda_{jk}(W^h_{ik} - W^h_{ij})|$ uniformly in h, i, j, c, while the bounds (29) ensure that $s_j(a, c_-) \geq \eta$ and $s_j(a, c_+) \leq -\eta$ on I. For each adjacent pair of indices $i, i+1$, we evaluate the right-hand side of the discrete Bellman equation at c_- at grid point i, and at c_+ at grid point $i+1$. Because both controls must give nonpositive residuals (by the definition of W^h), we can rearrange and obtain

$$-\frac{C_0}{\eta} \leq \frac{W^h_{i+1,j} - W^h_{ij}}{h} \leq \frac{C_0}{\eta}, \qquad i = 0, \ldots, N_a - 1,$$

which is independent of h and shows that $\underline{W}$ is continuous at the endpoints and satisfies (32).

Finally, both $\underline{W}$ and $\overline{W}$ have the stated semicontinuity properties and satisfy $\underline{W} \leq \overline{W}$ simply by the definitions of liminf and limsup. ■

The following is the continuous-state analogue of the discrete comparison principle in Lemma 3.5.

Proposition C.2 (comparison)**.** *Suppose that U is a viscosity subsolution of (30) on I and V is a viscosity supersolution of (30) on I° that satisfies (32) (with V in place of $\underline{W}$), and that both assume values in $[u(c_-), u(c_+)]$. Then $U_j(a) \leq V_j(a)$ for every $a \in I$ and j and so there is at most one $[u(c_-), u(c_+)]$-valued, constrained viscosity solution.*

Proof. We first define $Z(W) := \ln u^{-1}(W)$ and $\Psi_\gamma(d) = (e^{(1-\gamma)d} - 1)/(1-\gamma)$ and consider the transformed Bellman residual

$$\widetilde{\mathcal{B}}_j(a, z, p) = \max_{c \in C_j(a)} l(c, z_j) + (ra + e^{y_j} - c)p + \sum_{k \neq j} \lambda_{jk} \Psi_\gamma(z_k - z_j)$$

where

$$l(c, z) = \begin{cases} \frac{\delta}{1-\rho}\left[(ce^{-z})^{1-\rho} - 1\right], & \rho \neq 1, \\ \delta(\ln c - z), & \rho = 1. \end{cases}$$

We define $z_U(a) := Z(U(a))$ and $z_V(a) := Z(V(a))$ and note that if $H(z) := u(e^z)$ then

$$\mathcal{B}_j(a, H(z), H'(z_j)p) = H'(z_j)\widetilde{\mathcal{B}}_j(a, z, p)$$

which together with $H'(z) > 0$ implies that z_U and z_V satisfy the subsolution and supersolution inequalities on I and I°, respectively, for the transformed Bellman residual. The benefit of working with the transformed Bellman residual $\widetilde{\mathcal{B}}$ instead of the original $\mathcal{B}$ is that for all $c > 0$ we have $l_z(c, z) = -\delta(ce^{-z})^{1-\rho}$ and so there exists $\kappa > 0$ such that

$$l_z(c, z) \leq -\kappa \qquad \text{for all } (c, z) \in [c_-, c_+] \times [\ln c_-, \ln c_+]. \tag{35}$$

That is, the function $l(c,z)$ is strictly decreasing on the relevant domain, even though the original function $f(c,w)$ is not. We now let (j^*,a^*) be a pair attaining the maximum in

$$M^* = \max_{\substack{j\in\{0,\dots,N_y\}\\ a\in I}} z_{Uj}(a) - z_{Vj}(a), \tag{36}$$

which exists by the semicontinuity properties of U and V and the compactness of I. We will assume that $M^* > 0$ and derive a contradiction. First note that if U and V (and hence z_U and z_V) are smooth and $a^* \in I^\circ$, then $z'_{Uj^*}(a^*) = z'_{Vj^*}(a^*)$ and

$$z_{Uk}(a^*) - z_{Uj^*}(a^*) \le z_{Vk}(a^*) - z_{Vj^*}(a^*) \quad \text{for every } k = 0,\dots,N_y. \tag{37}$$

Maximizing over the controls then yields

$$\begin{aligned} 0 &\le \widetilde{\mathcal{B}}_{j^*}(a^*, z_U(a^*), z'_{Uj^*}(a^*)) && (z_U \text{ is a subsolution}) \\ &\le \widetilde{\mathcal{B}}_{j^*}(a^*, z_V(a^*), z'_{Vj^*}(a^*)) - \kappa M^* && (35) \text{ and } (37) \\ &\le -\kappa M^* && (z_V \text{ is a supersolution}) \end{aligned} \tag{38}$$

which is a contradiction and establishes the comparison principle for *smooth* functions when the maximum in (36) is attained at an *interior point*. To deal with non-smooth functions and boundaries, we consider the following "doubling the variables" technique: instead of maximizing $z_{Uj}(a) - z_{Vj}(a)$, we fix $\epsilon > 0$ and let $(j_\epsilon, a_\epsilon, b_\epsilon)$ denote the values attaining the maximum in

$$M(P_\epsilon) := \max_{\substack{j\in\{0,\dots,N_y\}\\ (a,b)\in I^2}} z_{Uj}(a) - z_{Vj}(b) - P_\epsilon(a,b)$$

for convenient choices of "penalty functions" P_ϵ with the following properties:

1. $\lim_{\epsilon\downarrow 0} P_\epsilon(a_\epsilon, b_\epsilon) = \lim_{\epsilon\downarrow 0}(P_{\epsilon a} + P_{\epsilon b})(a_\epsilon, b_\epsilon) = 0$;
2. $b_\epsilon \in I^\circ$ for sufficiently small $\epsilon > 0$;
3. $P_{\epsilon b}(a_\epsilon, b_\epsilon)$ is bounded in $\epsilon > 0$.

For each choice of P_ϵ satisfying these three properties, we define the $C^1(I)$ functions

$$\begin{aligned} \phi_U(a) &= z_{Uj_\epsilon}(a_\epsilon) + P_\epsilon(a, b_\epsilon) - P_\epsilon(a_\epsilon, b_\epsilon) \\ \phi_V(b) &= z_{Vj_\epsilon}(b_\epsilon) - P_\epsilon(a_\epsilon, b) + P_\epsilon(a_\epsilon, b_\epsilon) \end{aligned}$$

and note that because $b_\epsilon \in I^\circ$, the definitions of subsolutions and supersolutions imply

$$\widetilde{\mathcal{B}}_{j_\epsilon}(a_\epsilon, z_U(a_\epsilon), \phi'_U(a_\epsilon)) \ge 0, \qquad \widetilde{\mathcal{B}}_{j_\epsilon}(b_\epsilon, z_V(b_\epsilon), \phi'_V(b_\epsilon)) \le 0. \tag{39}$$

For interior points, $a^* \in I^\circ$, we define

$$P_\epsilon(a,b) = \frac{(a-b)^2}{\epsilon} + (b-a^*)^2.$$

For boundary points, if $a^* \in \{\underline{a}, \overline{a}\}$, we consider two separate penalty functions that force one of the maximizing points into the interior: let $\sigma = 1$ if $a^* = \underline{a}$ and $\sigma = -1$ if $a^* = \overline{a}$ and define

$$P_\epsilon(a,b) = \frac{(a-b)^2}{\epsilon} + \left[\left(\frac{\sigma(b-a)}{\epsilon} - 1\right)^-\right]^2 + (b-a^*)^2. \tag{40}$$

We now verify that these penalty functions P_ϵ satisfy the above three properties. For Property 1, note that if $a^* \in I^\circ$, evaluating the objective in $M(P_\epsilon)$ at (j^*, a^*, a^*) gives $M(P_\epsilon) \geq M^*$, while if $a^* \in \{\underline{a}, \overline{a}\}$, the condition (32) implies a sequence $\epsilon_n \downarrow 0$ such that $b_n^0 := a^* + \sigma\epsilon_n \in I^\circ$ and $z_{Vj^*}(b_n^0) \to z_{Vj^*}(a^*)$, which gives $M(P_{\epsilon_n}) \geq M^* - o(1)$. In either case, if $d_\epsilon := z_{Uj_\epsilon}(a_\epsilon) - z_{Vj_\epsilon}(b_\epsilon)$ then $d_\epsilon - P_\epsilon(a_\epsilon, b_\epsilon) \geq M^* - o(1)$ and hence $a_\epsilon - b_\epsilon \to 0$. Semicontinuity then gives $\limsup_{\epsilon\to 0} d_\epsilon \leq M^*$ and hence $P_\epsilon(a_\epsilon, b_\epsilon) \to 0$ and $d_\epsilon \to M^*$, while the $(b-a^*)^2$ term gives $b_\epsilon \to a^*$ and $a_\epsilon \to a^*$. Finally, $(P_{\epsilon a} + P_{\epsilon b})(a_\epsilon, b_\epsilon) = 2(b_\epsilon - a^*) \to 0$. For Property 2, if $a^* \in I^\circ$ this follows from $b_\epsilon \to a^*$, while if $a^* \in \{\underline{a}, \overline{a}\}$, the middle term in (40) gives $\sigma(b_\epsilon - a_\epsilon) \geq \epsilon(1 - o(1)) > 0$, and hence $b_\epsilon > \underline{a}$ when $a^* = \underline{a}$ and $b_\epsilon < \overline{a}$ when $a^* = \overline{a}$. For Property 3, if C is a constant such that $\widetilde{\mathcal{B}}_j(a,z,p) \geq \eta|p| - C$ for all $a \in I^\circ$ and $z \in [\ln c_-, \ln c_+]^{N_y+1}$, then $|P_{\epsilon b}(a_\epsilon, b_\epsilon)| \leq C/\eta$.

To produce a contradiction to $M^* > 0$ and complete the proof, we mimic the reasoning in the chain of inequalities (38) and combine with bounds on changes in $\widetilde{\mathcal{B}}$. For any $a \in I$, $b \in I^\circ$ and $z_U, z_V \in [\ln c_-, \ln c_+]^{N_y+1}$ with

$$d := z_{Uj} - z_{Vj} = \max_{k\in\{0,\ldots,N_y\}} z_{Uk} - z_{Vk} > 0,$$

the monotonicity of Ψ_γ and the bound (35) together imply that for any $p, q \in \mathbb{R}$, we have

$$\widetilde{\mathcal{B}}_j(a, z_U, p) - \widetilde{\mathcal{B}}_j(b, z_V, q) \leq -\kappa d + r|a-b||q| + S|p-q| \tag{41}$$

where $S := \sup_{j,a\in I, c\in[c_-,c_+]} |s_j(a,c)|$. Combining (39) with the inequality (41) then implies

$$0 \leq -\kappa d_\epsilon + r|a_\epsilon - b_\epsilon||P_{\epsilon b}(a_\epsilon, b_\epsilon)| + S|P_{\epsilon a}(a_\epsilon, b_\epsilon) + P_{\epsilon b}(a_\epsilon, b_\epsilon)|$$

and the right-hand side tends to $-\kappa M^*$ as $\epsilon \downarrow 0$, a contradiction which completes the proof. ■

The following is the completion of the Barles–Souganidis argument, tailored to our setting.[14]

Corollary C.3 (uniform convergence)**.** *The grid solutions converge uniformly to the unique, continuous, $[u(c_-), u(c_+)]$-valued, constrained viscosity solution W:* $\lim_{h\to 0} \max_{i,j} |W_{ij}^h - W_j(a_i)| = 0$.

Proof. Proposition C.1 verifies that the functions in (31) satisfy the assumptions of Proposition C.2, and so $\underline{W} \leq \overline{W} \leq \underline{W}$ and their common value W is continuous and inherits the required viscosity inequalities. Equality of the half-relaxed limits, compactness of I, and finiteness of the income set give uniform grid convergence. ■

[14] See, e.g., Theorem 2.1 of Barles and Souganidis (1991) and Proposition 3.3 of Briani et al. (2012).