Highlights

The Method of Moderation[⋆]

Christopher D. Carroll, Karsten Chipeniuk, Alan Lujan, Kiichi Tokuoka, Weifeng Wu

- A realist's consumption lies between a pessimist's rule and an optimist's
- Interpolating the logit of where it lies keeps the approximation inside both bounds
- Five gridpoints suffice, and beyond them the method is ten times as accurate as EGM
- The bracket carries over to value, artificial constraints, risky returns, and assets
- For the portfolio share the log of the moderation ratio replaces the logit

The Method of Moderation

Christopher D. Carroll[a,b], Karsten Chipeniuk[c], Alan Lujan[a,b,*], Kiichi Tokuoka[d], Weifeng Wu[e]

[a] *Johns Hopkins University, Krieger School of Arts and Sciences Baltimore, MD,*
[b] *Econ-ARK*
[c] *Reserve Bank of New Zealand*
[d] *Japanese Ministry of Finance*
[e] *Fannie Mae, Washington, DC,*

Abstract

In a risky world, a pessimist assumes the worst; someone who ignores risk is an optimist. Both problems are easy, because pessimist and optimist alike behave as if the world were riskless. A realist, responding optimally to risk, spends between the two. We use that to change what is interpolated: not consumption, but the logit of a moderation ratio recording where between the two rules the realist lies. The result cannot escape the bounds the two rules form, is accurate on five gridpoints, and beyond the last of them is more than an order of magnitude more accurate than the endogenous gridpoints method. It carries over to value and to risky returns, and to the portfolio share with the logarithm in place of the logit.

Keywords: Dynamic Stochastic Optimization, Consumption-Saving Models, Numerical Methods, Endogenous Gridpoints Method, Precautionary Saving, Bounded Approximation, Extrapolation
JEL: D14; C61; G11

1. Introduction

Solving a consumption, investment, portfolio choice, or similar intertemporal optimization problem using numerical methods generally requires the modeler to choose how to represent a policy or value function. In the stochastic case, where analytical solutions are generally not available, a common approach is to

⋆This work was supported by the Alfred P. Sloan Foundation under grant G-2017-9832.

We thank Matthew N. White for many thoughtful conversations about the ideas in this paper. The views presented here are those of the authors and should not be attributed to the Japanese Ministry of Finance, the Reserve Bank of New Zealand, or Fannie Mae.

Every figure and table in this paper is generated by the replication package at https://github.com/econ-ark/method-of-moderation, which also recomputes each number cited in the text.

*Corresponding author

Email address: alujan@jhu.edu (Alan Lujan)

use low-order polynomial splines that exactly match the function (and maybe some derivatives) at a finite set of gridpoints, and then to assume that interpolated or extrapolated versions of that spline represent the function well at the continuous infinity of unmatched points. Outside the range of the gridpoints, however, this trust is often misplaced (Ma and Toda, 2022; White, 2015), and extending the grid does not restore it, because beyond whatever gridpoint is last a linearly extrapolated consumption rule eventually predicts negative precautionary saving, which theory rules out for any prudent consumer.

This paper argues that a better approach is to rely upon the fact that without uncertainty, the optimal consumption function has a simple analytical solution. The key insight is that, under standard assumptions, the consumer who faces an uninsurable labor income risk will consume less than a consumer with the same path for expected income but who does not perceive any uncertainty. The 'realist' consumer, who *does* perceive the risks, will engage in 'precautionary saving' (Leland, 1968; Sandmo, 1970; Kimball, 1990), so the perfect foresight riskless solution provides an upper bound to the solution that will actually be optimal. A lower bound is provided by the behavior of a consumer who has the subjective belief that the future level of income will be the worst that it can possibly be. This consumer, too, behaves according to the convenient analytical perfect foresight solution, but their certainty is that of a pessimist perfectly confident in their pessimism. The realist's consumption lies strictly between these two immoderate benchmarks. We use this fact to change the object being interpolated, so that rather than consumption itself we interpolate a moderation ratio that measures where the realist falls between the pessimist and the optimist. Because that ratio is bounded and smooth and its logit is asymptotically linear, the logit interpolates and extrapolates far more accurately than when the level of the consumption ratio is approximated, and by construction the result can never violate the bounds.

On a standard calibration the method is more accurate than the benchmark endogenous gridpoints method (EGM) in every interval between gridpoints, and by more than an order of magnitude in all but the first. Beyond the last gridpoint, where EGM's default linear extrapolation fails outright, it is roughly forty-six times more accurate over the range we evaluate, and even when EGM is lent the same bounds to hold its extrapolation inside them, it remains less accurate there, by a factor of seventeen.

We build on bounds for the consumption function and limiting marginal propensities to consume established in buffer-stock theory (Stachurski and Toda, 2019; Ma et al., 2020; Carroll, 2009; Ma and Toda, 2021) for the standard buffer-stock problem (Carroll, 1997). Using results from Carroll and Shanker (2026), we show how to use these upper and lower bounds to tightly constrain the shape and characteristics of the solution to the problem of the 'realist.' Imposition of these constraints can clarify and speed the solution of the realist's problem.[1]

[1]An early version of the paper was presented at the 2012 Society for Economic Dynamics meetings by Wu, Tokuoka, and Carroll, and Carroll (2020b) discusses the method in lecture-

After showing how to use the method in the baseline case, we show how to refine it to encompass an even tighter theoretical bound, and how to extend it to solve a problem in which the consumer faces both labor income risk and rate-of-return risk. Section 2 states the realist's problem and the patience conditions that guarantee a finite solution, and Section 3 sets out the endogenous gridpoints benchmark. Section 4 develops the method of moderation for the consumption function and, in Section 4.3, for the value function; Section 5 quantifies its accuracy against endogenous gridpoints; and Section 6 takes up four extensions: the tighter bound near the borrowing constraint (Section 6.1), Hermite interpolation that matches the Euler equation at each gridpoint (Section 6.2), an artificial borrowing constraint (Section 6.3), and a problem stated in assets rather than market resources (Section 6.4). Section 7 makes the return risky and then makes the risky share a choice, which is where the method first meets a bound that the solution actually reaches at a finite level of wealth instead of merely approaching.

2. The Realist's Problem

The truly optimal behavior in the problem facing the consumer who understands all their risks is captured by

$$\max\ \mathbb{E}_t \left[\sum_{n=0}^{T-t} \beta^n \mathbf{u}(\mathbf{c}_{t+n}) \right] \tag{1}$$

where the utility function is CRRA with risk aversion parameter $\rho > 0$:

$$\mathbf{u}(c) = \begin{cases} \frac{c^{1-\rho}}{1-\rho} & \text{if } \rho \neq 1 \\ \log c & \text{if } \rho = 1. \end{cases} \tag{2}$$

Maximization is subject to the budget constraints

$$\begin{aligned} \mathbf{a}_t &= \mathbf{m}_t - \mathbf{c}_t \\ \mathbf{p}_{t+1} &= \mathbf{p}_t \mathcal{G}_{t+1} \\ \mathbf{y}_{t+1} &= \mathbf{p}_{t+1} \xi_{t+1} \\ \mathbf{m}_{t+1} &= \mathbf{a}_t \mathrm{R}_{t+1} + \mathbf{y}_{t+1} \end{aligned} \tag{3}$$

where the variables are defined as

note form. This paper supersedes both accounts.

$$
\begin{aligned}
\beta &\text{ - pure time discount factor} \\
\mathbf{a}_t &\text{ - assets at the end of period } t \\
\mathbf{c}_t &\text{ - consumption in period } t \\
\mathbf{m}_t &\text{ - `market resources' available for consumption} \\
\mathbf{p}_{t+1} &\text{ - `permanent labor income' in period } t+1 \\
\mathrm{R}_{t+1} &\text{ - gross interest rate from period } t \text{ to } t+1 \\
\mathbf{y}_{t+1} &\text{ - noncapital income in period } t+1.
\end{aligned} \tag{4}
$$

The exogenous variables evolve according to the *Friedman-Muth Income Process*[2]:

$$
\begin{aligned}
\mathcal{G}_{t+1} &= G_{t+1}\psi_{t+1} \\
\xi_{t+1} &= \begin{cases} \underline{\xi} & \text{with probability } \wp > 0 \\ \left(\frac{1-\wp\,\underline{\xi}}{1-\wp}\right)\theta_{t+1} & \text{with probability } (1-\wp) \end{cases}
\end{aligned} \tag{5}
$$

where G_{t+1} is the deterministic permanent income growth factor, and the permanent shocks to income ψ_{t+1} are independently and identically distributed with mean $\mathbb{E}[\psi_{t+1}] = 1$ and support $[\underline{\psi}, \bar{\psi}]$ where $0 < \underline{\psi} \leq 1 \leq \bar{\psi} < \infty$.[3] The worst transitory draw $\underline{\xi} \geq 0$ arrives with probability $\wp$. The Friedman-Muth unemployment case sets $\underline{\xi} = 0$; a positive $\underline{\xi}$ is the same unemployment event paying a benefit, which leaves income at a positive floor rather than at zero. Everything below requires only that $\underline{\xi}$ is indeed the worst draw, $\underline{\xi} \leq \left(\frac{1-\wp\,\underline{\xi}}{1-\wp}\right)\underline{\theta}$, which holds automatically when $\underline{\xi} = 0$. Transitory shocks other than the worst draw, θ_{t+1}, are independently distributed with mean $\mathbb{E}[\theta_{t+1}] = 1$ and bounded support $\underline{\theta} \leq \theta_{t+1} \leq \bar{\theta}$ where $0 \leq \underline{\theta} \leq 1 \leq \bar{\theta} < \infty$; the scaling factor in (5) keeps $\mathbb{E}[\xi_{t+1}] = 1$.

It turns out (see Carroll (2020b) for a proof) that this problem can be rewritten in a more convenient form in which choice and state variables are normalized by the level of permanent income, e.g., using nonbold font for normalized variables, $m_t = \mathbf{m}_t/\mathbf{p}_t$. When that is done, the Bellman equation for the transformed version of the consumer's problem is

$$
\begin{aligned}
\mathbf{v}_t(m_t) &= \max_{c_t} \quad \mathbf{u}(c_t) + \beta\mathbb{E}_t[\mathcal{G}_{t+1}^{1-\rho}\mathbf{v}_{t+1}(m_{t+1})] \\
&\text{s.t.} \\
a_t &= m_t - c_t \\
m_{t+1} &= (\mathrm{R}/\mathcal{G}_{t+1})\, a_t + \xi_{t+1},
\end{aligned} \tag{6}
$$

[2]The permanent-transitory decomposition is Friedman (1957); the stochastic process form is Muth (1960).

[3]Bounded support ($\bar{\psi} < \infty$) gives the consumption function well-defined upper and lower bounds. Results for unbounded shocks with finite moments (lognormal, for example) exist, and we do not pursue that extension.

and because we have not imposed a liquidity constraint, the solution satisfies the Euler equation

$$\mathbf{u}'(c_t) = \beta \mathrm{R} \mathbb{E}_t[\mathcal{G}_{t+1}^{-\rho} \mathbf{u}'(c_{t+1})]. \tag{7}$$

We define the *absolute patience factor* $\text{Þ} \equiv (\beta \mathrm{R})^{1/\rho}$, the factor by which a perfect foresight consumer's consumption grows from one period to the next. Whether the problem has a well-behaved solution depends on a set of conditions established and named in Carroll and Shanker (2026), whose names we adopt and whose content we state here, because we appeal to them repeatedly below.

Three conditions govern the solution. The *finite value of autarky condition* (FVAC), $0 < \beta G^{1-\rho} \mathbb{E}[\psi^{1-\rho}] < 1$, says that a consumer who simply consumed their income in every period ('autarky') would obtain finite lifetime value, so that the objective being maximized is finite to begin with. The *absolute impatience condition* (AIC), $\text{Þ} < 1$, says that the perfect foresight consumer plans a falling rather than a rising consumption path; it is what makes the consumer want to run a buffer stock down rather than accumulate without limit. The *return impatience condition* (RIC), $\text{Þ}/\mathrm{R} < 1$, says that planned consumption growth falls short of the rate at which resources accumulate, so the perfect foresight consumer spends a strictly positive fraction of their resources ($\underline{\boldsymbol{\kappa}} = 1 - \text{Þ}/\mathrm{R} > 0$ in (9) below) instead of deferring spending indefinitely.

Two further conditions relate spending to the growth of income. The *growth impatience condition* (GIC), $\text{Þ}/G < 1$, says that consumption grows more slowly than permanent income, which is what makes the ratio of wealth to permanent income revert to a target instead of drifting away. The *finite human wealth condition* (FHWC), $G/\mathrm{R} < 1$, says that the present discounted value of future labor income is finite, so that the 'optimist' rule we construct below exists at all.

Together these conditions ensure existence of the upper and lower bounds on consumption that the method of moderation is built from (Carroll, 2009; Stachurski and Toda, 2019) and pin down the limiting MPCs (Ma and Toda, 2021).

We simplify the exposition by setting $G = 1$ and assuming $\psi_{t+n} = 1$ with probability 1 for all $n > 0$ (no permanent income growth or shocks), and drop time subscripts except where context requires, working with the infinite-horizon formulation.[4] Under these simplifications, FVAC becomes $0 < \beta < 1$, the GIC coincides with the AIC, and the FHWC reduces to $\mathrm{R} > 1$. All results apply equally to finite-horizon models via backward recursion from terminal period T, and to models with permanent income growth by appropriately adjusting the patience conditions above. The generalization to the case with permanent shocks is straightforward.[5]

[4] We assume $\rho \neq 1$ throughout, because the transformations of the value function below ($\bar{\boldsymbol{\Lambda}}$, $\hat{\boldsymbol{\Lambda}}$) divide by $1 - \rho$. Log utility needs parallel derivations, with the same economics.

[5] Restoring permanent shocks (a five-point discretization at standard deviation 0.1) moves every error reported below by less than a factor of two, on that calibration's own endogenous

3. Benchmark: The Method of Endogenous Gridpoints

For comparison to our new solution method, we use the endogenous gridpoints solution to the microeconomic problem presented in Carroll (2006). That method computes the level of consumption at a set of gridpoints for market resources m that are determined endogenously using the Euler equation. The consumption function is then constructed by linear interpolation among the gridpoints thus found. Extensions of this method handle multi-dimensional problems (Barillas and Fernández-Villaverde, 2007), occasionally binding constraints (Hintermaier and Koeniger, 2010), non-smooth and non-concave problems (Fella, 2014), and discrete-continuous choice models (Iskhakov et al., 2017), while White (2015) treats the method's theory and practice. None of them repairs behavior outside the solved grid, which is the problem we take up here. The closest work on that problem takes a different route, since Ma and Toda (2022) justify linearly extrapolating consumption itself at its asymptotic MPC, and Gouin-Bonenfant and Toda (2023) build wealth-tail analysis on the same asymptotics. Extrapolating the level enforces the correct slope only in the limit, whereas the method below extrapolates a bounded transform, which keeps the approximation between both theoretical bounds at every finite wealth level whatever slope the boundary supplies.

Carroll (2020b) describes a specific calibration of the model and constructs a solution using five gridpoints chosen to capture the structure of the consumption function reasonably well at values of m near the infinite-horizon target value (See those notes for details). The five gridpoints are for illustration (applied work typically uses 30-80 gridpoints for accurate solutions), and the comparisons below show what happens to each method when the grid is this sparse.

Throughout the numerical illustrations we adopt this calibration: relative risk aversion $\rho = 2$, discount factor $\beta = 0.96$, and gross risk-free return $\mathrm{R} = 1.02$, with no permanent income growth or shocks ($G = 1$, $\psi \equiv 1$) and mean-one transitory shocks of log standard deviation 1 discretized on a seven-point equiprobable grid (the quadrature schemes of Tauchen (1986) and Tauchen and Hussey (1991) are the standard alternatives). The calibration has no zero-income atom, since the worst draw $\underline{\xi}$ is the smallest point of the (strictly positive) discretized grid, so the general worst-shock case of (5) applies with $\wp$ equal to that point's probability, and the natural borrowing constraint is negative. Unless otherwise noted, the figures and the accuracy comparison below display the next-to-last period $T-1$ of the finite-horizon problem, for which the analytical last-period rule furnishes exact boundary conditions (human wealth one period before the end is $\bar{h}_{T-1} = 1/\mathrm{R} \approx 0.98$, in contrast to the infinite-horizon value of 50), and both methods use linear interpolation, of consumption for endogenous gridpoints and of the logit for moderation. The Hermite refinement is treated separately in Section 6.2.

gridpoints, and leaves the ordering of the two methods intact. The replication package's verification script carries the switch.

The endogenous gridpoints solution method says nothing about how the policy function should be extrapolated beyond the grid at which the problem has been solved, and that silence allows many interpretations. Here we assume the natural choice of linear extension, which leads to badly wrong outcomes. (Other common solution methods are no better outside their own predefined ranges.) Figure 1 demonstrates the point by plotting the amount of precautionary saving implied by a linear extrapolation of our approximated consumption rule (the consumption of the perfect foresight consumer $\bar{\mathbf{c}}_{T-1}$ minus our approximation to optimal consumption under uncertainty, $\grave{\mathbf{c}}_{T-1}$). Although theory proves that precautionary saving is always positive, the linearly extrapolated numerical approximation eventually predicts negative precautionary saving (at the point in the figure where the extrapolated locus crosses the horizontal axis).

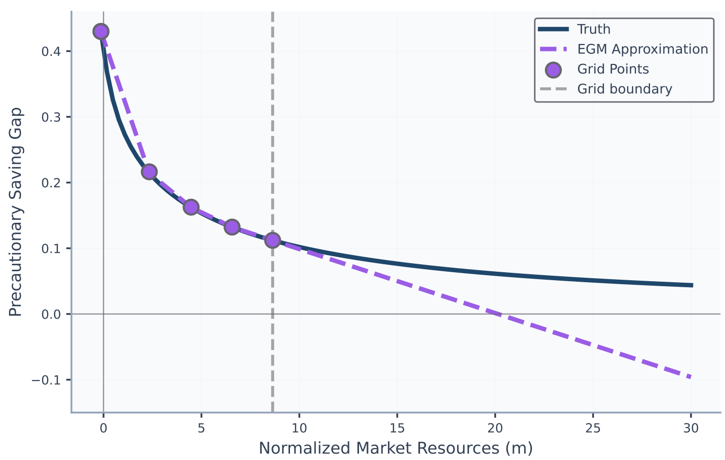

Figure 1: For Large Enough m_{T-1}, Predicted Precautionary Saving is Negative (Oops!)

This error cannot be fixed by extending the upper gridpoint; in the presence of serious uncertainty, the consumption rule will need to be evaluated outside of *any* prespecified grid (because starting from the top gridpoint, a large enough realization of the uncertain variable will push next period's realization of assets above that top; a similar argument applies below the bottom gridpoint). While a judicious extrapolation technique can prevent this problem from being fatal (for example by carefully excluding negative precautionary saving), the problem is often dealt with using inelegant methods whose implications for the accuracy of the solution are difficult to gauge. A fix chosen because it is convenient can restore the sign of precautionary saving while leaving its magnitude ungrounded in the theory, which matters because the convenient fixes are the ones a modeler tries first.[6]

[6] The `HARK` toolkit does not (yet) extend the last linear segment, and instead shrinks the gap above the top gridpoint to the perfect-foresight limit exponentially, which keeps precautionary

4. The Method of Moderation

4.1. The Optimist, the Pessimist, and the Realist

As a preliminary to our solution, define $\bar{h}$ as end-of-period human wealth (the present discounted value of future labor income) for a perfect foresight version of the problem of a 'risk optimist:' a consumer who believes with perfect confidence that the shocks will always take their expected value of 1, $\xi_{t+n} = \mathbb{E}[\xi] = 1$ for all $n > 0$. Under the FHWC this present value converges to $\bar{h} = G/(\mathrm{R} - G)$, which is $1/(\mathrm{R} - 1)$ in our calibration with $G = 1$. The solution to a perfect foresight problem of this kind takes the form[7]

$$\bar{\mathbf{c}}(m) = (m + \bar{h})\underline{\boldsymbol{\kappa}} \tag{8}$$

for a constant minimal marginal propensity to consume, which in the infinite-horizon case is

$$\underline{\boldsymbol{\kappa}} = 1 - \text{Þ}/\mathrm{R}, \tag{9}$$

strictly positive under the RIC.[8] We similarly define $\underline{h}$ as 'minimal human wealth,' the present discounted value of labor income if the shocks were to take on their worst value $\underline{\xi} \geq 0$ (zero for uninsured unemployment, positive when unemployment pays a benefit) in every future period, $\xi_{t+n} = \underline{\xi}$ for all $n > 0$ (which we define as corresponding to the beliefs of a 'pessimist'). The same limit gives $\underline{h} = \underline{\xi}G/(\mathrm{R} - G)$, which is zero when the worst shock is unemployment.

We will call a 'realist' the consumer who correctly perceives the true probabilities of the future risks and optimizes accordingly.

For the realist, a lower bound for the level of market resources is $\underline{m} = -\underline{h}$, because if m equalled this value then there would be a positive finite chance (however small) of receiving $\xi_{t+n} = \underline{\xi}$ in every future period, which would require the consumer to set c to zero to guarantee that the intertemporal budget constraint holds. Since consumption of zero yields infinite marginal utility, the solution to the realist consumer's problem is not well defined for values of $m \leq \underline{m}$ (Zeldes, 1989; Deaton, 1991), and the limiting value of the realist's c is zero as $m \downarrow \underline{m}$.

---

saving positive. That is the easy and obvious choice, the one we made early in the toolkit's development, when the limiting bounds were known but the theory behind the approach to them was not yet worked out. It is also not how the true gap closes, and nothing in the calibration says how it should, so the accuracy it delivers at any wealth level is not knowable in advance. We therefore compare against linear extrapolation, which is the generic practice, whereas the method below takes its tail from the solved gridpoints.

[7]For a derivation, see Carroll and Shanker (2026); $\underline{\boldsymbol{\kappa}}$ is defined therein as the MPC of the perfect foresight consumer with horizon $T - t$.

[8]Each of these constants has a finite-horizon counterpart, obtained by the same backward recursion and summing to the same closed form. With horizon $T - t$: $\bar{h}_T = 0$ and $\bar{h}_t = (G/\mathrm{R})(1 + \bar{h}_{t+1})$, so $\bar{h}_t = \sum_{n=1}^{T-t}(G/\mathrm{R})^n$; $\underline{h}$ obeys the same recursion with $\underline{\xi}$ in place of the unit income draw, giving $\underline{h}_t = \underline{\xi}\sum_{n=1}^{T-t}(G/\mathrm{R})^n$; and $\underline{\boldsymbol{\kappa}}_t = \underline{\boldsymbol{\kappa}}_{t+1}/(\underline{\boldsymbol{\kappa}}_{t+1} + \text{Þ}/\mathrm{R})$ with $\underline{\boldsymbol{\kappa}}_T = 1$, so $\underline{\boldsymbol{\kappa}}_t = (\sum_{n=0}^{T-t}(\text{Þ}/\mathrm{R})^n)^{-1}$, converging to (9).

Given this result, it will be convenient to define 'excess' market resources as the amount by which actual resources exceed the lower bound, and 'excess' human wealth as the amount by which mean expected human wealth exceeds guaranteed minimum human wealth:[9]

$$
\begin{aligned}
\Delta m &= m + \overbrace{\underline{h}}^{=-\underline{m}} \\
\Delta h &= \bar{h} - \underline{h}.
\end{aligned}
\tag{10}
$$

We can now transparently define the optimal consumption rules for the two perfect foresight problems, those of the 'optimist' and the 'pessimist.' The 'pessimist' perceives human wealth to be equal to its minimum feasible value $\underline{h}$ with certainty, so consumption is given by the perfect foresight solution

$$
\begin{aligned}
\underline{\mathbf{c}}(m) &= (m + \underline{h})\underline{\boldsymbol{\kappa}} \\
&= \Delta m \underline{\boldsymbol{\kappa}}.
\end{aligned}
\tag{11}
$$

The 'optimist,' on the other hand, pretends that there is no uncertainty about future income, and therefore consumes

$$
\begin{aligned}
\bar{\mathbf{c}}(m) &= (m + \underline{h} - \underline{h} + \bar{h})\underline{\boldsymbol{\kappa}} \\
&= (\Delta m + \Delta h)\underline{\boldsymbol{\kappa}} \\
&= \underline{\mathbf{c}}(m) + \Delta h \underline{\boldsymbol{\kappa}}.
\end{aligned}
\tag{12}
$$

4.2. The Consumption Function

It seems obvious that the spending of the realist will be strictly greater than that of the pessimist and strictly less than that of the optimist. The realist reoptimizes as uncertainty resolves, while the pessimist must finance every future period out of saving, and because the adverse outcome remains possible, even a wealthy realist is never completely self-insured. Figure 2 illustrates the proposition for the consumption rule in period $T-1$.

The proof is more difficult than might be imagined, but the necessary work is done in Carroll and Shanker (2026): under shocks with nondegenerate bounded support (whose worst realization $\underline{\xi}$ may be zero), the consumption function is strictly increasing and concave (Carroll and Kimball, 1996) and lies strictly between the two perfect foresight rules. We therefore take the proposition as a fact and proceed by manipulating the inequality:

$$
\underline{\mathbf{c}}(\underline{m} + \Delta m) < \hat{\mathbf{c}}(\underline{m} + \Delta m) < \bar{\mathbf{c}}(\underline{m} + \Delta m). \tag{13}
$$

Subtracting $\underline{\mathbf{c}}(\underline{m} + \Delta m)$ in each of these inequalities and using Equations (11) and (12) gives

[9] Here Δ denotes excess above minimum, not a time difference.

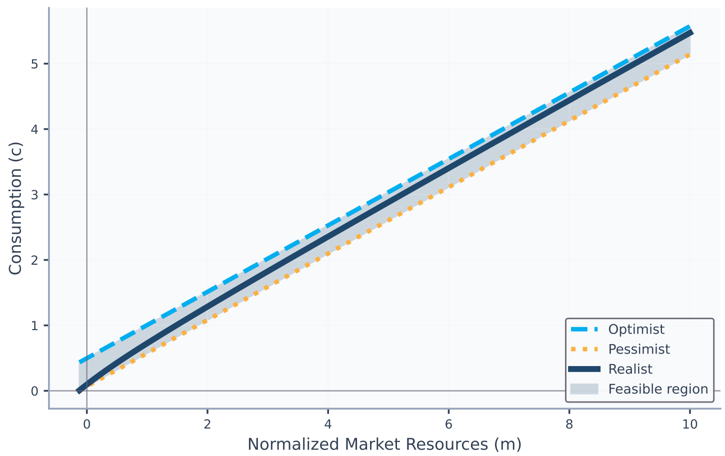

Figure 2: Moderation Illustrated: $\underline{\mathbf{c}} < \hat{\mathbf{c}} < \bar{\mathbf{c}}$

$$\begin{aligned} 0 <& \quad \hat{\mathbf{c}}(\underline{m} + \Delta m) - \underline{\mathbf{c}}(\underline{m} + \Delta m) && < \Delta h \underline{\boldsymbol{\kappa}} \\ 0 <& \quad \underbrace{\left(\frac{\hat{\mathbf{c}}(\underline{m} + \Delta m) - \underline{\mathbf{c}}(\underline{m} + \Delta m)}{\Delta h \underline{\boldsymbol{\kappa}}} \right)}_{\equiv \boldsymbol{\omega}} && < 1, \end{aligned} \tag{14}$$

where the fraction in the middle of the last inequality is the moderation ratio measuring how close the realist's consumption is to the optimist's behavior (the numerator is the gap between the realist and pessimist) relative to the maximum possible gap between optimist and pessimist. When $\boldsymbol{\omega} = 0$, the realist behaves like the pessimist (maximum precautionary saving); when $\boldsymbol{\omega} = 1$, the realist behaves like the optimist (no precautionary saving). Under bounded shocks neither endpoint is attained, so $\boldsymbol{\omega}$ lies strictly inside $(0, 1)$ for all $m > \underline{m}$, and the equivalent form $\hat{\mathbf{c}} = \underline{\mathbf{c}} + \boldsymbol{\omega} \Delta h \underline{\boldsymbol{\kappa}}$ places the realist strictly between the pessimist and the optimist.

Figure 3 shows why we work in this quantity rather than in consumption itself. The bracket has a constant width of $\underline{\boldsymbol{\kappa}} \Delta h$, so in levels it shrinks against the rule it brackets, from more than half of the optimist's consumption near the constraint to under a tenth by $m = 10$, which is why the three agents of Figure 2 crowd onto one line. The same information rescaled by that width fills its range, as $\boldsymbol{\omega}$ climbs from zero at the constraint through 0.5 near $m = 2.3$ and past 0.8 by $m = 15$, with the five gridpoints spread across the climb rather than bunched.

Defining $\boldsymbol{\mu} = \log \Delta m$ (which can range from $-\infty$ to ∞), the object in the middle of the last inequality is

$$\boldsymbol{\omega}(\boldsymbol{\mu}) \equiv \left(\frac{\hat{\mathbf{c}}(\underline{m} + e^{\boldsymbol{\mu}}) - \underline{\mathbf{c}}(\underline{m} + e^{\boldsymbol{\mu}})}{\Delta h \underline{\boldsymbol{\kappa}}} \right), \tag{15}$$

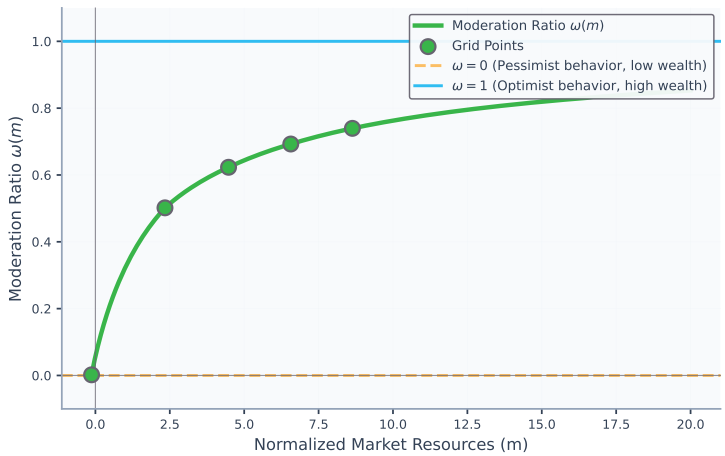

Figure 3: Rescaling by the Bracket Width Turns a Vanishing Gap Into a Full-Range Object

and we now define

$$\begin{aligned}\boldsymbol{\chi}(\boldsymbol{\mu}) &= \log\left(\frac{\boldsymbol{\omega}(\boldsymbol{\mu})}{1-\boldsymbol{\omega}(\boldsymbol{\mu})}\right) \\ &= \log(\boldsymbol{\omega}(\boldsymbol{\mu})) - \log(1-\boldsymbol{\omega}(\boldsymbol{\mu}))\end{aligned} \tag{16}$$

which carries the bounded moderation ratio to the whole real line, just as the log carried $\Delta m \in (0,\infty)$ to $\boldsymbol{\mu} \in (-\infty,\infty)$. The inverse is the sigmoid $\boldsymbol{\omega} = 1/(1+\exp(-\boldsymbol{\chi}))$, with $\boldsymbol{\chi} \to +\infty$ as the realist approaches the optimist and $\boldsymbol{\chi} \to -\infty$ as the realist approaches the pessimist.

The transformed function suits linear interpolation because it straightens at both ends of its domain. As $\boldsymbol{\mu} \to -\infty$ the realist's MPC approaches a finite limit above $\underline{\boldsymbol{\kappa}}$ (the bound $\bar{\boldsymbol{\kappa}}$ of Section 6.1), hence $\boldsymbol{\omega}$ vanishes in proportion to Δm and $\boldsymbol{\chi}(\boldsymbol{\mu})$ approaches a line with slope one. As $\boldsymbol{\mu} \to +\infty$, $\boldsymbol{\chi}$ is *asymptotically linear*,[10] so we extrapolate it linearly using the slope at the upper boundary gridpoint. Because any linear extrapolation of the logit maps back through the sigmoid into $\boldsymbol{\omega} \in (0,1)$, the approximation satisfies $\underline{\mathbf{c}} < \grave{\mathbf{c}} < \bar{\mathbf{c}}$ throughout the extrapolation domain. Equation (16) splits the logit into two pieces, $\log \boldsymbol{\omega}$ and $-\log(1-\boldsymbol{\omega})$, and either piece alone would enforce only one of the two bounds.

Each piece straightens the tail at which its own bound is approached, $-\log(1-\boldsymbol{\omega})$ where $\boldsymbol{\omega} \to 1$ and $\log \boldsymbol{\omega}$ where $\boldsymbol{\omega} \to 0$. Here both bounds are approached only asymptotically, the pessimist's as $\boldsymbol{\mu} \to -\infty$ and the optimist's as $\boldsymbol{\mu} \to +\infty$, so both singularities of the logit lie at the ends of the domain. The logit needs bounds that are approached and never reached, though it does not care from

[10]Under the GIC the limiting slope $\alpha = \lim_{\boldsymbol{\mu}\to+\infty} \frac{\partial \boldsymbol{\chi}}{\partial \boldsymbol{\mu}} \geq 0$ exists. It may equal zero in theory, but the boundary slope on any finite grid is strictly positive, and the bounds are preserved whichever slope is used.

which side, since reorienting the ratio only changes the sign of the transform, $\text{logit}(1-\boldsymbol{\omega}) = -\text{logit}(\boldsymbol{\omega})$. Where a bound is actually reached at some finite level of wealth, the piece belonging to it diverges at the edge of the working range, and instead of straightening the approach it bends it, so the right transform there is the one-sided piece for the bound that is only approached. The portfolio share below presents exactly that case.

Figure 4 plots the object the method actually interpolates: $\boldsymbol{\chi}(\boldsymbol{\mu})$ for the five-gridpoint solution of the calibration above. The function is smooth, gently curved in the middle of the grid, and close to linear at both ends.

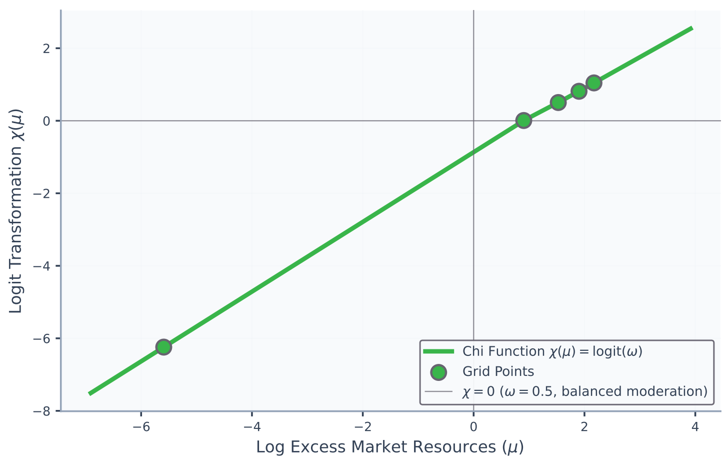

Figure 4: The Interpolated Object $\boldsymbol{\chi}(\boldsymbol{\mu})$ Is Smooth and Close to Linear at Both Ends

Given $\boldsymbol{\chi}$, the consumption function can be recovered from

$$\hat{\mathbf{c}} = \underline{\mathbf{c}} + \overbrace{\frac{1}{1+\exp(-\boldsymbol{\chi})}}^{=\boldsymbol{\omega}} \Delta h \underline{\boldsymbol{\kappa}}. \tag{17}$$

Thus, the procedure is to calculate $\boldsymbol{\chi}$ at the points $\boldsymbol{\mu}$ corresponding to the log of the Δm points defined above, and then using these to construct an interpolating approximation $\grave{\boldsymbol{\chi}}$ from which we indirectly obtain our approximated consumption rule $\grave{\mathbf{c}}$ (an approximation to the true $\hat{\mathbf{c}}$) by substituting $\grave{\boldsymbol{\chi}}$ for $\boldsymbol{\chi}$ in Equation (17).

Because this method relies upon the fact that the problem is easy to solve if the decision maker has unreasonable views (either in the optimistic or the pessimistic direction), and because the correct solution is always between these immoderate extremes, we call our solution procedure the 'method of moderation.'

Figure 5 is visually indistinguishable from the true rule; a reader with very good eyesight might detect the barest hint of a discrepancy between the Truth and the Approximation at the far right-hand edge of the figure, a stark contrast with the calamitous divergence evident in Figure 1.

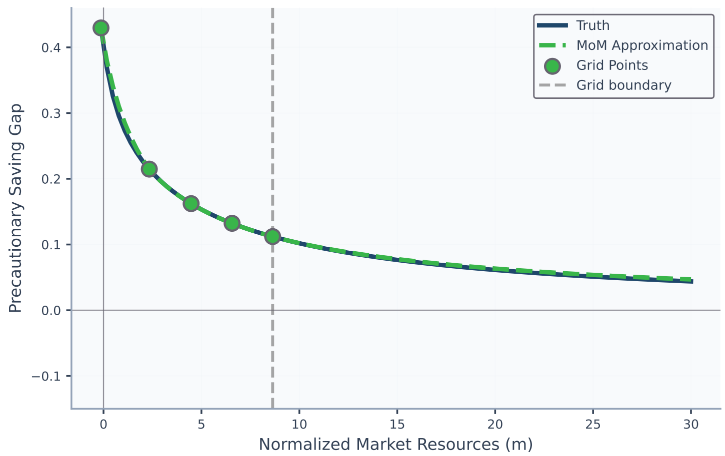

Figure 5: Extrapolated $\grave{\mathbf{c}}_{T-1}$ Stays on the True Rule Where Linear EGM Extrapolation Diverged

4.3. The Value Function

Often it is useful to know the value function as well as the consumption rule. Fortunately, many of the tricks used when solving for the consumption rule have a direct analogue in approximation of the value function.

Consider the perfect foresight (or "optimist's") problem in period $T-1$. Using the fact that in a perfect foresight model the growth factor for consumption is constant, we can use $c_t = \text{Þ} \cdot c_{t-1}$ to calculate the value function in period $T-1$:

$$\begin{aligned} \bar{\mathbf{v}}_{T-1}(m_{T-1}) &\equiv \mathbf{u}(c_{T-1}) + \beta \mathbf{u}(c_T) \\ &= \mathbf{u}(c_{T-1}) \left(1 + \beta \text{Þ}^{1-\rho}\right) \\ &= \mathbf{u}(c_{T-1}) \left(1 + \text{Þ}/\text{R}\right) \\ &= \mathbf{u}(c_{T-1}) \underbrace{\text{PDV}_{T-1}^{T}(c)/c_{T-1}}_{\equiv \mathcal{C}_{T-1}^{T}} \end{aligned} \tag{18}$$

where $\mathcal{C}_t^T = \text{PDV}_t^T(c)/c_t$ is the present discounted value of consumption, normalized by current consumption. Using the fact demonstrated in Carroll and Shanker (2026) that $\mathcal{C}_t^T = \underline{\boldsymbol{\kappa}}_t^{-1}$[11], a similar function can be constructed recursively for earlier periods, yielding the general expression

[11] Under perfect foresight, consumption grows at rate Þ: $\mathbf{c}_{t+n} = \mathbf{c}_t \text{Þ}^n$. Discounting yields $(\text{PDV}_t^T(\mathbf{c})/\mathbf{c}_t) = \sum_{n=0}^{T-t} (\text{Þ}/\text{R})^n = \underline{\boldsymbol{\kappa}}_t^{-1}$, so $\mathcal{C}_t^T = \underline{\boldsymbol{\kappa}}_t^{-1}$ (unchanged for normalized variables). In the infinite-horizon limit, we write simply $\mathcal{C} = \underline{\boldsymbol{\kappa}}^{-1}$.

$$
\begin{aligned}
\bar{\mathbf{v}}(m) &= \mathbf{u}(\bar{\mathbf{c}}(m))\mathcal{C} \\
&= \mathbf{u}(\bar{\mathbf{c}}(m))\underline{\boldsymbol{\kappa}}^{-1} \\
&= \mathbf{u}((\Delta m + \Delta h)\underline{\boldsymbol{\kappa}})\underline{\boldsymbol{\kappa}}^{-1} \\
&= \left[(\Delta m + \Delta h)^{1-\rho}/(1-\rho)\right] \cdot \left[\underline{\boldsymbol{\kappa}}^{1-\rho} \cdot \underline{\boldsymbol{\kappa}}^{-1}\right] \\
&= \mathbf{u}(\Delta m + \Delta h)\underline{\boldsymbol{\kappa}}^{-\rho}.
\end{aligned} \tag{19}
$$

This can be transformed as

$$
\begin{aligned}
\bar{\boldsymbol{\Lambda}} &\equiv ((1-\rho)\bar{\mathbf{v}})^{1/(1-\rho)} \\
&= c\,\mathcal{C}^{1/(1-\rho)} \\
&= (\Delta m + \Delta h)\underline{\boldsymbol{\kappa}}^{-\rho/(1-\rho)}.
\end{aligned} \tag{20}
$$

The transformation $\boldsymbol{\Lambda} \equiv ((1-\rho)\mathbf{v})^{1/(1-\rho)}$ applies the inverse utility function to value, $\boldsymbol{\Lambda} = \mathbf{u}^{-1}(\mathbf{v})$, expressing value in consumption units. Despite the customary name 'inverse value function,' it is this composition, not the functional inverse of $\mathbf{v}$ itself. The pessimist's inverse value follows by the same steps, with $\underline{\mathbf{c}} = \Delta m\underline{\boldsymbol{\kappa}}$:

$$
\underline{\boldsymbol{\Lambda}} = \underline{\mathbf{c}}\,\mathcal{C}^{1/(1-\rho)} = \Delta m\,\underline{\boldsymbol{\kappa}}^{-\rho/(1-\rho)}, \tag{21}
$$

hence $\bar{\boldsymbol{\Lambda}} - \underline{\boldsymbol{\Lambda}} = \Delta h\,\underline{\boldsymbol{\kappa}}\,\mathcal{C}^{1/(1-\rho)}$, the denominator that normalizes the value moderation ratio in (23) below.

We apply the same transformation to the value function for the problem with uncertainty (the "realist's" problem):

$$
\hat{\boldsymbol{\Lambda}} = ((1-\rho)\hat{\mathbf{v}}(m))^{1/(1-\rho)} \tag{22}
$$

and the value function can be approximated by calculating the values of $\hat{\boldsymbol{\Lambda}}$ at the same gridpoints used by the consumption function approximation, and interpolating among those points.

However, as with the consumption approximation, we can do even better if we realize that the $\bar{\boldsymbol{\Lambda}}$ function for the optimist's problem is an upper bound for the $\boldsymbol{\Lambda}$ function in the presence of uncertainty, and the value function for the pessimist is a lower bound. The ordering $\underline{\mathbf{v}} < \hat{\mathbf{v}} < \bar{\mathbf{v}}$ holds because the realist's true income process stochastically dominates the pessimist's worst-case income and is dominated, for a risk-averse agent, by the optimist's certain expected income. Because the inverse-value transform is monotonic, the same ordering carries over to $\underline{\boldsymbol{\Lambda}} < \hat{\boldsymbol{\Lambda}} < \bar{\boldsymbol{\Lambda}}$. Analogously to (15), define an upper-case

$$
\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu}) = \left(\frac{\hat{\boldsymbol{\Lambda}}(\underline{m} + e^{\boldsymbol{\mu}}) - \underline{\boldsymbol{\Lambda}}(\underline{m} + e^{\boldsymbol{\mu}})}{\Delta h\underline{\boldsymbol{\kappa}}\,\mathcal{C}^{1/(1-\rho)}}\right) \tag{23}
$$

and an upper-case version of the $\boldsymbol{\chi}$ equation in (16):

$$
\begin{aligned}
\hat{\boldsymbol{X}}(\boldsymbol{\mu}) &= \log\left(\frac{\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})}{1 - \hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})}\right) \\
&= \log(\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})) - \log(1 - \hat{\boldsymbol{\Omega}}(\boldsymbol{\mu}))
\end{aligned} \tag{24}
$$

and if we approximate these objects then invert them (as above with the $\boldsymbol{\omega}$ and $\boldsymbol{\chi}$ functions) we obtain an approximation to the inverted value function that respects both bounds by construction, at the same points for which we have our approximated value function:

$$\hat{\mathbf{\Lambda}} = \underline{\mathbf{\Lambda}} + \overbrace{\left(\frac{1}{1+\exp(-\hat{\boldsymbol{X}})}\right)}^{=\hat{\mathbf{\Omega}}} \Delta h \underline{\boldsymbol{\kappa}}\, \mathcal{C}^{1/(1-\rho)} \tag{25}$$

from which we obtain our approximation to the value function as

$$\begin{aligned} \hat{\mathbf{v}} &= \mathbf{u}(\hat{\mathbf{\Lambda}}) \\ \hat{\mathbf{v}}' &= \mathbf{u}'(\hat{\mathbf{\Lambda}})\, \hat{\mathbf{\Lambda}}'. \end{aligned} \tag{26}$$

5. Numerical Accuracy

The figures show the failure and the repair but not the size of the gain, so we compare the method of moderation against the benchmark endogenous gridpoints method on the calibration above, solving both on the same sparse five-point grid. We report in Table 1 the maximum absolute consumption error on a dense subgrid of each interval between adjacent gridpoints m_j, m_{j+1}, and of the extrapolation region running out to $\overline{m} = 30$, taking a high-precision endogenous gridpoints solution (500 gridpoints) as the truth. The moderation column applies the near-constraint envelope of Section 6.1, which is slack over the tabulated range under the natural constraint used here. Except in the first interval, where the advantage is about three and a half times, the method of moderation is more than an order of magnitude more accurate than EGM, and we obtain this on a fixed sparse grid without optimizing the grid itself (Chipeniuk, 2020). The contrast is starkest in the extrapolation region beyond the top gridpoint, exactly where Figure 1 showed linear EGM extrapolation failing, and there the method of moderation is roughly forty-six times more accurate. That ratio is a property of the horizon $\overline{m} = 30$ as much as of the two methods, since EGM's linear extrapolation diverges and any ratio one likes can be had by evaluating far enough out. The fairer comparison, against an EGM rule handed the same bounds, is seventeen, and we come to it below.[12]

Figure 6 shows that the table's intervals are error humps, with both approximations nearly exact at the sparse gridpoints and less accurate between them, and the method of moderation's humps sit below EGM's throughout. The two methods part company beyond the top gridpoint (marked), where the EGM error rises steeply as its linear extrapolation drifts away from the truth, whereas moderation stays flat.

[12]The replication package regenerates the tables and the numbers quoted in this section and checks each accuracy claim against the code.

Table 1: Maximum approximation errors by interval for endogenous gridpoints (EGM) and the method of moderation (MoM) on the same five-point grid with linear interpolation: absolute consumption errors against the reference solution (upper panel) and the largest Euler-equation residual in consumption-equivalent units, which needs no reference solution (lower panel). Gridpoints $m_0 = -0.13$, $m_1 = 2.34$, $m_2 = 4.47$, $m_3 = 6.57$, $m_4 = 8.64$; evaluation horizon $\overline{m} = 30$. Orders of magnitude in parentheses.

Method	$[m_0, m_1]$	$[m_1, m_2]$	$[m_2, m_3]$	$[m_3, m_4]$	$[m_4, \overline{m}]$
Absolute error					
EGM	5.4(-2)	4.2(-3)	1.6(-3)	8.6(-4)	1.4(-1)
MoM	1.5(-2)	2.5(-4)	1.4(-4)	7.4(-5)	3.0(-3)
Euler residual					
EGM	7.7(-1)	4.6(-3)	1.1(-3)	4.2(-4)	1.8(-2)
MoM	1.7(-1)	2.7(-4)	9.8(-5)	3.6(-5)	4.1(-4)

An absolute error is hard to interpret near the borrowing constraint, where consumption itself is close to zero, so in the lower panel of the table we score the same two rules by a yardstick that depends neither on the scale of consumption nor on our reference solution: the Euler-equation residual, the proportion by which consumption would have to change for the Euler equation to hold exactly (Santos, 2000). By the residual, as by the absolute errors, the method of moderation is more accurate in every interval, by a factor of four and a half in the first, eleven to seventeen in the interior, and forty-four beyond the top gridpoint. Both methods do poorly in the first interval by this yardstick, since near the constraint even a small absolute error is a large proportion of consumption.

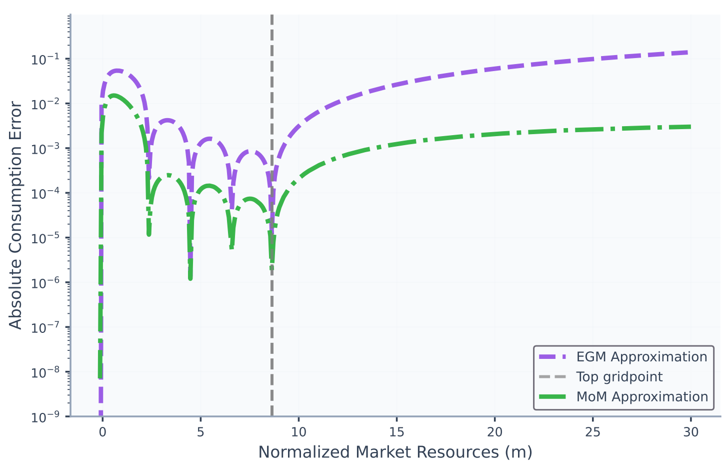

Figure 6: Moderation's Error Humps Sit Below EGM's Between Gridpoints and Stay Flat Beyond the Top Gridpoint

The two halves of the figure have different causes. Inside the grid we gain because χ is the smoother object, and beyond it because linear extrapolation of χ cannot leave the bracket, while linear EGM extrapolation eventually does.

Grid design governs beyond-grid accuracy but never the bounds, and what matters is how far the grid reaches rather than how many points it holds (Figure 7). Raising the count from five to thirty without moving the top gridpoint does not improve the beyond-grid error at all, while moving the top out to twice human wealth improves it by roughly an order of magnitude, and to four times human wealth by a further factor of two; the two spacings the figure plots differ little by comparison. In practice we would therefore pair the 30-80 gridpoints noted earlier with a top gridpoint at a small multiple of human wealth, which the exhibits above already satisfy, human wealth one period before the end being close to one.

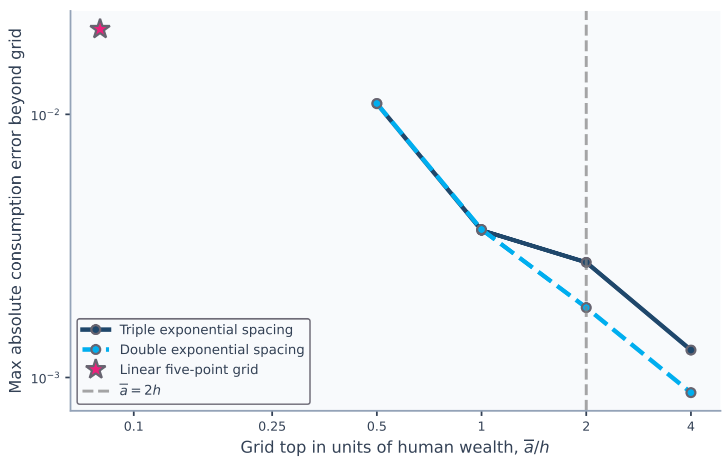

Figure 7: Extending the Grid Top Improves Beyond-Grid Accuracy, While the Bounds Hold Everywhere

A skeptical reader will want to know how much of the comparison of Table 1 is the design rather than the method, so we ran two further comparisons. First, on a grid of the size used in practice (twenty points, triple-exponential spacing, same top gridpoint) both methods gain roughly two orders of magnitude of interior accuracy, compressing moderation's worst-case interior advantage to about four times (the median interval ratio is eleven), while beyond the grid the ratio is untouched by refinement, forty-six at five points and at twenty. Second, the extrapolation advantage does not come from denying EGM the bounds, for when we impose the same optimist and pessimist bounds on EGM's rule every interior interval is unchanged (neither bound binds inside the grid) and its beyond-grid error falls only from 1.4(-1) to 6.1(-2), against 3.0(-3) for moderation. Nor does that comparison depend on the evaluation horizon $\overline{m} = 30$. The error of EGM without the bounds grows without limit as the horizon extends, but the errors of the bounded EGM rule and of moderation both level off, at 6.1(-2) and 3.6(-3), so their ratio converges to seventeen however far the comparison runs. Imposing the bounds also concedes the premise, since the bounds EGM borrows are exactly what the method of moderation is built on, and the bounded rule is kinked wherever a bound binds, whereas moderation approaches its bounds smoothly by construction.

6. Extensions

6.1. A Tighter Upper Bound

Carroll and Shanker (2026) derives an upper limit $\bar{\boldsymbol{\kappa}}$ for the MPC as m approaches its lower bound. Recalling that $\wp$ is the probability of the worst income draw, $\wp = \Pr(\xi = \underline{\xi})$ (one seventh in the calibration above, where the worst draw is one point of a seven-point equiprobable grid), the infinite-horizon limit is

$$\bar{\boldsymbol{\kappa}} = 1 - \wp^{1/\rho}(\text{Þ}/\text{R}), \tag{27}$$

which nests the Friedman-Muth unemployment case $\underline{\xi} = 0$ of Carroll and Toche (2009).[13] Using this fact plus the strict concavity of the consumption function yields the proposition that

$$\hat{\mathbf{c}}(\underline{m} + \Delta m) < \bar{\boldsymbol{\kappa}}\Delta m. \tag{28}$$

Near the constraint the optimist's bound is loose, because it is calibrated to the low MPC that prevails at high wealth. Moderating between the two perfect-foresight rules alone therefore does not guarantee that approximated consumption will respect this tighter constraint between gridpoints, and a failure to respect it can occasionally cause computational problems in solving or simulating the model. Here, we describe a method for constructing an approximation that always satisfies the constraint. It is applied throughout, so every consumption number reported in this paper respects the tighter bound, except where we drop the envelope on purpose to show what it does.

Defining m^* as the *cusp* point where the two upper bounds intersect (where $\Delta m^* \equiv m^* - \underline{m}$):

$$\begin{aligned} \left(\Delta m^* + \Delta h\right) \underline{\boldsymbol{\kappa}} &= \bar{\boldsymbol{\kappa}}\, \Delta m^* \\ \Delta m^* &= \frac{\underline{\boldsymbol{\kappa}}\, \Delta h}{\bar{\boldsymbol{\kappa}} - \underline{\boldsymbol{\kappa}}} \\ m^* &= -\underline{h} + \frac{\underline{\boldsymbol{\kappa}}\left(\bar{h} - \underline{h}\right)}{\bar{\boldsymbol{\kappa}} - \underline{\boldsymbol{\kappa}}}, \end{aligned} \tag{29}$$

this intersection occurs in the feasible region ($m^* > \underline{m}$), because $\bar{\boldsymbol{\kappa}} > \underline{\boldsymbol{\kappa}}$ under the stated conditions and so $\Delta m^* > 0$.

[13] With horizon $T - t$: backward recursion $\bar{\boldsymbol{\kappa}}_t = \bar{\boldsymbol{\kappa}}_{t+1}/(\bar{\boldsymbol{\kappa}}_{t+1} + \wp^{1/\rho}(\text{Þ}/\text{R}))$ with $\bar{\boldsymbol{\kappa}}_T = 1$; equivalently the forward sum $\bar{\boldsymbol{\kappa}}_t = \left(\sum_{n=0}^{T-t}\left(\wp^{1/\rho}(\text{Þ}/\text{R})\right)^n\right)^{-1}$, which converges to the closed form in (27).

The two bounds therefore change places at m^*, the tighter bound binding below it and the optimist's above it. Since the realist's rule respects both everywhere, it respects whichever is smaller,

$$\hat{\mathbf{c}}(m) \leq \min\big(\underline{\boldsymbol{\kappa}}(\Delta m + \Delta h),\ \bar{\boldsymbol{\kappa}}\,\Delta m\big), \tag{30}$$

and we build the approximation to respect it too, by taking the same lower envelope of the moderated rule and the tighter bound. This just says that at each level of wealth we use whichever of the two is the more appropriate, without needing to know in advance which that will be.

Doing so cannot make the approximation worse, because where the moderated rule already lies below the tighter bound, the envelope leaves it alone. If instead the rule lies above the bound, the truth lies below it, so the bound is closer than the moderated value was. At wealth levels where the constraint actually binds, the envelope is not an approximation at all, because there the realist consumes $\bar{\boldsymbol{\kappa}}\Delta m$ exactly.[14]

This is the lower-envelope construction Carroll and Shanker (2026) use for a liquidity constraint, with the constrained branch in closed form joined to the unconstrained branch. The bound imposed here is $\bar{\boldsymbol{\kappa}}\Delta m$, which for $\bar{\boldsymbol{\kappa}} < 1$ lies strictly inside the budget line Δm.

We now construct an upper-bound value function, denoted $\grave{\check{\mathbf{v}}}$ to distinguish it from the optimist's own value $\bar{\mathbf{v}}$, implied for a consumer whose spending behavior is consistent with the refined upper-bound consumption rule.

For $m \geq m^*$, the refined consumption rule coincides with the optimist's, so there we set $\grave{\check{\mathbf{v}}} = \bar{\mathbf{v}}$. However, for values $m < m^*$, matters are slightly more complicated.

Start with the fact that at the cusp point,

$$\begin{aligned} \bar{\mathbf{v}}(m^*) &= \mathbf{u}(\bar{\mathbf{c}}(m^*))\mathcal{C} \\ &= \mathbf{u}(\Delta m^* \bar{\boldsymbol{\kappa}})\mathcal{C}. \end{aligned} \tag{31}$$

For every m, value equals current utility plus a continuation term,

$$\bar{\mathbf{v}}(m) = \mathbf{u}(\bar{\mathbf{c}}(m)) + \mathbf{w}(m - \bar{\mathbf{c}}(m)), \tag{32}$$

where $\mathbf{w}$ is the end-of-period value as a function of the assets retained at the end of the period. Below the cusp point the refined rule spends $\bar{\boldsymbol{\kappa}}\Delta m$, which leaves assets of $m - \bar{\boldsymbol{\kappa}}\Delta m = \underline{m} + (1 - \bar{\boldsymbol{\kappa}})\Delta m$, so for $m < m^*$

$$\grave{\check{\mathbf{v}}}(m) = \mathbf{u}(\bar{\boldsymbol{\kappa}}\Delta m) + \mathbf{w}(\underline{m} + (1 - \bar{\boldsymbol{\kappa}})\Delta m), \tag{33}$$

[14] The envelope also fixes behavior at the constraint itself, since $\bar{\boldsymbol{\kappa}}\,\Delta m$ vanishes at $m = \underline{m}$. Under an artificial constraint $a \geq 0$, EGM cannot place a gridpoint below the kink, because inverting the Euler equation from an asset grid gives $m = a + \hat{\mathbf{c}}(a)$ with $a \geq 0$, so the smallest m it reaches is $\hat{\mathbf{c}}(0)$, the kink itself, and no refinement of the grid creates gridpoints in a region that has no interior solution to invert.

which can be computed directly because $\mathbf{w}(a_t) = \beta \bar{\mathbf{v}}_{t+1}(a_t \mathcal{R}_{t+1} + 1)$, where $\mathcal{R}_{t+1} \equiv \mathrm{R}/\mathcal{G}_{t+1}$ is the growth-normalized gross return and $\bar{\mathbf{v}}$ is as defined above, because a consumer who ends the current period with assets exceeding the lower bound will not expect to be constrained next period. (Recall again that we are merely constructing an object that is guaranteed to be an *upper bound* for the value that the 'realist' consumer will experience. When the worst draw is unemployment, $\underline{m} = 0$ and the retained-asset expression reduces to $(1-\bar{\boldsymbol{\kappa}})\Delta m$.) At the gridpoints defined by the solution of the consumption problem we can then construct

$$\grave{\check{\mathbf{\Lambda}}}(m) = ((1-\rho)\grave{\check{\mathbf{v}}}(m))^{1/(1-\rho)} \tag{34}$$

which yields the appropriate vector for constructing $\grave{\boldsymbol{X}}$ and $\grave{\boldsymbol{\Omega}}$. The rest of the procedure mirrors the one performed for the consumption rule, and delivers a value-function approximation that respects the tighter bound below the cusp and the optimist's bound above it.

Figure 8 shows why the refinement is worth the trouble. To the left of the cusp the optimist's rule permits spending the realist would never choose, and only the tighter bound keeps the approximation out of that region.

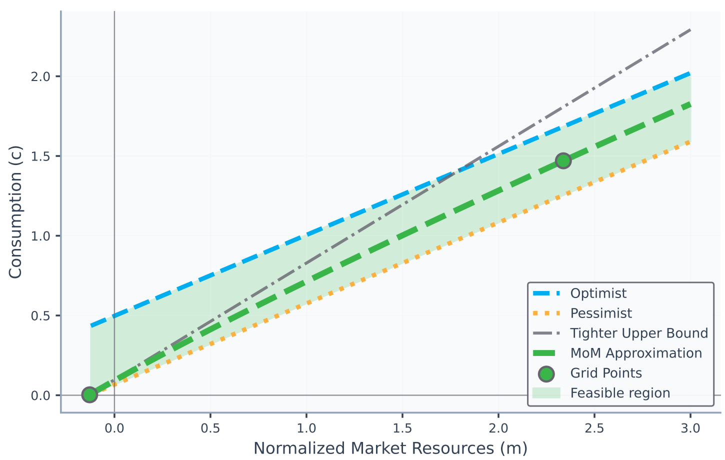

Figure 8: The Tighter Bound $\bar{\boldsymbol{\kappa}}\,\Delta m$ Binds Below the Cusp, the Optimist's Bound Above It

6.2. Hermite Interpolation

Although linear interpolation that matches the level of $\hat{\mathbf{c}}$ at the gridpoints is simple, Hermite interpolation (Fritsch and Carlson, 1980; Fritsch and Butland, 1984; Hyman, 1983) matches both the level and the derivative of $\hat{\mathbf{c}}$ there. That derivative is supplied by the envelope condition (Benveniste and Scheinkman, 1979; Milgrom and Segal, 2002) together with the EGM Euler equation, so the interpolated rule matches, at each solved gridpoint, the marginal propensity to consume that the solution implies as well as the level.[15]

The slopes we need come from the moderation ratio $\boldsymbol{\omega}$ itself. Since its argument $\boldsymbol{\mu}$ is the log of excess market resources, the derivative measures how quickly the realist approaches the optimist as resources increase:

$$\frac{\partial \boldsymbol{\omega}}{\partial \boldsymbol{\mu}} = \frac{\Delta m(\partial \hat{\mathbf{c}}/\partial m - \underline{\boldsymbol{\kappa}})}{\underline{\boldsymbol{\kappa}}\Delta h}. \tag{35}$$

The interpolant carries the logit of (16) rather than $\boldsymbol{\omega}$ itself, and the chain rule gives its derivative:

$$\frac{\partial \boldsymbol{\chi}}{\partial \boldsymbol{\mu}} = \frac{\partial \boldsymbol{\omega}/\partial \boldsymbol{\mu}}{\boldsymbol{\omega}(1-\boldsymbol{\omega})}. \tag{36}$$

It turns out that this is all the slope data a cubic Hermite interpolant needs.[16] Differentiating (17) yields a moderation form for the MPC:[17]

$$\frac{\partial \hat{\mathbf{c}}}{\partial m} = (1-\boldsymbol{\eta})\,\underline{\boldsymbol{\kappa}} + \boldsymbol{\eta}\,\bar{\boldsymbol{\kappa}} \tag{37}$$

where

$$\boldsymbol{\eta} = \frac{\underline{\boldsymbol{\kappa}}}{\bar{\boldsymbol{\kappa}} - \underline{\boldsymbol{\kappa}}} \cdot \frac{\Delta h}{\Delta m} \cdot \partial \boldsymbol{\omega}/\partial \boldsymbol{\mu}. \tag{38}$$

Read the other way, this says we need no second approximation for the marginal propensity to consume. It is the derivative of the object we have already approximated, rescaled so that $\boldsymbol{\eta}$ places the slope between $\underline{\boldsymbol{\kappa}}$ and $\bar{\boldsymbol{\kappa}}$ exactly as $\boldsymbol{\omega}$ places the level between the pessimist and the optimist. Matching slopes as well as levels at each gridpoint therefore gives a consumption rule whose implied MPC agrees with the solved one at every node.

The pairing also explains why Figure 9 falls where Figure 3 climbs. The extra factor $\Delta h/\Delta m$ keeps the two from being mirror images. As $\boldsymbol{\omega}$ levels off near one its derivative falls toward zero, and the factor $\Delta h/\Delta m$ falls with it.

At gridpoints the weight $\boldsymbol{\eta}$ lies in $[0, 1]$ by construction. Equation (37) rearranges to $\boldsymbol{\eta} = (\partial \hat{\mathbf{c}}/\partial m - \underline{\boldsymbol{\kappa}})/(\bar{\boldsymbol{\kappa}} - \underline{\boldsymbol{\kappa}})$, and theory guarantees $\underline{\boldsymbol{\kappa}} \leq \partial \hat{\mathbf{c}}/\partial m \leq \bar{\boldsymbol{\kappa}}$ wherever the Euler equation is solved. Between gridpoints the property is not guaranteed, and we rely on the grid being fine enough to preserve it.

[15]Approximation quality between gridpoints is what ultimately limits any such method. Santos (2000) gauges it through Euler-equation residuals, as the lower panel of Table 1 does, and Judd et al. (2017) derives lower bounds on the resulting approximation errors.

[16]Cubic Hermite interpolation of the transformed ratio takes node values from the transformation and node slopes $\partial \boldsymbol{\chi}/\partial \boldsymbol{\mu}$. A monotone scheme (de Boor, 2001) (Fritsch-Carlson or Fritsch-Butland) treats those slopes as targets and adjusts them to enforce monotonicity.

[17]Differentiating (17) with respect to m and applying the chain rule: $\partial \hat{\mathbf{c}}/\partial m = \underline{\boldsymbol{\kappa}}\,(1 + \Delta h/\Delta m \cdot \partial \boldsymbol{\omega}/\partial \boldsymbol{\mu})$. The moderation form follows by factoring: $\boldsymbol{\eta}(\bar{\boldsymbol{\kappa}} - \underline{\boldsymbol{\kappa}}) = \underline{\boldsymbol{\kappa}} \cdot \Delta h/\Delta m \cdot \partial \boldsymbol{\omega}/\partial \boldsymbol{\mu}$. Note that $\partial \boldsymbol{\omega}/\partial \boldsymbol{\mu} = \boldsymbol{\omega}(1-\boldsymbol{\omega}) \cdot \partial \boldsymbol{\chi}/\partial \boldsymbol{\mu}$ from the chain rule, where $\boldsymbol{\omega}(1-\boldsymbol{\omega}) = \partial \boldsymbol{\omega}/\partial \boldsymbol{\chi}$ is the standard sigmoid derivative. Working directly with $\partial \boldsymbol{\omega}/\partial \boldsymbol{\mu}$ keeps the computation on the bounded scale of $\boldsymbol{\omega} \in [0, 1]$.

Table 2: Maximum absolute approximation errors by interval, as in Table 1 but with cubic Hermite interpolation in place of linear for both methods. Orders of magnitude in parentheses.

Method	$[m_0, m_1]$	$[m_1, m_2]$	$[m_2, m_3]$	$[m_3, m_4]$	$[m_4, \overline{m}]$
EGM	8.5(-3)	1.8(-4)	2.5(-5)	7.3(-6)	1.1(-1)
MoM	2.9(-3)	4.3(-6)	6.6(-7)	1.3(-7)	2.4(-3)

The weight reflects precautionary intensity, increasing when market resources are low relative to human wealth ($\Delta h/\Delta m$ large) and when the moderation ratio responds sharply to changes in log excess resources ($\partial\boldsymbol{\omega}/\partial\boldsymbol{\mu}$ large). As $\Delta m \to \infty$, $\boldsymbol{\eta} \to 0$ and the MPC approaches the optimist's minimal value $\underline{\boldsymbol{\kappa}}$; as $\Delta m \to 0$, $\boldsymbol{\eta} \to 1$ and the MPC approaches the maximal value $\bar{\boldsymbol{\kappa}}$ at the borrowing constraint.

Figure 9 shows both limits on the solved rule, with the moderated MPC inside the band at every solved gridpoint.

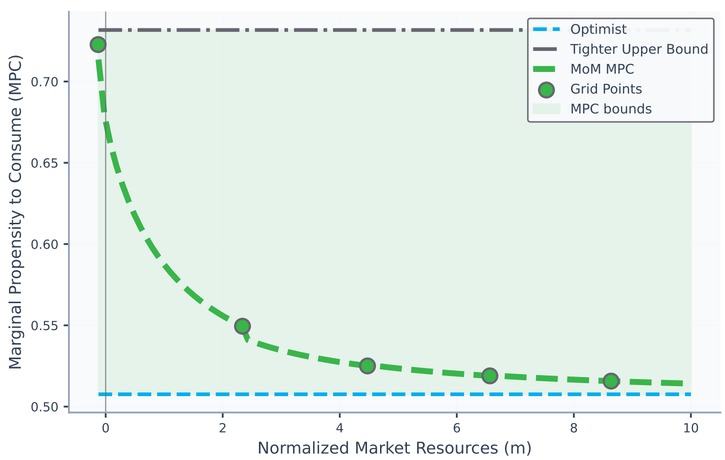

Figure 9: The Moderated MPC Is a Weighted Average of Its Two Limits at Every Solved Gridpoint

We now repeat the experiment of Table 1 in Table 2 with cubic Hermite in place of linear interpolation, for both methods on the same five-point grid.

Moderation's interior errors fall by a factor of five in the first interval and by roughly two to three orders of magnitude in the rest. EGM improves too, but by less beyond the first interval, where moderation's advantage grows from eleven or seventeen times under linear interpolation to between thirty-eight and fifty-six times under cubic, while in the first interval it slips from about three and a half times to three. The higher order gains more on the transformed moderation ratio than on consumption itself, so refining the interpolation widens the gap between the two methods.

Beyond the top gridpoint the errors barely move, 3.0(-3) to 2.4(-3) for moderation, and the ratio between the methods is forty-six in one table and forty-five in the other. Both moderation variants extrapolate the logit linearly there, so the beyond-grid advantage comes entirely from the change of variables, since interpolation order matters only where there are gridpoints to interpolate between.

The value function needs the same two ingredients, and we already have both in hand. Since $\mathcal{C}$ is constant under perfect foresight, the optimist's inverse value function has a constant slope with respect to cash-on-hand:

$$\begin{aligned} \bar{\mathbf{\Lambda}}' &= \mathcal{C}^{1/(1-\rho)} \underline{\boldsymbol{\kappa}} \\ &= \underline{\boldsymbol{\kappa}}^{-\rho/(1-\rho)}. \end{aligned} \tag{39}$$

The derivative of the value moderation ratio (23) with respect to the log-gap argument is

$$\frac{\partial \hat{\mathbf{\Omega}}}{\partial \boldsymbol{\mu}} = \frac{\Delta m(\hat{\mathbf{\Lambda}}' - \bar{\mathbf{\Lambda}}')}{\Delta h \bar{\mathbf{\Lambda}}'} \tag{40}$$

where $\bar{\mathbf{\Lambda}}'$ is given by (39) and $\hat{\mathbf{\Lambda}}'$ is the derivative of the realist's inverse value function. The pessimist's inverse value function has the same derivative $\bar{\mathbf{\Lambda}}'$, since both are linear perfect foresight solutions.

We apply the same transformation to the value-based moderation ratio, which converts the bounded ratio into an unconstrained slope:

$$\frac{\partial \hat{\boldsymbol{X}}}{\partial \boldsymbol{\mu}} = \frac{\partial \hat{\mathbf{\Omega}} / \partial \boldsymbol{\mu}}{\hat{\mathbf{\Omega}}(1 - \hat{\mathbf{\Omega}})}. \tag{41}$$

Note that $\hat{\mathbf{\Lambda}} = \mathbf{u}^{-1}(\hat{\mathbf{v}})$, so the chain rule supplies their derivatives and nothing further need be solved for. The first derivative is:

$$\hat{\mathbf{\Lambda}}' = ((1-\rho)\hat{\mathbf{v}}(m))^{-1+1/(1-\rho)} \, \hat{\mathbf{v}}'(m). \tag{42}$$

The first- and second-derivative connections are:

$$\begin{aligned} \hat{\mathbf{v}}' &= \mathbf{u}'(\hat{\mathbf{\Lambda}}) \, \hat{\mathbf{\Lambda}}' \\ \hat{\mathbf{v}}'' &= \mathbf{u}''(\hat{\mathbf{\Lambda}}) \, (\hat{\mathbf{\Lambda}}')^2 + \mathbf{u}'(\hat{\mathbf{\Lambda}}) \, \hat{\mathbf{\Lambda}}''. \end{aligned} \tag{43}$$

If we also match the second derivative in (43) with a higher-order Hermite polynomial (with $\hat{\mathbf{\Lambda}}''$ obtained by differentiating the expression for $\hat{\mathbf{\Lambda}}'$ once more), the approximation matches the marginal propensity to consume at the gridpoints as well.[18]

[18] The consumption function generated from the value function would then match both the level of consumption and the marginal propensity to consume at the gridpoints; within the grid, its numerical differences from the one constructed earlier are negligible.

Table 3: Maximum absolute consumption errors for $\Delta m \leq 1$ as the artificial borrowing constraint tightens, EGM and MoM on the same five-point grid against a dense reference solution. $\bar{\boldsymbol{\kappa}}$ is the period-$T-1$ value; the last column is the interval over which the realist's rule coincides with the tighter bound.

$\underline{a}$	$\bar{\boldsymbol{\kappa}}$	EGM	MoM	rule coincides with bound
natural	0.732	5.4(-2)	1.5(-2)	nowhere
-0.10	1.000	4.5(-2)	6.9(-3)	$\Delta m \leq 0.087$
-0.05	1.000	3.4(-2)	7.3(-3)	$\Delta m \leq 0.203$
0.00	1.000	2.6(-2)	7.5(-3)	$\Delta m \leq 0.303$
0.10	1.000	1.6(-2)	7.0(-3)	$\Delta m \leq 0.473$

6.3. Artificial Borrowing Constraints

Nothing in (30) requires the borrowing constraint to be the natural one. A modeler who imposes a constraint $a \geq \underline{a}$ tighter than the natural one changes both $\underline{m}$ and $\bar{\boldsymbol{\kappa}}$, and we need change nothing else: when the artificial constraint binds, $\bar{\boldsymbol{\kappa}} = 1$ (the derivation behind (27) presumes the natural constraint; against an artificial one the consumer spends every unit above the floor) and the tighter bound is the budget line itself; when the constraint is natural, $\bar{\boldsymbol{\kappa}} < 1$ and the bound lies strictly inside it. One expression covers both cases, which is fortunate, because a modeler who has just tightened a constraint is rarely in the mood to rederive an approximation.

The two bounds always cross, so the question for us is whether our realist ever reaches the tighter one. Under a natural constraint they approach $\bar{\boldsymbol{\kappa}}\Delta m$ only as $m \downarrow \underline{m}$ and touch it at the single point where both are zero, so the envelope has nothing to do. Under an artificial constraint they spend every resource they have over a whole interval above the constraint, up to the first kink point where it ceases to bind, and over that interval the envelope stops approximating and reproduces their rule.

We report what happens as we tighten the constraint in Table 3. The region over which the rule coincides with the bound grows, and moderation remains more accurate than endogenous gridpoints throughout. The advantage is largest where the constrained region is smallest, because that is where a method without the bound has the most room to overshoot it. We should be candid about how much rests on the envelope here. Drop it, and the moderated rule breaches the budget constraint in every row with an artificial constraint, spending resources the consumer does not have. The breach grows as the constraint tightens, and in the two tightest rows it is large enough to reverse the ordering (at $\underline{a} = 0$ the breach reaches 0.045, and the worst-case error of $4.5(-2)$ is now the larger of the two). Only the natural row, where the envelope has nothing to do, is unaffected.

6.4. Assets as the State Variable

We have written the consumption problem in market resources, while the standard incomplete-markets models of Huggett (1993), Aiyagari (1994), and

Krusell and Smith (1998) are written in assets and spend most of their computational budget solving the household problem we have been bracketing. Everything above transfers, because the bounds are perfect-foresight rules that hold at every m, and m is a function of whatever state the model uses.

To see that transfer, write down what those models actually carry. The state is the pair (a, z): assets a carried into the period, which together with consumption determine next period's assets as $a' = \mathrm{R}\, a + z - c$, and income z, a finite Markov chain with transition matrix Π. Since $m = \mathrm{R}\, a + z$, both bounds are affine in a with the common slope $\mathrm{R}\underline{\boldsymbol{\kappa}}$:

$$
\begin{aligned}
\bar{\mathbf{c}}(a, z) &= \underline{\boldsymbol{\kappa}}\big(\mathrm{R}\, a + z + \bar{h}(z)\big), \quad & \bar{h}(z) &= \frac{1}{\mathrm{R}} \sum_{z'} \Pi(z, z')\big[z' + \bar{h}(z')\big], \\
\underline{\mathbf{c}}(a, z) &= \underline{\boldsymbol{\kappa}}\big(\mathrm{R}\, a + z + \underline{h}\big), & \underline{h} &= \frac{\underline{z}}{\mathrm{R} - 1}.
\end{aligned}
\tag{44}
$$

The optimist's human wealth is expected discounted income along the chain, and carries the state because the chain does, while the pessimist's is the lowest income forever, the constant whose negative is the natural borrowing limit of Aiyagari (1994). Current income cancels from the gap, so the bracket has width $\underline{\boldsymbol{\kappa}}\big(\bar{h}(z) - \underline{h}\big)$, a quantity formed once per income state and constant in a within it, which is all the moderation ratio requires. The transformation and interpolation above are unchanged, applied to $\Delta m = \mathrm{R}\, a + z - \underline{m}$ within each income state.

The household problem transfers, though the layers above it do not, and on those (the fixed point in prices, the stationary distribution, the forecasting rule that closes a model with aggregate shocks) we say nothing.

7. Risky Returns and Portfolio Choice

The variations so far all concern how the consumption rule is approximated, while the asset structure has stayed the same, one riskless asset earning a known R. We now relax that in two steps, first making the return stochastic, so that the consumer's savings are entirely invested in a risky asset, and then letting the fraction so invested be a choice.

The two problems without a portfolio choice are the corners of that choice, since holding the risky share at zero recovers the riskless model of the preceding sections, and holding it at one recovers the model of Section 7.1. Making the return risky changes the *bounds* of the consumption rule, since the limiting marginal propensity to consume is no longer R-based, but it leaves the object being moderated where it was. Making the share a choice opens the interior between the corners and adds a second rule to approximate. Only the second step requires a new ingredient, because the transform that linearizes consumption does not linearize the share.

7.1. Idiosyncratic Rate-of-Return Risk

Thus far we have assumed that the interest factor is constant at R. Extending the previous derivations to allow for a perfectly forecastable time-varying interest factor would be trivial. Allowing for a stochastic interest factor is not trivial: it changes the two constants the bracket is built from, $\underline{\boldsymbol{\kappa}}$ and $\bar{h}$, though nothing else. We treat return risk that is idiosyncratic, the return on a household-specific asset such as a private business or housing, drawn independently across households and over time (Benhabib and Bisin, 2018), so the household's problem acquires no aggregate state and the interest factor is i.i.d.,[19]

$$\log \mathbf{R}_{t+n} \sim \mathcal{N}(r + \pi - \sigma_{\mathbf{r}}^2/2, \sigma_{\mathbf{r}}^2) \ \forall \ n > 0. \tag{45}$$

Both constants come from the high-wealth limit, where labor income is negligible relative to financial wealth, because in this case Merton (1969) and Samuelson (1969) showed that for a consumer without labor income the consumption function is linear, with an MPC that is the $\underline{\boldsymbol{\kappa}}$ of the bracket,

$$\underline{\boldsymbol{\kappa}} = 1 - \left(\beta \mathbb{E}[\mathbf{R}_{t+1}^{1-\rho}]\right)^{1/\rho}, \tag{46}$$

positive under the stochastic-return impatience condition $\beta \mathbb{E}[\mathbf{R}^{1-\rho}] < 1$, the i.i.d.-return analogue of $\text{Þ}/\text{R} < 1$ (Carroll, 2020a; Benhabib and Bisin, 2018; Chipeniuk et al., 2022). Human wealth needs a rate to discount at, and with no traded riskless asset there is none for us to appeal to, but the Euler equation supplies one, the implicit riskless return of the economy, which is the reciprocal of the expected stochastic discount factor (Campbell and Viceira, 2002). In the same limit consumption grows at $\mathbf{R}_{t+1}(1 - \underline{\boldsymbol{\kappa}})$, so that discount factor is $\beta(\mathbf{R}_{t+1}(1 - \underline{\boldsymbol{\kappa}}))^{-\rho}$, and substituting (46) for $(1 - \underline{\boldsymbol{\kappa}})^{\rho}$ cancels the preference parameters, leaving the shadow riskless return $\tilde{\text{R}}$,

$$\mathbb{E}\left[\beta \left(\mathbf{R}_{t+1}(1 - \underline{\boldsymbol{\kappa}})\right)^{-\rho}\right] = \frac{\mathbb{E}[\mathbf{R}^{-\rho}]}{\mathbb{E}[\mathbf{R}^{1-\rho}]} \equiv 1/\tilde{\text{R}}, \tag{47}$$

so we obtain the two human wealths by writing their riskless expressions with $\tilde{\text{R}}$ in place of R, $\bar{h} = G/(\tilde{\text{R}} - G)$ and $\underline{h} = \underline{\xi} G/(\tilde{\text{R}} - G)$. The two constants use the return distribution differently, the MPC needing the single moment $\mathbb{E}[\mathbf{R}^{1-\rho}]$ and human wealth the ratio of two, and a deterministic return collapses that ratio onto R, which is why we needed only one rate when the return was certain. Finite human wealth requires $\mathbb{E}[\mathbf{R}^{-\rho}] < \mathbb{E}[\mathbf{R}^{1-\rho}]$, the return-risk counterpart of the finite-human-wealth condition (the riskless-return conditions are in Carroll and Shanker (2026)).

With the $\underline{\boldsymbol{\kappa}}$ of (46) and the $\tilde{\text{R}}$ of (47) in place, the earlier construction applies unchanged. But the two sides of the bracket rest on different footing. Carroll and Lujan (2026) prove at full generality that the pessimist's rule bounds optimal consumption from below whether or not returns are risky. The upper

[19] Here r is the log of the earlier calibration's riskless return and π the premium of the mean log return over it. Since no riskless asset trades in this subsection, the pair only parameterizes the distribution in the earlier units. The $-\sigma_{\mathbf{r}}^2/2$ term makes $\mathbb{E}[\mathbf{R}] = \exp(r+\pi)$, so raising $\sigma_{\mathbf{r}}^2$ is a mean-preserving spread of the level return, and the lognormal moment generating function gives $\mathbb{E}[\mathbf{R}^{1-\rho}] = \exp\left((1-\rho)\left(r + \pi - \rho\sigma_{\mathbf{r}}^2/2\right)\right)$, the expression the code evaluates.

bound is open there, because the Jensen step that delivers it under deterministic returns fails once the return enters the marginal-utility argument. We check the upper bound numerically, and with the mean return equal to the riskless return of the earlier calibration and a level standard deviation of 0.08, the realist's rule stays strictly inside the bracket at every one of two thousand evaluation points between the borrowing constraint and $m = 10$, the moderation ratio ranging over $[0.005, 0.770]$.[20]

7.2. Portfolio Choice

Every application so far has moderated a consumption rule between two perfect-foresight consumption rules. The share of savings held in the risky asset is a sterner test, being a different choice variable with bounds of a different kind.

The share has a single limiting agent instead of two. As wealth grows, labor income becomes negligible and the optimal share approaches the constant ς^* that solves the portfolio problem of an investor with financial wealth alone (Samuelson, 1969; Merton, 1969), the first-order condition

$$\mathbb{E}\big[(\mathbf{R} - \mathrm{R})\, \mathbf{u}'\big(\mathrm{R} + \varsigma^*(\mathbf{R} - \mathrm{R})\big)\big] = 0. \tag{48}$$

This is a single-period problem, which is why Campbell and Viceira (2002) call its solution the 'myopic' share, and with i.i.d. returns it is the same at every wealth level. Their log-linear approximation $\pi/(\rho\sigma_{\mathbf{r}}^2)$ gives 0.28 at the calibration below, whereas we use the 0.31 solved numerically on the discretized return process, for which it is exact. That limit is also a floor, because human wealth is a bond-like position, so that a household holding it tilts financial wealth toward the risky asset, and the share lies above ς^* at every finite wealth level (Campbell and Viceira, 2002; Viceira, 2001). That floor depends on income shocks being independent of returns, as they are here. When labor income is cointegrated with the stock market, the young want to short the market (Benzoni et al., 2007), and ς^* is no floor at all.

The two characters of the consumption problem are of no use here. A pessimist who expected the worst return every period would hold no stock at all (the worst return is below the riskless one), and an optimist who ignored the risk would hold nothing else, so a bracket built on the pair would be the constraint set $[0, 1]$ itself, within which moderating the share is the same thing as interpolating it. The upper edge is instead a constraint: the prohibition on borrowing to invest, $\varsigma \leq 1$. This is the upper half of the share constraint $\varsigma \in [0, 1]$ that the portfolio literature imposes alongside $a \geq 0$ (Cocco et al., 2005), as we do here. Together the two rule out debt of any kind, so next period's resources are at least the income floor whatever the return turns out to be. They also supply a ceiling the share would otherwise not have, since a household with labor income to lean on would, if unconstrained, put an ever larger multiple of its financial

[20] The test suite asserts strict containment, not containment to a tolerance, and first checks that the spread lies below the volatility (near 0.10 here) at which $\mathbb{E}[\mathbf{R}^{-\rho}]$ overtakes $\mathbb{E}[\mathbf{R}^{1-\rho}]$, so that the condition above holds.

wealth into the risky asset as that wealth shrank toward zero. We call the upper half the leverage constraint, to keep it distinct from the constraint on assets.

The bracket is therefore $[\varsigma^*, 1]$, and its width $1 - \varsigma^*$ is constant in the state, as the width of the consumption bracket is. Bodie et al. (1992) obtain a tighter ceiling, $\varsigma^*(1 + \bar{h}/m)$, for riskless labor income, but the solved share crosses it here, because that ceiling's gap above ς^* is proportional to $\bar{h}/m$, while the solved share's gap closes more slowly. So we bracket against the constraint, which cannot be crossed.

The illustrations solve the infinite-horizon problem with relative risk aversion $\rho = 5$, discount factor 0.9, survival probability 0.98 (which scales the discount factor), riskless return $\mathrm{R} = 1.03$, permanent income growth factor 1.01, permanent and transitory income shocks each with log standard deviation 0.1, a five percent chance of an unemployment draw paying thirty percent of permanent income, and a risky asset with mean gross return 1.08 and log standard deviation 0.184. The construction does not depend on these choices.

The moderation ratio for the share, $\boldsymbol{\omega}_\varsigma = (\varsigma - \varsigma^*)/(1 - \varsigma^*)$, records where the solved share lies in its bracket, as $\boldsymbol{\omega}$ does for consumption. The share approaches its limiting value only as wealth grows, so $\boldsymbol{\omega}_\varsigma \to 0$ asymptotically and $\log \boldsymbol{\omega}_\varsigma$ straightens that tail. Its other bound, the leverage constraint, is attained at finite wealth, where $-\log(1 - \boldsymbol{\omega}_\varsigma)$ diverges, so the logit of Section 4.2 is the wrong transform here and the logarithm alone is the right one.

Figure 10 shows the consequence, with each candidate transformation rescaled to a common range and drawn against the straight chord through its own endpoints. Over $m \in [10, 10^3]$ the untransformed share departs from its chord by 39 percent of the plotted range and the logit by 14 percent, while $\log \boldsymbol{\omega}_\varsigma$ departs by 1.4 percent. Scored on accuracy against a converged reference with five nodes, the logarithm is roughly a hundred times more accurate than interpolating the share itself, the complementary log-log, $\log(-\log(1 - \boldsymbol{\omega}_\varsigma))$, eight times, and the logit four times. Nothing here fixes how the share approaches its limit, since the interpolant is fitted to the solved gridpoints, so a calibration whose tail behaves differently is served by the same construction.

At low wealth the leverage constraint binds and the solved share is exactly one. The point at which it stops binding is a kink in the share function, and unlike the cusp of (29) it has no closed form. The cusp was the intersection of two straight lines whose coefficients we knew, whereas the kink is the level of wealth at which the investor's own first-order condition first holds with the share at one, and that depends on next period's marginal value, which we do not know until the problem has been solved. So we take the kink from the solved rule, in which the investor leaves the constraint at $m = 7.25$ (Figure 11). The approximation is exact on the constraint below it, moderated above it, and joined across it by a Hermite segment matched in level and slope at the two adjacent gridpoints. A cubic matched at both ends can overshoot the endpoint it starts from, so the approximation takes the lower envelope against the constraint, as the consumption rule does against $\bar{\boldsymbol{\kappa}}\Delta m$, since the solved share never exceeds one.

Clear of the kink point, moderation is between seventy and a hundred and

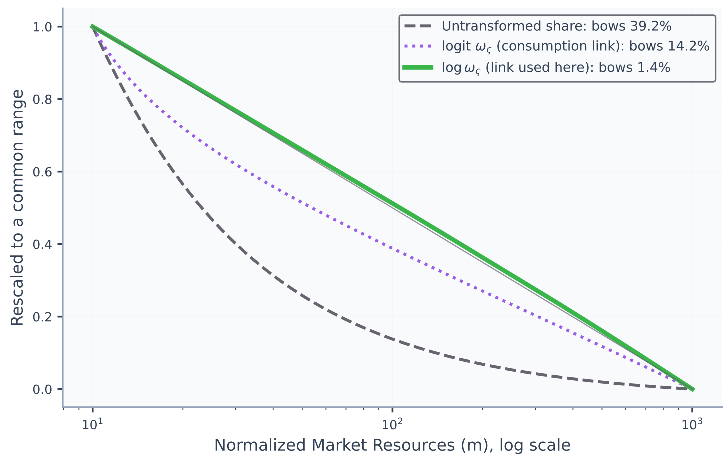

Figure 10: The Logarithm of the Share's Moderation Ratio Is Nearly Straight Where the Logit Bows

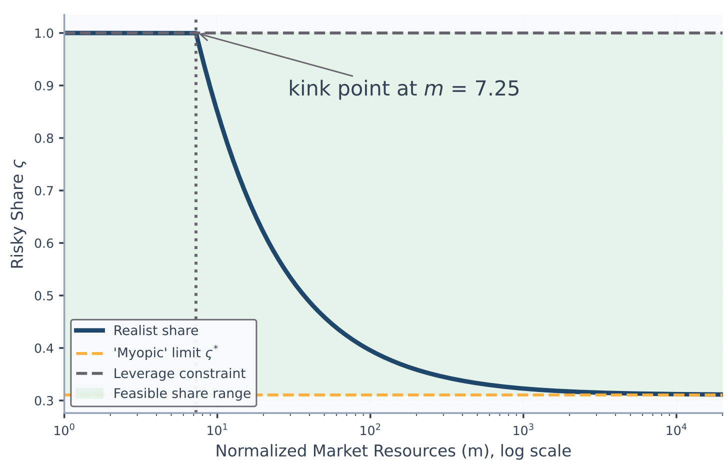

Figure 11: The Share Is Fully Invested Until Its Kink Point, Then Descends Toward Its Limit

fifty times more accurate than interpolating the share directly on the same four to seven nodes, with the advantage growing as the grid coarsens. On a grid spanning the constrained region the advantage is smaller, and the reader is entitled to the whole picture. In the tail it runs from twelve to over a hundred times, across the transition it is about twice, and on a grid coarse enough that the kink point falls in a wide interval our joining segment is worse than plain interpolation there.

Beyond the top gridpoint, Figure 12 is the portfolio analogue of Figure 1. Linear extrapolation carries the share below its limiting value, a portfolio no optimizing investor would hold at any wealth, while the moderated rule continues the descent recorded at the solved gridpoints and approaches the limit from above.

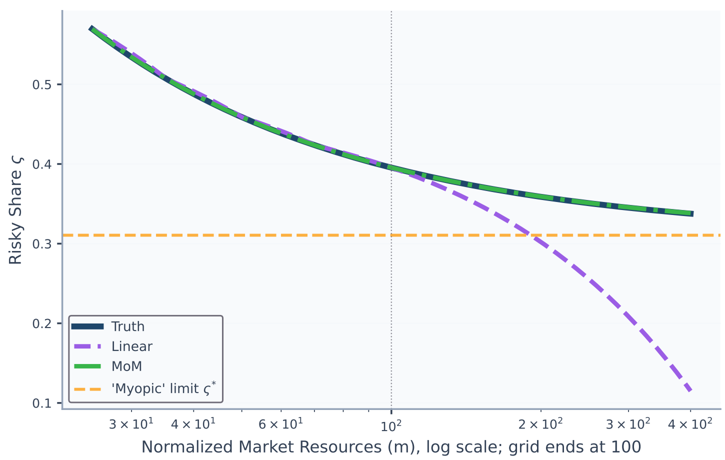

Figure 12: Beyond the Grid, Linear Extrapolation Breaches the Limiting Share and Moderation Does Not

Before we claim too much, we should say that the share is an output of each period's optimization rather than an input to the next period's, so a moderated share does not propagate through the backward induction the way a moderated consumption rule does. The construction supplies a representation, one in which the expensive per-node optimization can be performed on far fewer nodes because the share between them is moderated instead of interpolated.

8. Conclusion

The method proposed here is not universally applicable. For example, the method cannot be used for problems for which upper and lower bounds to the 'true' solution are not known. Those bounds need not come in pairs of perfect-foresight agents, as the portfolio application shows by bracketing the share between the 'myopic' investor and the leverage constraint. But many problems do have obvious upper and lower bounds, and in those cases (as in the consumption and portfolio examples used in the paper), the method may result in substantial improvements in accuracy and stability of solutions.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude Code (Anthropic) to assist in editing and revising the prose and the code, and in checking the numerical claims in the text against the replication package. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Aiyagari, S.R., 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. Quarterly Journal of Economics 109, 659–684. doi:10.2307/2118417.

Barillas, F., Fernández-Villaverde, J., 2007. A Generalization of the Endogenous Grid Method. Journal of Economic Dynamics and Control 31, 2698–2712. doi:10.1016/j.jedc.2006.08.005.

Benhabib, J., Bisin, A., 2018. Skewed Wealth Distributions: Theory and Empirics. Journal of Economic Literature 56, 1261–1291. doi:10.1257/jel.20161390.

Benveniste, L.M., Scheinkman, J.A., 1979. On the Differentiability of the Value Function in Dynamic Models of Economics. Econometrica 47, 727–732. doi:10.2307/1910417.

Benzoni, L., Collin-Dufresne, P., Goldstein, R.S., 2007. Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated. The Journal of Finance 62, 2123–2167. doi:10.1111/j.1540-6261.2007.01271.x.

Bodie, Z., Merton, R.C., Samuelson, W.F., 1992. Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model. Journal of Economic Dynamics and Control 16, 427–449.

de Boor, C., 2001. A Practical Guide to Splines. volume 27 of *Applied Mathematical Sciences*. Revised ed., Springer.

Campbell, J.Y., Viceira, L.M., 2002. Strategic Asset Allocation: Portfolio Choice for Long-Term Investors. Oxford University Press, Oxford.

Carroll, C.D., 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. Quarterly Journal of Economics 112, 1–55. doi:10.1162/003355397555109.

Carroll, C.D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. Economics Letters 91, 312–320. doi:10.1016/j.econlet.2005.09.013.

Carroll, C.D., 2009. Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income. Journal of Monetary Economics 56, 780–790. doi:10.1016/j.jmoneco.2009.06.016.

Carroll, C.D., 2020a. CRRA utility with rate risk. techreport. Johns Hopkins University. Lecture notes.

Carroll, C.D., 2020b. Solving microeconomic dynamic stochastic optimization problems. techreport. Johns Hopkins University. URL: https://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs.pdf.

Carroll, C.D., Kimball, M.S., 1996. On the Concavity of the Consumption Function. Econometrica 64, 981–992. doi:10.2307/2171853.

Carroll, C.D., Lujan, A., 2026. State-Dependent Buffer-Stock Theory. Working paper.

Carroll, C.D., Shanker, A., 2026. Theoretical Foundations of Buffer Stock Saving. URL: https://llorracc.github.io/BufferStockTheory/BufferStockTheory.pdf. revise and Resubmit, Quantitative Economics.

Carroll, C.D., Toche, P., 2009. A Tractable Model of Buffer Stock Saving. Working Paper 15265. National Bureau of Economic Research. doi:10.3386/w15265.

Chipeniuk, K.O., 2020. Optimal Grid Selection for the Numerical Solution of Dynamic Stochastic Optimization Problems. Computational Economics 56, 883–928. doi:10.1007/s10614-019-09953-4.

Chipeniuk, K.O., Katz, N.H., Walker, T.B., 2022. Households, Auctioneers, and Aggregation. European Economic Review 141, 103997. doi:10.1016/j.euroecorev.2021.103997.

Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and Portfolio Choice over the Life Cycle. Review of Financial Studies 18, 491–533. doi:10.1093/rfs/hhi017.

Deaton, A., 1991. Saving and Liquidity Constraints. Econometrica 59, 1221–1248. doi:10.2307/2938366.

Fella, G., 2014. A Generalized Endogenous Grid Method for Non-smooth and Non-concave Problems. Review of Economic Dynamics 17, 329–344. doi:10.1016/j.red.2013.07.001.

Friedman, M., 1957. Theory of the Consumption Function. Princeton University Press, Princeton, NJ. doi:10.1515/9780691188485.

Fritsch, F.N., Butland, J., 1984. A Method for Constructing Local Monotone Piecewise Cubic Interpolants. SIAM Journal on Scientific and Statistical Computing 5, 300–304. doi:10.1137/0905021.

Fritsch, F.N., Carlson, R.E., 1980. Monotone Piecewise Cubic Interpolation. SIAM Journal on Numerical Analysis 17, 238–246. doi:10.1137/0717021.

Gouin-Bonenfant, É., Toda, A.A., 2023. Pareto Extrapolation: An Analytical Framework for Studying Tail Inequality. Quantitative Economics 14, 201–233. doi:10.3982/QE1817.

Hintermaier, T., Koeniger, W., 2010. The Method of Endogenous Gridpoints with Occasionally Binding Constraints among Endogenous Variables. Journal of Economic Dynamics and Control 34, 2074–2088. doi:10.1016/j.jedc.2010.05.002.

Huggett, M., 1993. The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. Journal of Economic Dynamics and Control 17, 953–969. doi:10.1016/0165-1889(93)90024-m.

Hyman, J.M., 1983. Accurate Monotonicity-Preserving Cubic Interpolation. SIAM Journal on Scientific and Statistical Computing 4, 645–654. doi:10.1137/0904045.

Iskhakov, F., Jørgensen, T.H., Rust, J., Schjerning, B., 2017. The Endogenous Grid Method for Discrete-Continuous Dynamic Choice Models with (or without) Taste Shocks. Quantitative Economics 8, 317–365. doi:10.3982/qe643.

Judd, K.L., Maliar, L., Maliar, S., 2017. Lower Bounds on Approximation Errors to Numerical Solutions of Dynamic Economic Models. Econometrica 85, 991–1012. doi:10.3982/ecta12791.

Kimball, M.S., 1990. Precautionary Saving in the Small and in the Large. Econometrica 58, 53–73. doi:10.2307/2938334.

Krusell, P., Smith, Jr., A.A., 1998. Income and Wealth Heterogeneity in the Macroeconomy. Journal of Political Economy 106, 867–896.

Leland, H.E., 1968. Saving and Uncertainty: The Precautionary Demand for Saving. Quarterly Journal of Economics 82, 465–473. doi:10.2307/1879518.

Ma, Q., Stachurski, J., Toda, A.A., 2020. The Income Fluctuation Problem and the Evolution of Wealth. Journal of Economic Theory 187, 105003. doi:10.1016/j.jet.2020.105003.

Ma, Q., Toda, A.A., 2021. A Theory of the Saving Rate of the Rich. Journal of Economic Theory 192, 105193. doi:10.1016/j.jet.2021.105193.

Ma, Q., Toda, A.A., 2022. Asymptotic Linearity of Consumption Functions and Computational Efficiency. Journal of Mathematical Economics 98, 102562. doi:10.1016/j.jmateco.2021.102562.

Merton, R.C., 1969. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. Review of Economics and Statistics 51, 247–257. doi:10.2307/1926560.

Milgrom, P., Segal, I., 2002. Envelope Theorems for Arbitrary Choice Sets. Econometrica 70, 583–601. doi:10.1111/1468-0262.00296.

Muth, J.F., 1960. Optimal Properties of Exponentially Weighted Forecasts. Journal of the American Statistical Association 55, 299–306. doi:10.1080/01621459.1960.10482064.

Samuelson, P.A., 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming. Review of Economics and Statistics 51, 239–246. doi:10.2307/1926559.

Sandmo, A., 1970. The Effect of Uncertainty on Saving Decisions. Review of Economic Studies 37, 353–360. doi:10.2307/2296725.

Santos, M.S., 2000. Accuracy of Numerical Solutions Using the Euler Equation Residuals. Econometrica 68, 1377–1402. doi:10.1111/1468-0262.00165.

Stachurski, J., Toda, A.A., 2019. An Impossibility Theorem for Wealth in Heterogeneous-Agent Models with Limited Heterogeneity. Journal of Economic Theory 182, 1–24. doi:10.1016/j.jet.2019.04.001. provides an explicit linear lower bound on consumption; Corrigendum 2020.

Tauchen, G., 1986. Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. Economics Letters 20, 177–181. doi:10.1016/0165-1765(86)90168-0.

Tauchen, G., Hussey, R., 1991. Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models. Econometrica 59, 371–396. doi:10.2307/2938261.

Viceira, L.M., 2001. Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. The Journal of Finance 56, 433–470.

White, M.N., 2015. The Method of Endogenous Gridpoints in Theory and Practice. Journal of Economic Dynamics and Control 60, 26–41. doi:10.1016/j.jedc.2015.08.001.

Zeldes, S.P., 1989. Consumption and Liquidity Constraints: An Empirical Investigation. Journal of Political Economy 97, 305–346. doi:10.1086/261605.