

THE METHOD OF MODERATION

PREPRINT, COMPILED AUGUST 18, 2025

Christopher D. Carroll ^{1*}, Alan Lujan ^{1†}, Kiichi Tokuoka^{2, 3‡}, and Weifeng Wu⁴

¹Johns Hopkins University

²European Central Bank

³International Monetary Fund

⁴Fannie Mae

ABSTRACT

In a risky world, a pessimist assumes the worst will happen. Someone who ignores risk altogether is an optimist. Consumption decisions are mathematically simple for both the pessimist and the optimist because both behave as if they live in a riskless world. A realist (someone who wants to respond optimally to risk) faces a much more difficult problem, but (under standard conditions) will choose a level of spending somewhere between pessimist's and the optimist's. We use this fact to redefine the space in which the realist searches for optimal consumption rules. The resulting solution accurately represents the numerical consumption rule over the entire interval of feasible wealth values with remarkably few computations.

Keywords Dynamic Stochastic Optimization

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 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.

This paper argues that a better approach in the standard consumption problem 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 as being attached to that future income. The 'realistic' consumer, who *does* perceive the risks, will engage in 'precautionary saving' [1, 2, 3], 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 his certainty is that of a pessimist perfectly confident in his pessimism.

We build on bounds for the consumption function and limiting MPCs established in buffer-stock theory and related work [4, 5]. Using results from Carroll and Shanker [6], we show how to use these upper and lower bounds to tightly constrain the shape and characteristics of the solution to problem of the 'realist' [7]. Imposition of these constraints can clarify and speed the solution of the realist's problem.

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.

2 THE REALIST'S PROBLEM

We assume that truly optimal behavior in the problem facing the consumer who understands all his risks is captured by⁴

$$\max \mathbb{E}_t \left[\sum_{n=0}^{T-t} \beta^n \mathbf{u}(\mathbf{c}_{t+n}) \right] \quad (2)$$

subject to

$$\begin{aligned} \mathbf{a}_t &= \mathbf{m}_t - \mathbf{c}_t \\ \mathbf{p}_{t+1} &= \mathbf{p}_t \mathcal{G}_{t+1} \\ \mathbf{y}_{t+1} &= \mathbf{p}_{t+1} \boldsymbol{\xi}_{t+1} \\ \mathbf{m}_{t+1} &= \mathbf{a}_t \mathbf{R}_{t+1} + \mathbf{y}_{t+1} \end{aligned} \quad (3)$$

where

⁴Where the utility function is CRRA with risk aversion parameter $\rho > 0$

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

- β - pure time discount factor
- \mathbf{a}_t - assets at the end of period t
- \mathbf{c}_t - consumption in period t
- \mathbf{m}_t - ‘market resources’ available for consumption (4)
- \mathbf{p}_{t+1} - ‘permanent labor income’ in period $t + 1$
- \mathbf{R}_{t+1} - interest factor $(1 + r_{t+1})$ from period t to $t + 1$
- \mathbf{y}_{t+1} - noncapital income in period $t + 1$.

and the exogenous variables evolve according to the *Friedman-Muth Income Process*⁵

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

where the permanent shocks ψ_{t+1} are independently and identically distributed with mean $\mathbb{E}[\psi_{t+1}] = 1$ and support on $[\underline{\psi}, \bar{\psi}]$, with $0 < \underline{\psi} \leq 1 \leq \bar{\psi} < \infty$, and the transitory shocks θ_{t+1} are independently distributed with mean $\mathbb{E}[\theta_{t+1}] = 1$ and bounded support $\underline{\theta} \leq \theta_{t+1} \leq \bar{\theta}$, with $0 \leq \underline{\theta} \leq 1 \leq \bar{\theta} < \infty$.

It turns out (see Carroll [10] 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

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

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

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

For the remainder of the paper we will assume that permanent income \mathbf{p}_t grows by the permanent growth shock $\mathcal{G}_{t+1} = G_{t+1}\psi_{t+1}$, where G_{t+1} is the deterministic permanent growth factor and ψ_{t+1} is the idiosyncratic permanent shock. For convenience, we define the absolute patience factor as $\Phi \equiv (\beta\mathbf{R})^{1/\rho}$. A finite solution requires [6]: (i) the finite-value-of-autarky

⁵Friedman [8] introduced the distinction between permanent and transitory components of income, where permanent income reflects long-term earning capacity and transitory income captures temporary fluctuations. Muth [9] developed the stochastic foundations for modeling these income components as random processes with specific distributional properties. Our Friedman-Muth specification combines Friedman’s economic insight about income persistence with Muth’s rigorous stochastic framework, allowing for realistic modeling of both unemployment risk and permanent income growth.

condition (FVAC) $0 < \beta G^{1-\rho} \mathbb{E}[\psi^{1-\rho}] < 1$; (ii) the absolute-impatience condition (AIC) $\Phi < 1$; (iii) the return-impatience condition (RIC) $\Phi/\mathbf{R} < 1$; (iv) the growth-impatience condition (GIC) $\Phi/G < 1$; and (v) the finite-human-wealth condition (FHW) $G/\mathbf{R} < 1$. These conditions ensure existence of upper and lower bounds on consumption [11, 4] and pin down limiting MPCs [12]. Under perfect foresight with $\mathcal{G} = 1$, FVAC becomes $0 < \beta < 1$, the GIC coincides with the AIC, and the FHW reduces to $\mathbf{R} > 1$.

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 [13]. 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 include multi-dimensional problems [14], occasionally binding constraints [15], non-smooth and non-concave problems [16], discrete-continuous choice models [17], and comprehensive treatments of theory and practice [18]. For shock discretization in numerical solutions, standard methods include [19, 20].

Carroll [10] 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).⁶

Unfortunately, this endogenous gridpoints solution is not very well-behaved outside the original range of gridpoints targeted by the solution method. (Though 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{c}_{T-1} minus our approximation to optimal consumption under uncertainty, \hat{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).

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.

⁶The five gridpoints are used for illustration; production codes typically use 30-80 gridpoints for accurate solutions.

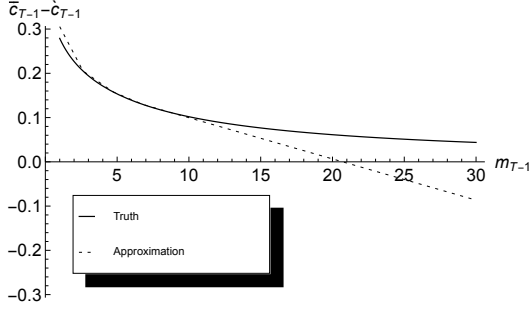


Figure 1: For Large Enough m_{T-1} , Predicted Precautionary Saving is Negative (Oops!)

4 THE METHOD OF MODERATION

4.1 The Optimist, the Pessimist, and the Realist

As a preliminary to our solution, define \bar{h}_t ⁷ 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 period- t consumer who believes with perfect confidence that the shocks will always take their expected value of 1, $\xi_{t+n} = \mathbb{E}[\xi] = 1 \forall n > 0$. The solution to a perfect foresight problem of this kind takes the form⁸

$$\bar{c}_t(m_t) = (m_t + \bar{h}_t)\kappa_t \quad (8)$$

for a constant minimal marginal propensity to consume κ_t given below.⁹ We similarly define \underline{h}_t ¹⁰ as ‘minimal human wealth,’ the present discounted value of labor income if the shocks were to take on their worst value in every future period $\xi_{t+n} = 0$ or $\underline{\xi} \forall n > 0$ (which we define as corresponding to the beliefs of a ‘pessimist’).

⁷In normalized variables, with constant interest factor $R > 1$ and deterministic permanent growth factor G , and setting $\xi_{t+n} = 1$ (the optimist’s assumption), human wealth can be calculated in three equivalent ways: (i) finite-horizon backward recursion: $\bar{h}_T = 0$, $\bar{h}_t = (G/R)(1 + \bar{h}_{t+1})$ for $t = T-1, T-2, \dots$; (ii) finite-horizon forward sum: $\bar{h}_t = \sum_{n=1}^{T-t} (G/R)^n$; (iii) infinite-horizon limit: $\bar{h}_\infty = G/(R - G)$ provided $R > G$. When $G = 1$ (as used in the main text), these expressions simplify to $\bar{h}_t = \sum_{n=1}^{T-t} (1/R)^n$ and $\bar{h}_\infty = 1/(R - 1)$.

⁸For a derivation, see Carroll and Shanker [6]; κ_t is defined therein as the MPC of the perfect foresight consumer with horizon $T - t$.

⁹This represents the MPC of the perfect foresight consumer with horizon $T - t$, corresponding to the optimist’s consumption behavior when market resources are very large. Equivalent expressions: (i) backward recursion $\kappa_t = \kappa_{t+1}/(\kappa_{t+1} + \Phi/R)$ with terminal condition $\kappa_T = 1$; (ii) finite-horizon forward sum $\kappa_t = (\sum_{n=0}^{T-t} (\Phi/R)^n)^{-1}$; (iii) closed-form infinite-horizon limit $\kappa_\infty = 1 - \Phi/R$.

¹⁰In normalized variables, with constant interest factor $R > 1$ and deterministic permanent growth factor G , and setting $\xi_{t+n} = \underline{\xi} \forall n > 0$ (the pessimist’s assumption), minimal human wealth can be calculated in three equivalent ways. In the Friedman-Muth process, $\underline{\xi} \equiv 0$ since $\varphi > 0$ represents the probability of unemployment (zero transitory income). The three calculation methods are: (i) finite-horizon backward recursion: $\underline{h}_T = 0$, $\underline{h}_t = (G/R)(\underline{\xi} + \underline{h}_{t+1})$ for $t = T-1, T-2, \dots$; (ii) finite-horizon forward sum: $\underline{h}_t = \underline{\xi} \sum_{n=1}^{T-t} (G/R)^n$; (iii) infinite-horizon limit: $\underline{h}_\infty = \underline{\xi} G/(R - G)$ provided $R > G$. When $\underline{\xi} = 0$ (unemployment possible), these expressions reduce to $\underline{h}_t = 0$.

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

A first useful point is that, for the realist, a lower bound for the level of market resources is $\underline{m}_t = -\underline{h}_t$ [21, 22], because if m_t equalled this value then there would be a positive finite chance (however small) of receiving $\xi_{t+n} = 0$ or $\underline{\xi}$ in every future period, which would require the consumer to set c_t to zero in order to guarantee that the intertemporal budget constraint holds. Since consumption of zero yields negative infinite utility, the solution to the realist consumer’s problem is not well defined for values of $m_t < \underline{m}_t$ [23, 24], and the limiting value of the realist’s c_t is zero as $m_t \downarrow \underline{m}_t$.

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:

$$\begin{aligned} \dot{m}_t &= m_t + \overbrace{\underline{h}_t}^{=-\underline{m}} \\ \dot{h}_t &= \bar{h}_t - \underline{h}_t. \end{aligned} \quad (9)$$

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}_t with certainty, so consumption is given by the perfect foresight solution

$$\begin{aligned} \underline{c}_t(m_t) &= (m_t + \underline{h}_t)\kappa_t \\ &= \dot{m}_t\kappa_t. \end{aligned} \quad (10)$$

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

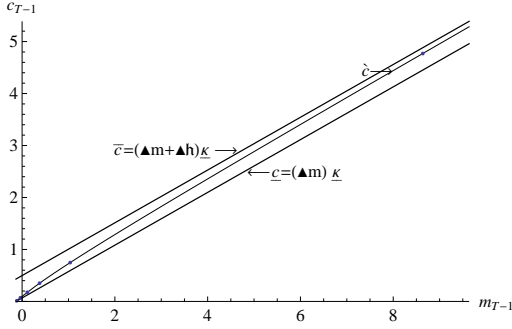
$$\begin{aligned} \bar{c}_t(m_t) &= (m_t + \underline{h}_t - \underline{h}_t + \bar{h}_t)\kappa_t \\ &= (\dot{m}_t + \dot{h}_t)\kappa_t \\ &= \underline{c}_t(m_t) + \dot{h}_t\kappa_t. \end{aligned} \quad (11)$$

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. 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 [6]¹¹ so we will take the proposition as a fact and proceed by manipulating the inequality:

¹¹Under prudence ($\mathbf{u}''' > 0$) and bounded shocks with strictly positive support (nondegenerate lower and upper bounds), the consumption function is strictly increasing and concave, and the moderation ratio lies in $(0, 1)$; see Carroll and Shanker [6].

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

$$\begin{aligned}
 \bar{c}_t(m_t + \dot{m}_t) &> \hat{c}_t(m_t + \dot{m}_t) > \underline{c}_t(m_t + \dot{m}_t) \\
 -\bar{c}_t(m_t + \dot{m}_t) &< -\hat{c}_t(m_t + \dot{m}_t) < -\underline{c}_t(m_t + \dot{m}_t) \\
 0 &< \frac{\bar{c}_t(m_t + \dot{m}_t) - \hat{c}_t(m_t + \dot{m}_t)}{\dot{h}_t \underline{c}_t} < 1 \\
 0 &< \underbrace{\left(\frac{\bar{c}_t(m_t + \dot{m}_t) - \hat{c}_t(m_t + \dot{m}_t)}{\dot{h}_t \underline{c}_t} \right)}_{\equiv \bar{\omega}_t} < 1
 \end{aligned} \tag{12}$$

where the fraction in the middle of the last inequality is the ratio of actual precautionary saving (the numerator is the difference between perfect-foresight consumption and optimal consumption in the presence of uncertainty) to the maximum conceivable amount of precautionary saving (the amount that would be undertaken by the pessimist who consumes nothing out of any future income beyond the perfectly certain component).¹² Defining $\mu_t = \log \dot{m}_t$ (which can range from $-\infty$ to ∞), the object in the middle of the last inequality is

$$\bar{\omega}_t(\mu_t) \equiv \left(\frac{\bar{c}_t(m_t + e^{\mu_t}) - \hat{c}_t(m_t + e^{\mu_t})}{\dot{h}_t \underline{c}_t} \right), \tag{13}$$

and we now define

$$\begin{aligned}
 \bar{\chi}_t(\mu_t) &= \log \left(\frac{1 - \bar{\omega}_t(\mu_t)}{\bar{\omega}_t(\mu_t)} \right) \\
 &= \log(1/\bar{\omega}_t(\mu_t) - 1)
 \end{aligned} \tag{14}$$

which has the virtue that it is *asymptotically linear* in the limit as μ_t approaches $+\infty$.¹³

¹²This ratio is strictly between 0 and 1 for all $m_t > \underline{m}_t$ under strict monotonicity of the consumption function, which is ensured by prudence ($u''' > 0$) and bounded shocks [25], as demonstrated in Carroll and Shanker [6]. The denominator $\dot{h}_t \underline{c}_t$ is constant in m_t within a given period, so the ratio is well-defined. Equivalently, precautionary saving satisfies $\bar{c}_t - \hat{c}_t = \bar{\omega}_t \dot{h}_t \underline{c}_t$ with $\bar{\omega}_t \in (0, 1)$, which ensures it is strictly positive.

¹³Under prudence ($u''' > 0$) and GIC, $\bar{\chi}_t(\mu_t)$ is asymptotically linear: its slope converges to a constant $\alpha_t \geq 0$ as $\mu_t \rightarrow +\infty$. At the infinite limit α_t may equal 0; on any finite grid, the boundary slope used in computation is strictly positive. For practical numerical work on finite grids, we extrapolate $\bar{\chi}_t$ linearly using that boundary slope. This approach preserves $\bar{\omega}_t \in (0, 1)$ and hence $\underline{c}_t < \hat{c}_t < \bar{c}_t$, ensuring that precautionary saving $\bar{c}_t - \hat{c}_t$ remains strictly positive over the extrapolation domain.

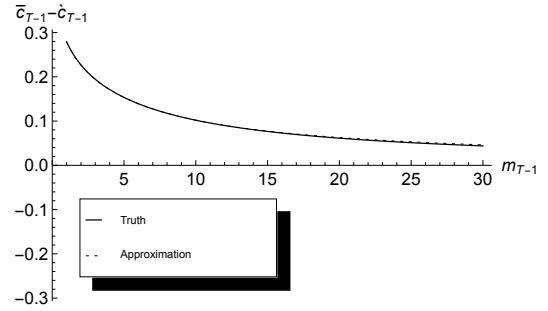
Given $\bar{\chi}$, the consumption function can be recovered from

$$\hat{c}_t = \bar{c}_t - \frac{\overbrace{1}^{\equiv \bar{\omega}_t}}{1 + \exp(\bar{\chi}_t)} \dot{h}_t \underline{c}_t. \tag{15}$$

Thus, the procedure is to calculate $\bar{\chi}_t$ at the points μ_t corresponding to the log of the \dot{m}_t points defined above, and then using these to construct an interpolating approximation $\hat{\chi}_t$ from which we indirectly obtain our approximated consumption rule \hat{c}_t by substituting $\hat{\chi}_t$ for $\bar{\chi}$ in equation (15).

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.’

Results are shown in Figure 3; a reader with very good eyesight might be able to detect the barest hint of a discrepancy between the Truth and the Approximation at the far righthand edge of the figure – a stark contrast with the calamitous divergence evident in Figure 1.

Figure 3: Extrapolated \hat{c}_{t-1} Constructed Using the Method of Moderation

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 $c_t = \Phi \cdot c_{t-1}$ to calculate the value function in period $T - 1$:

$$\begin{aligned}
 \bar{v}_{T-1}(m_{T-1}) &\equiv u(c_{T-1}) + \beta u(c_T) \\
 &= u(c_{T-1}) (1 + \beta \Phi^{1-\rho}) \\
 &= u(c_{T-1}) (1 + \Phi/R) \\
 &= u(c_{T-1}) \underbrace{\text{PDV}_{T-1}^T(c)/c_{T-1}}_{\equiv C_{T-1}^T}
 \end{aligned} \tag{16}$$

where $C_t^T = \text{PDV}_t^T(c)/c_t$ is the present discounted value of consumption, normalized by current consumption. Using the fact

demonstrated in Carroll and Shanker [6] that $C_t^T = \underline{\kappa}_t^{-1}$ ¹⁴, a similar function can be constructed recursively for earlier periods, yielding the general expression

$$\begin{aligned}\bar{\mathbf{v}}_t(m_t) &= \mathbf{u}(\bar{\mathbf{c}}_t(m_t))C_t^T \\ &= \mathbf{u}(\bar{\mathbf{c}}_t)\underline{\kappa}_t^{-1} \\ &= \mathbf{u}((\ddot{m}_t + \ddot{h}_t)\underline{\kappa}_t)\underline{\kappa}_t^{-1} \\ &= [(\ddot{m}_t + \ddot{h}_t)^{1-\rho}/(1-\rho)] \cdot [\underline{\kappa}_t^{1-\rho} \cdot \underline{\kappa}_t^{-1}] \\ &= \mathbf{u}(\ddot{m}_t + \ddot{h}_t)\underline{\kappa}_t^{-\rho}.\end{aligned}\quad (17)$$

This can be transformed as

$$\begin{aligned}\bar{\lambda}_t &\equiv ((1-\rho)\bar{\mathbf{v}}_t)^{1/(1-\rho)} \\ &= c_t(C_t^T)^{1/(1-\rho)} \\ &= (\ddot{m}_t + \ddot{h}_t)\underline{\kappa}_t^{-\rho/(1-\rho)}\end{aligned}\quad (18)$$

We apply the same transformation to the value function for the problem with uncertainty (the “realist’s” problem):

$$\hat{\lambda}_t = ((1-\rho)\hat{\mathbf{v}}_t(m_t))^{1/(1-\rho)} \quad (19)$$

and an excellent approximation to the value function can be obtained by calculating the values of $\hat{\lambda}$ 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 $\bar{\lambda}$ function for the optimist’s problem is an upper bound for the λ function in the presence of uncertainty, and the value function for the pessimist is a lower bound. Analogously to (13), define an upper-case

$$\hat{\Omega}_t(\mu_t) = \left(\frac{\bar{\lambda}_t(\underline{m}_t + e^{\mu_t}) - \hat{\lambda}_t(\underline{m}_t + e^{\mu_t})}{\ddot{h}_t \underline{\kappa}_t (C_t^T)^{1/(1-\rho)}} \right) \quad (20)$$

and an upper-case version of the $\bar{\chi}$ equation in (14):

$$\begin{aligned}\hat{\chi}_t(\mu_t) &= \log \left(\frac{1 - \hat{\Omega}_t(\mu_t)}{\hat{\Omega}_t(\mu_t)} \right) \\ &= \log \left(1/\hat{\Omega}_t(\mu_t) - 1 \right)\end{aligned}\quad (21)$$

and if we approximate these objects then invert them (as above with the $\bar{\omega}$ and $\bar{\chi}$ functions) we obtain a very high-quality approximation to our inverted value function at the same points for which we have our approximated value function:

$$\hat{\lambda}_t = \bar{\lambda}_t - \overbrace{\left(\frac{1}{1 + \exp(\hat{\chi}_t)} \right)}^{=\hat{\Omega}_t} \ddot{h}_t \underline{\kappa}_t (C_t^T)^{1/(1-\rho)} \quad (22)$$

¹⁴Under perfect foresight with time-invariant β and R , consumption grows at the constant gross rate Φ , so $\mathbf{c}_{t+n} = \mathbf{c}_t \Phi^n$. Discounting each term by R^{-n} yields the present discounted value ratio: $(PDV_t^T(\mathbf{c})/\mathbf{c}_t) = \sum_{n=0}^{T-t} (\Phi/R)^n$. Hence $C_t^T = \sum_{n=0}^{T-t} (\Phi/R)^n = \underline{\kappa}_t^{-1}$. When working with normalized variables, replace \mathbf{c} with c ; the identity remains unchanged.

from which we obtain our approximation to the value function as $\hat{\mathbf{v}}_t = \mathbf{u}(\hat{\lambda}_t)$.

5 EXTENSIONS

5.1 A Tighter Upper Bound

Carroll and Shanker [6] derives an upper limit $\bar{\kappa}_t$ ¹⁵ for the MPC as m_t approaches its lower bound, extending the explicit limiting MPC formulas established in buffer-stock theory [12]. Using this fact plus the strict concavity of the consumption function yields the proposition that

$$\hat{\mathbf{c}}_t(\underline{m}_t + \ddot{m}_t) < \bar{\kappa}_t \ddot{m}_t. \quad (23)$$

The solution method described above does not guarantee that approximated consumption will respect this constraint between gridpoints, and a failure to respect the constraint 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.

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

$$\begin{aligned}(\ddot{m}_t^* + \ddot{h}_t) \underline{\kappa}_t &= \bar{\kappa}_t \ddot{m}_t^* \\ \ddot{m}_t^* &= \frac{\underline{\kappa}_t \ddot{h}_t}{\bar{\kappa}_t - \underline{\kappa}_t} \\ m_t^* &= -\underline{h}_t + \frac{\underline{\kappa}_t (\bar{h}_t - \underline{h}_t)}{\bar{\kappa}_t - \underline{\kappa}_t},\end{aligned}\quad (24)$$

we want to construct a consumption function for $m_t \in (\underline{m}_t, m_t^*]$ that respects the tighter upper bound:

$$\begin{aligned}\ddot{m}_t \underline{\kappa}_t &< \hat{\mathbf{c}}_t(\underline{m}_t + \ddot{m}_t) < \bar{\kappa}_t \ddot{m}_t \\ \ddot{m}_t (\bar{\kappa}_t - \underline{\kappa}_t) &> \bar{\kappa}_t \ddot{m}_t - \hat{\mathbf{c}}_t(\underline{m}_t + \ddot{m}_t) > 0 \\ 1 &> \left(\frac{\bar{\kappa}_t \ddot{m}_t - \hat{\mathbf{c}}_t(\underline{m}_t + \ddot{m}_t)}{\ddot{m}_t (\bar{\kappa}_t - \underline{\kappa}_t)} \right) > 0.\end{aligned}\quad (25)$$

Again defining $\mu_t = \log \ddot{m}_t$, the object in the middle of the inequality is

$$\underline{\omega}_t(\mu_t) \equiv \frac{\bar{\kappa}_t - \hat{\mathbf{c}}_t(\underline{m}_t + e^{\mu_t})e^{-\mu_t}}{\bar{\kappa}_t - \underline{\kappa}_t}. \quad (26)$$

As m_t approaches $-\underline{m}_t$, $\underline{\omega}_t(\mu_t)$ converges to 0, while as m_t approaches $+\infty$, $\underline{\omega}_t(\mu_t)$ approaches 1.

As before, we can derive an approximated consumption function; call it $\check{\mathbf{c}}_t$. This function will clearly do a better job approximating the consumption function for low values of m_t while

¹⁵This represents the MPC when market resources approach the natural borrowing constraint, where the consumer faces maximum uncertainty about future income. In the Friedman-Muth process [26], φ is the probability that the transitory shock equals zero: $\varphi = \Pr(\xi = 0)$, representing the unemployment probability. Equivalent expressions: (i) backward recursion $\bar{\kappa}_t = 1 - \varphi^{1/\rho} \frac{\Phi}{R} (1 + \bar{\kappa}_{t+1})$ with terminal condition $\bar{\kappa}_T = 1$; (ii) finite-horizon forward sum $\bar{\kappa}_t = 1 - \varphi^{1/\rho} \frac{\Phi}{R} \sum_{n=0}^{T-t} (\varphi^{1/\rho} \frac{\Phi}{R})^n$; (iii) closed-form infinite-horizon limit $\bar{\kappa}_\infty = 1 - \varphi^{1/\rho} \frac{\Phi}{R}$.

the previous approximation will perform better for high values of m_t .

For middling values of m it is not clear which of these functions will perform better. However, an alternative is available which performs well. Define the highest gridpoint below m_t^* as \hat{m} and the lowest gridpoint above m_t^* as \check{m} . Then there will be a unique interpolating polynomial that matches the level and slope of the consumption function at these two points. Call this function $\check{c}_t(m)$.

Using indicator functions that are zero everywhere except for specified intervals,

$$\begin{aligned} \mathbf{1}_{\text{Lo}}(m) &= 1 \text{ if } m \leq \hat{m} \\ \mathbf{1}_{\text{Mid}}(m) &= 1 \text{ if } \hat{m} < m < \check{m} \\ \mathbf{1}_{\text{Hi}}(m) &= 1 \text{ if } \check{m} \leq m \end{aligned} \quad (27)$$

we can define a well-behaved approximating consumption function

$$\check{c}_t = \mathbf{1}_{\text{Lo}}\check{c}_t + \mathbf{1}_{\text{Mid}}\check{c}_t + \mathbf{1}_{\text{Hi}}\check{c}_t. \quad (28)$$

This just says that, for each interval, we use the approximation that is most appropriate. The function is continuous and once-differentiable everywhere, and is therefore well behaved for computational purposes.

We now construct an upper-bound value function implied for a consumer whose spending behavior is consistent with the refined upper-bound consumption rule.

For $m_t \geq m_t^*$, this consumption rule is the same as before, so the constructed upper-bound value function is also the same. However, for values $m_t < m_t^*$ matters are slightly more complicated.

Start with the fact that at the cusp point,

$$\begin{aligned} \bar{v}_t(m_t^*) &= \mathbf{u}(\bar{c}_t(m_t^*))C_t^T \\ &= \mathbf{u}(\check{m}_t^*\bar{\kappa}_t)C_t^T. \end{aligned} \quad (29)$$

But for *all* m_t ,

$$\bar{v}_t(m_t) = \mathbf{u}(\bar{c}_t(m_t)) + \mathbf{w}(m_t - \bar{c}_t(m_t)), \quad (30)$$

and we assume that for the consumer below the cusp point consumption is given by $\bar{\kappa}_t\check{m}_t$ so for $m_t < m_t^*$

$$\bar{v}_t(m_t) = \mathbf{u}(\bar{\kappa}_t\check{m}_t) + \mathbf{w}((1 - \bar{\kappa}_t)\check{m}_t), \quad (31)$$

which is easy to compute because $\mathbf{w}(a_t) = \beta\bar{v}_{t+1}(a_t\mathcal{R}_{t+1} + 1)$ where \bar{v}_t 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.) At the gridpoints defined by the solution of the consumption problem can then construct

$$\bar{\lambda}_t(m) = ((1 - \rho)\bar{v}_t(m))^{1/(1-\rho)} \quad (32)$$

which yields the appropriate vector for constructing \bar{X} and $\bar{\Omega}$. The rest of the procedure is analogous to that performed for the consumption rule and is thus omitted for brevity.

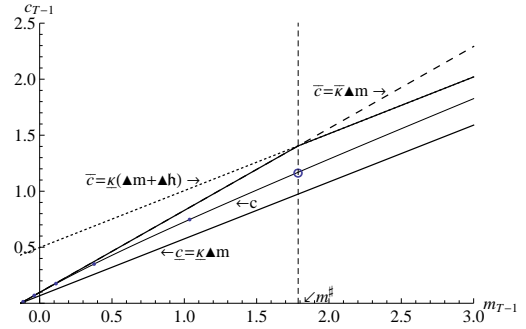


Figure 4: A Tighter Upper Bound

5.2 Hermite Interpolation

The numerical accuracy of the method of moderation depends critically on the quality of function approximation between gridpoints [27]. Our bracketing approach complements work that bounds numerical errors in dynamic economic models [28]. Although linear interpolation that matches the level of \hat{c} at the gridpoints is simple, Hermite interpolation [29, 30, 31] offers a considerable advantage. By matching both the level and the derivative of the \hat{c}_t function at the gridpoints, the consumption rule derived from such interpolation numerically satisfies the Euler equation at each gridpoint for which the problem has been solved [32, 33].

The theoretical foundation for this approach rests on the moderation ratio $\bar{\omega}$. This ratio captures how far the realist’s consumption lies between the perfect-foresight upper bound and the pessimistic lower bound. Since its log-gap argument μ moves with cash-on-hand relative to human wealth, the derivative measures how quickly the realist closes this gap as available resources shift:

$$\bar{\omega}_t^\mu = \frac{\check{m}_t(\kappa_t - \partial\hat{c}/\partial m_t)}{\kappa_t\check{h}_t}. \quad (33)$$

For numerical stability and interpretation on an unbounded scale, we apply the transformation defined in (14) to the moderation ratio. The derivative of this transformation is:

$$\chi_t^\mu = \frac{\bar{\omega}_t^\mu}{(\bar{\omega}_t - 1)\bar{\omega}_t}. \quad (34)$$

This expression provides the slope data required for cubic Hermite interpolation.¹⁶ To recover the realist consumption function from the interpolated transformation, we use the relationship established in (15). To recover the realist marginal propensity to consume, we differentiate with respect to the transformation:

¹⁶For cubic Hermite interpolation of the transformed moderation ratio, use node values from the transformation and node slopes χ_t^μ . For improved shape preservation, a monotone cubic Hermite scheme [34] can be used, where theoretical slopes serve as targets that may be adjusted to enforce monotonicity.

$$\frac{\partial \hat{c}_t}{\partial m_t} = \kappa_t \left(1 - \frac{\ddot{h}_t}{\dot{m}_t} \chi_t^\mu \right). \quad (35)$$

Since we have χ_t^μ from the Hermite interpolation, we can compute the realist MPC directly from the interpolated transformation derivatives. This expression reveals that the realist MPC is moderated away from κ_t by the transformation derivative, weighted by the ratio of human wealth to excess market resources. When uncertainty is low, the transformation derivative approaches zero and the realist MPC approaches the optimistic benchmark. Results showing identical MPCs under certain conditions [35] help explain why simple linear bounds work so well in many cases.

We can apply analogous techniques to the value function. Under perfect foresight, consumption grows at a constant rate, making C_t^T constant. This implies that the inverse value function for the optimist has a constant slope with respect to cash-on-hand:

$$\begin{aligned} \bar{\lambda}_t^m &= (C_t^T)^{1/(1-\rho)} \kappa_t \\ &= \kappa_t^{-\rho/(1-\rho)}. \end{aligned} \quad (36)$$

The result in (36) has important implications for the structure of the value function.¹⁷

Consider the value analogue of the moderation ratio, which compares the realist's value to the optimist's. The derivative of this ratio with respect to the log-gap argument is:

$$\hat{\Omega}_t^\mu = \frac{\ddot{m}_t(\bar{\lambda}_t^m - \hat{\lambda}_t^m)}{\dot{h}_t \bar{\lambda}_t^m} \quad (37)$$

where $\bar{\lambda}_t^m = \kappa_t^{-\rho/(1-\rho)}$ from (36) and $\hat{\lambda}_t^m$ is the derivative of the realist's inverse value function.

Applying the same transformation to the value-based moderation ratio converts the bounded ratio into an unconstrained slope:

$$\hat{\chi}_t^\mu = \frac{-\hat{\Omega}_t^\mu}{\hat{\Omega}_t(1 - \hat{\Omega}_t)}. \quad (38)$$

Since $\hat{\lambda}$ and \hat{v} are functional inverses, their derivatives are linked by chain-rule relationships. The first derivative is:

$$\hat{\lambda}_t^m = ((1 - \rho)\hat{v}_t(m_t))^{-1/(1-\rho)} \hat{v}_t^m(m_t). \quad (39)$$

The first- and second-derivative connections are:

$$\begin{aligned} \hat{\lambda}_t^m &= \mathbf{u}'(\hat{\lambda}_t) \hat{\lambda}_t^m \\ \hat{v}_t^{mm} &= \mathbf{u}''(\hat{\lambda}_t) (\hat{\lambda}_t^m)^2 + \mathbf{u}'(\hat{\lambda}_t) \hat{\lambda}_t^{mmm}. \end{aligned} \quad (40)$$

Moreover, if we use the double-derivative calculated in (40) to produce a higher-order Hermite polynomial, our approximation

¹⁷This confirms that $\bar{\lambda}_t$ is linear in m and highlights the role of κ_t in scaling marginal utility in the perfect-foresight benchmark. The linearity property simplifies both theoretical analysis and numerical implementation.

will also match the marginal propensity to consume at the gridpoints.¹⁸

These results provide the theoretical foundation for constructing high-quality cubic Hermite interpolants that preserve both the economic structure and numerical accuracy of the model between gridpoints.

5.3 Stochastic Rate of Return

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 \mathbf{R}_t would be trivial. Allowing for a stochastic interest factor is less trivial.

The easiest case is where the interest factor is i.i.d.,

$$\log \mathbf{R}_{t+n} \sim \mathcal{N}(r + \pi - \sigma_r^2/2, \sigma_r^2) \quad \forall n > 0 \quad (41)$$

because in this case Samuelson [36], Merton [37, 38] showed that for a consumer without labor income (or with perfectly forecastable labor income) the consumption function is linear, with an MPC.¹⁹

$$\kappa_t = 1 - (\beta \mathbb{E}_t[\mathbf{R}_{t+1}^{1-\rho}])^{1/\rho} \quad (42)$$

and in this case the previous analysis applies once we substitute this MPC for the one that characterizes the perfect foresight problem without rate-of-return risk. The more realistic case where the interest factor has some serial correlation is more complex, and thus left for future work.

In principle, this refinement should be combined with the previous one; further exposition of this combination is omitted here because no new insights spring from the combination of the two techniques.

6 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. But many problems do have obvious upper and lower bounds, and in those cases (as in the consumption example used in the paper), the method may result in substantial improvements in accuracy and stability of solutions.

¹⁸This would guarantee that the consumption function generated from the value function would match both the level of consumption and the marginal propensity to consume at the gridpoints, making the numerical differences between the newly constructed consumption function and the highly accurate one constructed earlier negligible within the grid.

¹⁹The Merton-Samuelson rule (for iid returns and no labor income, or perfectly forecastable income) implies a linear consumption function $c(m) = \kappa m$, where κ is given by Eq. (42). This result extends naturally to our framework by substituting the stochastic-return MPC for the perfect foresight MPC. At high wealth, the consumption function becomes asymptotically linear as the precautionary motive vanishes [39]. See Carroll [40] for a detailed derivation.

REFERENCES

- [1] Hayne E. Leland. Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics*, 82(3):465–473, 1968. doi: 10.2307/1879518.
- [2] Agnar Sandmo. The Effect of Uncertainty on Saving Decisions. *Review of Economic Studies*, 37(3):353–360, 1970. doi: 10.2307/2296725.
- [3] Miles S. Kimball. Precautionary Saving in the Small and in the Large. *Econometrica*, 58(1):53–73, 1990. doi: 10.3386/w2848.
- [4] John Stachurski and Alexis Akira Toda. An Impossibility Theorem for Wealth in Heterogeneous-Agent Models. *Journal of Economic Theory*, 2019. Provides an explicit linear lower bound on consumption; Corrigendum 2020.
- [5] Qingyin Ma, John Stachurski, and Alexis Akira Toda. The Income Fluctuation Problem and the Evolution of Wealth. *Journal of Economic Theory*, 187:105003, 2020. doi: 10.1016/j.jet.2020.105003.
- [6] Christopher D. Carroll and Akshay Shanker. Theoretical Foundations of Buffer Stock Saving (Revise and Resubmit, Quantitative Economics). Revise and Resubmit, Quantitative Economics, 6 2024. URL <https://llorracc.github.io/BufferStockTheory/BufferStockTheory.pdf>.
- [7] Christopher D. Carroll. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics*, 112(1):1–55, 1997. doi: 10.1162/003355397555109.
- [8] Milton Friedman. *A Theory of the Consumption Function*. Princeton University Press, Princeton, NJ, 1957.
- [9] John F. Muth. Optimal Properties of Exponentially Weighted Forecasts. *Journal of the American Statistical Association*, 55(290):299–306, 1960.
- [10] Christopher D Carroll. Solving microeconomic dynamic stochastic optimization problems. techreport, Johns Hopkins University, 2020. URL <https://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs.pdf>.
- [11] Christopher D. Carroll. Precautionary Saving and the Marginal Propensity to Consume. Working Paper 8233, NBER, 2001.
- [12] Qingyin Ma and Alexis Akira Toda. A Theory of the Saving Rate of the Rich. *Journal of Economic Theory*, 192: 105193, 2021. doi: 10.1016/j.jet.2021.105193.
- [13] Christopher D Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006.
- [14] Francisco Barillas and Jesús Fernández-Villaverde. A Generalization of the Endogenous Grid Method. *Journal of Economic Dynamics and Control*, 31(8):2698–2712, 2007.
- [15] Thomas Hintermaier and Winfried Koeniger. The Method of Endogenous Gridpoints with Occasionally Binding Constraints among Endogenous Variables. *Journal of Economic Dynamics and Control*, 34(10):2074–2088, 2010. doi: 10.1016/j.jedc.2010.05.002.
- [16] Giulio Fella. A Generalized Endogenous Grid Method for Non-smooth and Non-concave Problems. *Review of Economic Dynamics*, 17(2):329–344, 2014. doi: 10.1016/j.red.2013.07.001.
- [17] Fedor Iskhakov, Thomas H. Jørgensen, John Rust, and Bertel Schjerning. The Endogenous Grid Method for Discrete–Continuous Dynamic Choice Models with (or without) Taste Shocks. *Quantitative Economics*, 8(2): 317–365, 2017. doi: 10.3982/qe643.
- [18] Matthew N. White. The Method of Endogenous Gridpoints in Theory and Practice. *Journal of Economic Dynamics and Control*, 60:26–41, 2015. doi: 10.1016/j.jedc.2015.08.001.
- [19] George Tauchen and Robert Hussey. Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models. *Econometrica*, 59(2):371–396, 1991. doi: 10.2307/2938261.
- [20] George Tauchen. Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, 20(2):177–181, 1986. doi: 10.1016/0165-1765(86)90168-0.
- [21] S. Rao Aiyagari. Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics*, 109(3): 659–684, 1994. doi: 10.2307/2118417.
- [22] Mark Huggett. The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969, 1993. doi: 10.1016/0165-1889(93)90024-m.
- [23] Stephen P. Zeldes. Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, 97(2):305–346, 1989. doi: 10.1086/261605.
- [24] Angus Deaton. Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–1248, 1991. doi: 10.2307/2938366.
- [25] Christopher D. Carroll and Miles S. Kimball. On the Concavity of the Consumption Function. *Econometrica*, 64 (4):981–992, 1996. doi: 10.2307/2171853.
- [26] Christopher D. Carroll and Patrick Toche. A Tractable Model of Buffer Stock Saving. *Social Science Research Network*, 2009. doi: 10.3386/w15265.
- [27] Manuel S. Santos. Accuracy of Numerical Solutions Using the Euler Equation Residuals. *Econometrica*, 68(6): 1377–1402, 2000. doi: 10.1111/1468-0262.00165.
- [28] Kenneth L. Judd, Lilia Maliar, and Serguei Maliar. Lower Bounds on Approximation Errors to Numerical Solutions of Dynamic Economic Models. *Econometrica*, 85(3): 991–1020, 2017. doi: 10.3982/ecta12791.
- [29] F. N. Fritsch and R. E. Carlson. Monotone Piecewise Cubic Interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246, 1980. doi: 10.1137/0717021.
- [30] Fred N. Fritsch and J. Butland. A Method for Constructing Local Monotone Piecewise Cubic Interpolants. *SIAM Journal on Scientific and Statistical Computing*, 5 (2):300–304, 1984. doi: 10.1137/0905021.
- [31] James M. Hyman. Accurate Monotonicity-Preserving Cubic Interpolation. *SIAM Journal on Scientific and Statistical Computing*, 4(4):645–654, 1983. doi: 10.1137/0904045.

- [32] L. M. Benveniste and J. A. Scheinkman. On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica*, 47(3):727–732, 1979. doi: 10.2307/1910417.
- [33] Paul Milgrom and Ilya Segal. Envelope Theorems for Arbitrary Choice Sets. *Econometrica*, 70(2):583–601, 2002.
- [34] Carl de Boor. *A Practical Guide to Splines*. Springer, revised edition, 2001.
- [35] Karsten O. Chipeniuk, Nets Hawk Katz, and Todd B. Walker. Households, Auctioneers, and Aggregation. *European Economic Review*, 140, 2021. doi: 10.1016/j.euroecorev.2021.103997.
- [36] Paul A. Samuelson. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics*, 51(3):239–246, 1969.
- [37] Robert C. Merton. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics*, 51(3):247–257, 1969.
- [38] Robert C. Merton. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, 3(4):373–413, 1971.
- [39] Jess Benhabib, Alberto Bisin, and Shenghao Zhu. Skewed Wealth Distributions: Theory and Empirics. *Journal of Economic Literature*, 56(4):1331–1376, 2018. doi: 10.1257/jel.20161390.
- [40] Christopher D Carroll. Crra utility with rate risk. techreport, Johns Hopkins University, 2020. Lecture notes.