Theoretical Foundations of Buffer Stock Saving

2024-06-12

Christopher D. Carroll[1] and Akshay Shanker[2]

Abstract

This paper builds foundations for rigorous and intuitive understanding of 'buffer stock' saving models (Bewley [1977]-like models with a wealth target), pairing each theoretical result with quantitative illustrations. After describing conditions under which a consumption function exists, the paper articulates stricter 'Growth Impatience' conditions that guarantee alternative forms of 'target' saving — either at the population level, or for individual consumers. Together, the numerical tools and analytical results constitute a comprehensive toolkit for understanding buffer stock models.

Keywords Precautionary saving, buffer stock saving, marginal propensity to consume, permanent income hypothesis, income fluctuation problem

JEL codes D81, D91, E21

Powered by Econ-ARK

```
   REMARK: https://econ-ark.org/materials/bufferstocktheory
Dashboard: Click 'Launch Dashboard' Button
     html: https://llorracc.github.io/BufferStockTheory/
      PDF: BufferStockTheory.pdf
   Slides: BufferStockTheory-Slides.pdf
   GitHub: https://github.com/llorracc/BufferStockTheory
```

The paper's results can be automatically reproduced using the Econ-ARK toolkit by executing the notebook; for reference to the toolkit itself see Acknowleding Econ-ARK. Thanks to the Consumer Financial Protection Bureau for funding the original creation of the Econ-ARK toolkit; and to the Sloan Foundation for funding Econ-ARK's extensive further development that brought it to the point where it could be used for this project. The toolkit can be cited with its digital object identifier, https://doi.org/10.5281/zenodo.1001067, as is done in the paper's own references as Christopher D. Carroll et al. [2018]. Thanks to Will Du, James Feigenbaum, Joseph Kaboski, Miles Kimball, Qingyin Ma, Misuzu Otsuka, Damiano Sandri, John Stachurski, Adam Szeidl, Alexis Akira Toda, Metin Uyanik, Mateo Velásquez-Giraldo, Weifeng Wu, Jiaxiong Yao, and Xudong Zheng for comments on earlier versions of this paper, John Boyd for help in applying his weighted contraction mapping theorem, Ryoji Hiraguchi for extraordinary mathematical insight that improved the paper greatly, David Zervos for early guidance to the literature, and participants in a seminar at the Johns Hopkins University, a presentation at the 2009 meetings of the Society of Economic Dynamics for their insights, and at a presentation at the Australian National University.

[1]Contact: ccarroll@jhu.edu, Department of Economics, 590 Wyman Hall, Johns Hopkins University, Baltimore, MD 21218, https://www.econ2.jhu.edu/people/ccarroll, and National Bureau of Economic Research. [2]Contact: akshay.shanker@unsw.edu.au, Department of Economics, University of New South Wales, Sydney, Australia and Australian Research Council (ARC) Center for Excellence in Population and Ageing Research (CEPAR).

List of Figures

List of Tables

1 Introduction

{sec:intro}

The precautionary motive to save springs from the fact that extra resources give a consumer more ability to buffer spending against income shocks.[1] A consumer who, in the absence of shocks, would be 'impatient' enough to want to spend down resources, will (when there are shocks to worry about) experience an intensifying precautionary motive as their buffering capacity shrinks. If resources fall far enough, the consequence may be to make 'prudence' (Kimball [1993]) strong enough to counterbalance impatience. A consumer whose behavior is governed by this competition between impatience and prudence has been described, starting with Deaton [1991], as engaging in 'buffer stock saving.'

The logic of buffer stock saving underpins key findings in heterogeneous-agent macroeconomics. For example, it can explain why, during the Great Recession, middle-class consumers cut spending more than the poor or the rich [Krueger et al., 2016]. Buffer stock saving also can explain why consumption growth tracks income growth over much of the life cycle [Carroll, 1997], rather than being determined solely by preferences and interest rates as Irving Fisher [1930] proposed (and as log-linearized Euler equations erroneously suggest (Carroll [2001])).

Despite the central role that buffer stock saving plays in HA-macro, the literature lacks a formal theory mapping the conditions under which such behavior emerges. This paper establishes the required foundations by explaining the circumstances under which preferences, income growth, uncertainty, and the interest rate imply the existence of buffer-stock saving 'targets' both at the individual level and in the aggregate.

We formalize buffer stock saving using the Friedman-Muth(-Zeldes) income fluctuation model, incorporating realistic transitory and permanent shocks [Friedman, 1957, Muth, 1960, Zeldes, 1989],[2] constant relative risk aversion (CRRA) utility, and a 'natural borrowing constraint' of the kind first employed by Zeldes [1989].[3,4] In an infinite-horizon (or a Blanchard [1985] perpetual youth) version of this framework, we establish and explain the economic implications of two main results.

Our first contribution identifies conditions under which non-degenerate[5] infinite-horizon solutions exist without any 'artificial' liquidity constraint. The literature since Bewley [1977] and proceeding all the way through recent contributions by Ma et al. [2020], **?** has imposed an artificial liquidity constraint strictly tighter than the 'natural'

[1] Carroll and Kimball [2007].

[2] By which we mean, calibrated to micro empirical evidence.

[3] A natural borrowing constraint is the maximum amount a consumer will be willing to borrow under any circumstances. See Carroll [1997] or Gourinchas and Parker [2002] for arguments that models with only 'natural' constraints (see below) match a wide variety of facts; for a model with explicit constraints that produces very similar results, see, e.g. Cagetti [2003]. They are analytically convenient as the consumption function becomes twice continuously differentiable and also becomes arbitrarily close (cf. Section 2.4.3) to less tractable models with artificial liquidity constraints.

[4] The model permits separate transitory and permanent shocks (*a la* Muth [1960]) and permanent growth in income, which a large empirical literature finds to be of dominant importance in microdata. For example, MaCurdy [1982]; Abowd and Card [1989]; Carroll and Samwick [1997]; Jappelli and Pistaferri 2000; et. seq. Much of the literature instead incorporates highly 'persistent' but not completely permanent shocks, but Daly et al. [2016] show that when measurement problems are handled correctly, admin data yield serial correlation coefficients $0.98 - 1.00$; and Hryshko and Manovskii [2020] suggests that survey data support the same conclusion.

[5] Real valued solutions with strictly positive consumption as the planning horizon becomes arbitrarily distant.

constraint because this guarantees that the dynamic programming problem has a compact stationary Bellman operator (in the sense articulated below), permitting the use of contracting mapping arguments to show existence (**?**). Our proof of the case without artificial constraints holds even if a compact stationary Bellman operator may not exist.

Once we have established the existence of a non-degenerate solution, the second (and main) contribution of the paper is to identify conditions under which buffer stock 'targets' exist, for individual consumers or in the aggregate. Later we show that the 'artificial' case is a limiting version of the natural case.

The existence of nondegenerate solutions and buffer stock targets turn out to depend on various 'patience' conditions (combinations of parameters like the time preference rate, relative risk aversion, the rate of return, the underlying growth rate of income and the nature of income growth's stochastic elements). The simplest of these conditions (and one familiar from a large literature) corresponds to what we call 'absolute patience': A consumer who exhibits 'absolute impatience' is someone who, in the absence of a precautionary motive, would want to move future resources to the present to achieve a higher level of current consumption, because their pure rate of time preference makes them discount the future more than the interest rate encourages them to save (when $(\mathsf{R}\beta) < 1$). But, as Szeidl [2013] and **?** point out, the $(\mathsf{R}\beta) < 1$ condition is neither necessary nor sufficient for the model to have a nondegenerate solution.

Our patience conditions are most easily understood by comparison with the perfect foresight case with no borrowing constraint. For any given planning horizon, there will be a strictly positive marginal propensity to consume that optimally allocates current wealth between the present and the future. Of course, extending the horizon increases the number of future periods over which a given level of wealth can be spread, increasing the appeal of saving. The marginal propensity to save (and consume) in the limit depends on the term $\mathbf{Þ} = (\mathsf{R}\beta)^{1/\gamma}$, which we designate as the consumer's rate of absolute patience. The rate of absolute patience is a central concept in our paper and can be understood as the rate of consumption growth for a perfect foresight consumer with relative risk aversion of γ. In the perfect foresight case without borrowing constraints, what we show is necessary for a non-degenerate solution is 'return impatience': $\mathbf{Þ} < \mathsf{R}$. This condition guarantees that as the horizon extends arbitrarily, the limiting solution is one in which the consumer is not so patient that (in the limit) they will allocate all their wealth to the future (consuming nothing now).[6]

The first main result of the paper is that for our model with stochastic income shocks, the condition required for nondegeneracy is (surprisingly) *weaker* than the condition required for a nondegenerate solution in the perfect foresight case: What is needed is 'weak return impatience.' If a consumer is weak return impatient, then even if their wealth barely exceeds the limits imposed by the natural borrowing constraint, they will consume enough to make their expected resources fall. In addition to weak return impatience, we show that in the presence of permanent growth, the standard '$\beta < 1$'

[6]This condition is *not* required for nondegeneracy when there is an artificial borrowing constraint, so long as the consumer would wish to bring future resources to the present by borrowing – which occurs if income growth exceeds the interest rate.

requirement must be modified to take account of income growth uncertainty; we describe the new requirement as requiring 'finite value of autarky' (where here we think of autarky as perpetual consumption of your permanent income). If weak return impatience and finite value of autarky are satisfied, then we show that a nondegenerate solution exists.

Our method of proof uses a novel argument by utilizing the upper and lower bounds of consumers' marginal propensities of consume (MPCs) to show *'per-period' Bellman operators* are well-defined contraction maps. The finite horizon value functions are then shown to converge to the nondegenerate solution.[7]

The introduction of uncertainty brings a precautionary motive that enhances the consumer's preference for saving. So it is a surprise that in the stochastic setting, the 'patience' condition required for nondegeneracy is weaker than that required for the perfect foresight unconstrained case. The reason has to do with the other aspect of the introduction of uncertainty, which is the natural borrowing constraint imposed by the requirement that debts be repaid. Just as in the perfect foresight case, if the consumer would want to borrow against future income that is growing faster than the rate at which it is discounted, the natural borrowing constraint does what the artificial borrowing constraint does in the perfect foresight case: It prevents too much borrowing.

Turning to our results on the buffer stock target, the requirement for the existence of an individual target is 'strong growth impatience,' which prevents 'normalized market resources' (the ratio of market resources to permanent income) from growing without bound. In particular, strong growth impatience requires that as a consumer's resources grow, eventually a point will come at which the expected ratio of a consumer's absolute patience to the uncertainty-adjusted growth of permanent income is less than one. Under the condition, as a consumer's normalized market resources approach infinity, the *expected ratio* of market resources–incorporating optimal saving–to the uncertain permanent income growth factor, must eventually becomes less than one. That is, normalized market resources eventually revert back toward a target.

A weaker requirement, 'growth impatience,' ensures the existence of an aggregate buffer stock target even when individual target ratios are unbounded. Growth impatience requires the *ratio* of absolute patience to the *expected* growth factor of permanent income to be less than one. But recall that at the individual level, the ratio of wealth to permanent income might grow either as a result of an increase in wealth or as the result of a decrease in permanent income. Idiosyncratic uncertainty in permanent income means that in any given year there will be a portion of consumers whose permanent income has declined. Such consumers wll experience an increase in the ratio of wealth to permanent income, even if their absolute wealth has stagnated or even fallen (but by less than the decline in permanent income).

Nonetheless, when 'growth impatience', holds the ratio of average market resources to average permanent income converges back to a target. As Harmenberg [2021] points out, a stationary distribution of market resources, weighted by permanent income still exists under growth impatience. We develop the insight by Harmenberg [2021] and

[7] They are a Cauchy sequence in a complete weighted-norm space, converging to a nondegenerate solution as the planning horizon recedes.

demonstrate that the contribution to aggregate consumption of consumers who accumulate unbounded resources diminishes as they receive a smaller and smaller measure of permanent income. Thus in the aggregate, even with a fixed aggregate interest rate that differs from the time preference rate, a small open economy populated by buffer stock consumers has a balanced growth path in which growth rates of consumption, income, and wealth match the exogenous growth rate of permanent income (equivalent, here, to productivity growth). In the terms of Schmitt-Grohé and Uribe [2003], buffer stock saving is an appealing method of 'closing' a small open economy, because it requires no *ad-hoc* assumptions. Not even liquidity constraints.

An interesting implication is that the consumption function exists (and is not degenerate) even when the (exogenous) growth rate of income exceeds the (exogenous) interest rate. Many economic models impose an interest rate greater than the growth rate because if that condition does not hold then the risk-neutral present discounted value of future income is infinite. The presence of the precautionary motive short-circuits this logic, and implies that even if *in a risk-neutral sense* human wealth is infinite, the limiting solution is not $c = \infty$ as the horizon extends arbitrarily.

In the final section of the paper we show that for a consumer to have a non-degenerate value function in the limit, absolute patience cannot exceed *both* market returns and the growth rate of income – growth impatience must hold when the return impatience fails and vice-versa. The 'discipline' on patience is enforced by requiring time discounting, which ensures consumers' discounted sum of pay-offs remain bounded as the planning horizon recedes. For instance, take an excessively patient consumer who is not return impatient. A nondegenrate value function for the consumer requires a relatively high growth rate of income via growth impatience. Intuitively, a high enough growth rate of income gives an infinite net present 'market value' of human wealth, which satisfies the need for high consumption growth and prevents consumption today from falling to zero despite high patience. Such consumers' marginal propensity to save limits to one as they become wealthier, yet they exhibit buffer stock behaviour (atleast in the aggregate data) because their income growth limits the growth of normalized market resources. On the other hand, a buffer stock target (in the aggregate) can fail to exist (say growth impatience fails due to low income growth) only when market returns exceed absolute patience. In this case, the limiting MPC converges to be strictly less than one, but income growth may be so low that market resources grow at a high enough rate (due to the precautionary motive) and expected normalized resources diverge.

Relationship to Literature Buffer stock saving behaviour was recognized by Friedman [1957] and formally introduced by Carroll [1997] to account for consumption and income patterns in the data. The concept is closely linked to precautionary saving [Zeldes, 1989, Deaton, 1991, **?**] and the literature [Carroll and Kimball, 1996, **?**, Carroll et al., 2019, **?**] now provides theoretical results on how risk affects the consumption function in the presence of the precautionary motive. The class of models we use to study buffer stock saving, income fluctuation problems, are now pervasive, though the foundational contributions include Bewley [1977], **?**, Zeldes [1989], Deaton [1991], **?** and Aiyagari [1994]. Amongst this literature Zeldes [1989] was the first to calibrate a quantitatively

plausible example of permanent and transitory shocks and argue that the natural borrowing constraint was a quantitatively plausible alternative to 'artificial' or 'ad hoc' borrowing constraints.[8] The natural borrowing limit was also described by Aiyagari [1994], but implications for existence not discussed.[9]

On the technical front, traditional Bellman iteration approaches to showing existence rely on bounded pay-offs [Stokey et al., 1989].[10] Overcoming these restrictions to allow unbounded pay-offs, the literature on the one hand emphasized time iteration operators defined by Euler equations [Deaton, 1991, Li and Stachurski, 2014, Ma et al., 2020] and transformations of the Bellman equation [**?**]. The results by Ma et al. [2020], **?** are the most general we are aware of, and can be specialized to show existence in a model with the rate of return and discount factor shock structure arising from permanent and transitory shocks (once the model is normalized). However, the cited approaches impose an artificial liquidity constraint, thus cannot be applied here. Moreover, our growth and return patience concern economic mechanisms (rather than general assumptions) that arise in the presence of permanent income uncertainty and growth and to the best of our knowledge, have not been explored elsewhere.

Our approach to constructing the weighted-norm space of value functions uses results on unbounded dynamic programming by [Boyd, 1990].[11] However, our use of marginal propensities to consume to construct per-period bounds on the Bellman operator are novel. In a more abstract setting, our proofs address issues of compactness and continuity similar to that tackled by **?** and **?**. In contrast to the abstract methods, the proofs in this paper are directly applied to an income fluctuation problem and have the advantage that they employ standard concepts, such as Bellman iteration, that are straightforward to verify in practice.

Finally, our discussion on aggregate growth rates builds on Szeidl [2013] and Harmenberg [2021] who give results on the existence and convergence of stationary wealth distributions for the model presented here. Ma et al. [2020] also give results on stationarity, under the restrictions mentioned above. While conditions for stationarity resemble growth impatience and strong growth impatience our objective is to establish existence of stable buffer stock targets, which have empirical relevance, not prove stochastic stability.

---

[8] The same (numerical) point applies for infinite horizon models (calibrated to actual empirical data on household income dynamics); cf. Carroll [1992].

[9] Income fluctuation problems do not require existence of stable buffer stock targets, though such points will often exist.

[10] The CRRA utility function does not satisfy Bewley's assumption that u(0) is well-defined, or that $u'(0)$ is well-defined and finite. Our framework differs from Schectman and Escudero [1977] in that they impose liquidity constraints and positive minimum income. It differs from Deaton [1991] because liquidity constraints are absent; there are separate transitory and permanent shocks (*a la* Muth [1960]); and the transitory shocks here can occasionally cause income to reach zero. Similar restrictions are made in the well known papers by Scheinkman and Weiss [1986], Clarida [Clarida, 1987], and others Chamberlain and Wilson [2000]. For a related continuous-t model, see Toche [2005]. Alvarez and Stokey [1998] relaxed the bounds on the return function, but they address only the deterministic case with compact valued action sets. Matkowski and Nowak [2011] assume a framework with compact action sets, and real-valued pay-offs, which cannot handle CRRA utility unbounded below. See Yao [2012] for a detailed discussion of the reasons the existing literature up through Matkowski and Nowak [2011] cannot handle the problem described here. Martins-da Rocha and Vailakis [2010] provide a correction to Rincón-Zapatero and Rodríguez-Palmero [2003], but only addresses the deterministic case.

[11] Alvarez and Stokey [1998] showed how the approach could be used to address the homogeneous case (of which CRRA is an example) in a deterministic framework; later, Durán [2003] showed how to extend the Boyd [1990] approach to the stochastic case. See also exposition by **?**, Ch. 12.

2 Theoretical Foundations

This section formalizes our problem and derives formulae, for any period t earlier than the terminal period T, for the maximum and minimum MPCs as wealth approaches zero and infinity (these formulae are derived recursively backward from T). If the environment is that of an infinite-horizon 'income fluctuation problem,' our formulae yield the limiting upper and lower bounds of the nondegenerate stationary solution. {sec:Foundations}

2.1 Setup

{subsec:Setup}

We start by stating the problem with permanent income growth in levels and then normalize by permanent income.

Our time index t can take on values in $\{T, T-1, T-2, \dots\}$. We assume that our consumer has a Constant Relative Risk Aversion (CRRA) per-period utility function, $\mathrm{u}(c) = \frac{c^{1-\gamma}}{1-\gamma}$, where $\gamma > 1$. β is the (strictly positive) discount factor. In each period, the consumer faces income shocks, with the permanent shock $\psi_t \in \mathbb{R}_{++}$ and the transitory shock by $\boldsymbol{\xi}_t \in \mathbb{R}_{+}$.

In each t, value will be a function of 'market resources' $\boldsymbol{m}_t$ and permanent income $\boldsymbol{p}_t$, with $\boldsymbol{m}_t$ and $\boldsymbol{p}_t$ strictly positive real numbers ($\{\boldsymbol{m}_t, \boldsymbol{p}_t\} \in \mathbb{R}_{++} \times \mathbb{R}_{++}$).

Letting $\mathbf{v}_{T+1} = 0$, the finite-horizon value functions are recursively defined by:[12]

$$\mathbf{v}_t(\boldsymbol{m}_t, \boldsymbol{p}_t) = \max_{0 \leq \boldsymbol{c}_t \leq \boldsymbol{m}_t} \mathrm{u}(\boldsymbol{c}_t) + \beta \mathbb{E}_t \mathbf{v}_{t+1}(\boldsymbol{m}_{t+1}, \boldsymbol{p}_{t+1}) \qquad (\mathcal{P}_L)$$ {eq:levelRecProblem

where $\boldsymbol{c}_t$ is the level of consumption at time t. We assume the consumer cannot die in debt:

$$\boldsymbol{c}_T \leq \boldsymbol{m}_T. \qquad (1)$$ {eq:NoDebtAtDeath

For maximal clarity, we separately describe every step in the dynamic budget evolution that determines next period's $\boldsymbol{m}_{t+1}$ from this period's $\boldsymbol{m}_t$ and choice of $\boldsymbol{c}_t$:[13,14]

[12] Notation throughout follows guidelines specified for the Econ-ARK toolkit; see Notation in the ARK for rationales and details. Consequently, there is an exact mapping between objects in the paper and objects in the code.

[13] The steps are broken down also so that the notation of the paper will correspond exactly to the variable names in the toolkit, because it is required for solving life cycle problems.

[14] Allowing a stochastic interest factor would complicate the notation but not affect the points we want to address; however, see Benhabib et al. [2015], Ma and Toda [2020], and Ma et al. [2020] for the implications of capital income risk for the distribution of wealth and other interesting questions not considered here.

$$
\begin{aligned}
\boldsymbol{a}_t &= \boldsymbol{m}_t - \boldsymbol{c}_t \\
\boldsymbol{k}_{t+1} &= \boldsymbol{a}_t \\
\boldsymbol{p}_{t+1} &= \boldsymbol{p}_t \underbrace{\mathcal{G}\psi_{t+1}}_{:=\tilde{\mathcal{G}}_{t+1}} \\
\boldsymbol{m}_{t+1} &= \underbrace{\mathsf{R}\boldsymbol{k}_{t+1}}_{:=\boldsymbol{b}_{t+1}} + \underbrace{\boldsymbol{p}_{t+1}\boldsymbol{\xi}_{t+1}}_{:=\boldsymbol{y}_{t+1}}.
\end{aligned}
$$

The consumer's assets at the end of t, $\boldsymbol{a}_t$, translate one-for-one into capital $\boldsymbol{k}_{t+1}$ at the beginning of the next period. In turn, $\boldsymbol{k}_{t+1}$ is augmented by a fixed interest factor R to become the consumer's financial ('bank') balances $\boldsymbol{b}_{t+1} = \mathsf{R}\boldsymbol{k}_{t+1}$.[15] 'Market resources,' $\boldsymbol{m}_{t+1}$, are the sum of financial wealth $\mathsf{R}\boldsymbol{k}_{t+1}$ and noncapital income $\boldsymbol{y}_{t+1} = \boldsymbol{p}_{t+1}\boldsymbol{\xi}_{t+1}$ (permanent noncapital income $\boldsymbol{p}_{t+1}$ multiplied by the transitory-income-shock factor $\boldsymbol{\xi}_{t+1}$ described below). Permanent noncapital income $\boldsymbol{p}_{t+1}$ is derived from $\boldsymbol{p}_t$ by application of a growth factor $\mathcal{G}$,[16] modified by the permanent income shock ψ_{t+1}, and the resulting idiosyncratic growth factor for permanent income is compactly written as $\tilde{\mathcal{G}}_{t+1}$.

The finite-horizon problems furnish a sequence of value functions $\{\mathbf{v}_T, \mathbf{v}_{T-1}, \ldots, \mathbf{v}_{T-n}\}$ and associated consumption functions $\{\mathbf{c}_T, \mathbf{c}_{T-1}, \ldots, \mathbf{c}_{T-n}\}$. We define the infinite-horizon solution as the (limiting) first-period solution to a sequence of finite-horizon problems as the first period becomes arbitrarily distant from the terminal period (that is, as $n \to \infty$). An infinite-horizon solution will be 'nondegenerate' (or, 'sensible') if the limiting consumption function, denoted by $\mathbf{c}(\boldsymbol{m}, \boldsymbol{p}) = \lim_{n\to\infty} \mathbf{c}_{T-n}(\boldsymbol{m}, \boldsymbol{p})$, is neither $\mathbf{c} = 0$ everywhere (for all $(\boldsymbol{m}, \boldsymbol{p})$) nor $\mathbf{c} = \infty$ everywhere.[17]

The following assumption defines the income process.

{ass:shocks}

Assumption I.1. *(Friedman-Muth Income Process). The following holds for all t:*

1. *For $n > 0$, the permanent shocks ψ_t are independently and identically distributed (iid), with $\mathbb{E}_{t-n}[\psi_t] = 1$ and support $[\underline{\psi}, \overline{\psi}]$, where $0 < \underline{\psi} \leq 1$ and $1 \leq \overline{\psi} < \infty$.*

2. *The transitory shocks satisfy:*

$$
\boldsymbol{\xi}_t = \begin{cases} 0 & \text{with probability } q > 0 \\ \theta_t/(1-q) & \text{with probability } (1-q) \end{cases} \tag{2}
$$

{eq:TranShkDef}

[15] See below for a brief discussion of the case where returns are stochastic.

[16] A time-varying $\mathcal{G}$ has straightforward consequences for the analysis below; this is an option allowed for in the HARK toolkit.

[17] The traditional approach to study recursive problems defines an infinite-horizon maximization problem over stochastic recursive sequences [?]. Using the Bellman Principle of Optimality, the definition of degenerate solutions we use here can be shown to be equivalent to the optimal solution of a dynamic stochastic sequence problem (see Appendix A.1.3). The framing we use allows for a more direct link to life-cycle models (see Gourinchas and Parker [2002] for an instance where buffer stock saving is discussed in the context of a life-cycle model).

where θ_t is an iid random variable with $\mathbb{E}_{t-n}[\theta_t] = 1$ and $\underline{\theta} \leq \theta_t \leq \bar{\theta}$, where $\underline{\theta} > 0$ and $\underline{\theta} \leq 1 \leq \bar{\theta} < \infty$.

Following Zeldes [1989], the income process incorporates a small probability q that income will be zero (a 'zero-income event'). At date $T-1$, the (strictly positive) probability q of zero income in period T will prevent the consumer from spending all resources, because saving nothing would mean arriving in the following period with zero bank balances and thus facing the possibility of being required to consume 0, which would yield utility of $-\infty$. This logic holds recursively from $T-1$ back, so the consumer will never spend everything, giving rise to what Aiyagari [1994] dubbed a 'natural borrowing constraint.'[18] (Formally, this establishes that the upper-bound constraint on consumption in the problem ($\mathcal{P}_L$) will not bind.)

2.1.1 Normalized Problem

{subsubsec:ratio}

Let nonbold variables be the boldface counterpart normalized by $\boldsymbol{p}_t$, allowing us to reduce the number of states from two ($\boldsymbol{m}$ and $\boldsymbol{p}$) to one ($m = \boldsymbol{m}/\boldsymbol{p}$). Now, in a one-time deviation from the notational convention established in the last sentence, define nonbold 'normalized value' not as $\boldsymbol{v}_t/\boldsymbol{p}_t$ but as $v_t = \boldsymbol{v}_t/\boldsymbol{p}_t^{1-\gamma}$, because this allows us to write nonbold v_t to denote the 'normalized value function':

$$
\begin{aligned}
\mathrm{v}_t(m_t) = \max_{c_t \leq m_t} \ & \mathrm{u}(c_t) + \beta \mathbb{E}_t[\tilde{\mathcal{G}}_{t+1}^{1-\gamma} \mathrm{v}_{t+1}(m_{t+1})], \qquad m_t \in \mathbb{R}_{++} \\
& \text{s.t.} \\
a_t &= m_t - c_t \\
k_{t+1} &= a_t / \tilde{\mathcal{G}}_{t+1} \\
b_{t+1} &= k_{t+1} \mathsf{R} \\
m_{t+1} &= b_{t+1} + \boldsymbol{\xi}_{t+1}.
\end{aligned}
\tag{$\mathcal{P}_N$}
$$

{eq:veqnNrmRecBel

(Appendix A.1.1 explains how the solution to the original problem in levels can be recovered from the normalized problem.)

The time t normalized consumption *policy function* for the finite-horizon problem, c_t, is defined by:

$$
\mathrm{c}_t(m_t) = \arg\max_{c_t \leq m_t} \ \mathrm{u}(c_t) + \beta \mathbb{E}_t[\tilde{\mathcal{G}}_{t+1}^{1-\gamma} \mathrm{v}_{t+1}(m_{t+1})] \tag{3}
$$

{eq: cfunceq1}

The normalized problem's first order condition becomes:

$$
c_t^{-\gamma} = \mathsf{R} \beta \mathbb{E}_t[\tilde{\mathcal{G}}_{t+1}^{-\gamma} c_{t+1}^{-\gamma}]. \tag{4}
$$

{eq:scaledeuler}

We now formally define the limiting nondegenerate solution to the normalized problem, letting n index the planning horizon.

{def:nondegenerac

[18] We specify zero as the lowest-possible-income event without loss of generality, see for example the discussion by Aiyagari [1994].

Definition 1. *Problem $\mathcal{P}_N$ has a nondegenerate solution if there exists* c*, with* $\mathrm{c}\colon \mathbb{R}_{++} \to \mathbb{R}_{++}$*, such that* $\mathrm{c}(0) = 0$ *and:*

$$\mathrm{c}(m) = \lim_{n\to\infty} \mathrm{c}_{T-n}(m), \qquad m \in \mathbb{R}_{++}.$$

We will similarly use v without a subscript to refer to the pointwise limit of the value functions as the planning horizon recedes.

Having defined the limiting solution, we now explain why the standard dynamic programming approach for showing existence cannot be used here. Let $\mathbb{T}$ denote the mapping $\mathrm{v}_{t+1} \mapsto \mathrm{v}_t$ given by Problem $\mathcal{P}_N$:

$$\mathbb{T}\mathrm{v}_{t+1}(m) = \max_{c\in(0,m]} \left\{ \mathrm{u}(c) + \beta \mathbb{E} \tilde{\mathcal{G}}^{1-\gamma} \mathrm{v}_{t+1}(\tilde{\mathcal{R}}(m-c) + \boldsymbol{\xi}) \right\}, \quad m \in \mathbb{R}_{++}. \tag{5}$$ {eq:maintmap}

If $\mathbb{T}$ can be shown to be a contraction mapping operator, our finite-horizon value functions will converge to a nondegenerate solution. However, $\mathbb{T}$ will not a be well defined operator on a space of continuous functions due to the naturally arising liquidity constraint.

Remark 1. *(Challenges with Standard Dynamic Programming). Standard dynamic programming works by showing that* $\mathbb{T}$ *is a well-defined contraction map on a Banach space, implying the sequence of value functions given by Problem* $\mathcal{P}_N$ *will converge to a fixed point of* $\mathbb{T}$*, a non-degenerate solution. In our case, there are two challenges (further details in Appendix A.1.2). First, utility is unbounded, so we must construct a suitable metric space in which the sequential value functions remain bounded. Second, we cannot show* $\mathbb{T}$ *is a contraction mapping because we cannot immediately assert that* $\mathbb{T}$ *is a well-defined mapping on a suitable function space. Even if we verified that a maximizer exists for arbitrary continuous function* f*, that is,* $\mathrm{f} \mapsto \mathbb{T}\mathrm{f}$ *is well-defined, we cannot then assert* $\mathbb{T}$ *maps continuous functions to continuous functions since the feasible set on the RHS is not compact-valued. (Without compact valued feasible sets, we cannot use Berge's Theorem and* g*, with* $\mathrm{g} = \mathbb{T}\mathrm{f}$*, may not be continuous.)* **?** *show that the problem can be transformed to include zero consumption in the feasibility constraint; however, the operator continues to be ill-defined unless income has a strictly positive lower bound (or equivalently, there is an 'artificial' borrowing constraint is strictly tighter than the natural borrowing constraint).*[19]

{remark:stochdisc

Remark 2. *If we set* $q = 0$*, the normalized problem becomes a special case of the problem considered by Ma et al. [2020], with* $\tilde{\mathcal{R}}_{t+1} = \mathsf{R}/\tilde{\mathcal{G}}_{t+1}$ *corresponding to the stochastic rate of return on capital and* $\beta\tilde{\mathcal{G}}_{t+1}^{1-\gamma}$ *corresponding to the stochastic discount factor.*

[19] **?** generalize the requirement of continuity of feasibility correspondences to K-Inf-Compactness of the Bellman operator, yielding a mapping from semi-continuous to semi-continuous functions. **?** introduces a generalization, mild-Sup-compactness, which can be verified in the weak topology generated on the infinite dimensional product space of feasible random variables controlled by the consumers. Our approach, by contrast, has the advantage that it can be used to verify existence using more standard topological notions.

Notwithstanding remark 2, there are important economic consequences of the fact that in our problem $\tilde{\mathcal{R}}_{t+1} = \mathsf{R}/\tilde{\mathcal{G}}_{t+1}$ is tightly tied to 'normalized stochastic discount factor,' $\beta\tilde{\mathcal{G}}_{t+1}^{1-\gamma}$; we discuss these below.

2.2 Patience Conditions

{subsec:GICTheory

In order to have a central reference point for them, we now collect (without much explanation) conditions relating to growth factors and various different 'patience factors' that underpin results in the remainder of the paper. Assumptions L.1 - L.3 (finite value of autarky and return impatience) will be used to prove the existence of limiting solutions in Section 2.4, and Assumptions S.1 - S.2 (growth impatience and strong growth impatience) are required for existence of alternative definitions of a stable target buffer stock in Section 3.

We start by generalizing the standard $\beta < 1$ condition to our setting with permanent income growth and uncertainty.[20] The updated condition requires that the expected net discounted value of utility from consumption is finite under our definition of 'autarky' – where consumption is always equal to permanent income. A finite value of autarky helps guarantee that as the horizon extends, discounted value remains finite along *any* consumption path the consumer might choose. (See Appendix **??**).

{ass:FVAC}

Assumption L.1. *(Finite Value of Autarky).* $0 < \beta\mathcal{G}^{1-\gamma}\mathbb{E}(\psi^{1-\gamma}) < 1$.

The discount factor β defines the consumer's 'pure' rate of time preference – the relative weight of utility across time. The term 'patience' does not have a similarly clear definition in the literature. Part of our objective in this paper is to provide a taxonomy for each of various useful definitions of patience.

We start with 'absolute (im)patience.' We will say that an unconstrained perfect foresight consumer exhibits absolute impatience if they optimally choose to spend so much today that their consumption must decline in the future. The growth factor for consumption implied by the Euler equation of a perfect foresight model is $\boldsymbol{c}_{t+1}/\boldsymbol{c}_t = (\mathsf{R}\beta)^{1/\gamma}$,[21] which motivates our definition of an 'absolute patience factor' whose centrality (to everything that is to come later) justifies assigning to it a special symbol; we have settled on the archaic letter 'thorn':

$$\text{Þ} := (\mathsf{R}\beta)^{1/\gamma}. \tag{6}$$

{eq:APFac}

We will say that (in the perfect foresight problem) 'an absolutely impatient' consumer is one for whom $\text{Þ} < 1$; that is an absolutely impatient consumer prefers to consume more today than tomorrow (and vice versa for an 'absolutely patient' consumer, whose consumption will grow over time):

Assumption L.2. *(Absolute Impatience).* $\text{Þ} < 1$.

{ass:AIC}

[20] In light of Remark 2, Ma et al. [2020] Assumption 2.1 is a generalization of this discount condition, albeit in a context with artificial liquidity constraints.

[21] See (10) below.

Remark 3. *A consumer who is absolutely impatient, $\Ppatient < 1$, satisfies the standard impatience condition commonly used in the income fluctuation literature, $\beta\Rfree < 1$, which guarantees the existence of a stable distribution when there is no permanent income growth. However, as pointed out by Szeidl [2013] and* **?** *(henceforth, 'MST'), $\beta\Rfree < 1$ is not necessary for an infinite-horizon solution. The shock process of the normalized problem here can be seen as a special case of MST (albeit with their artificial liquidity constraint, which we drop).*

Recall now our earlier requirement that the limiting consumption function $\mathrm{c}(m)$ in our model must be 'sensible.' We will show below that for the perfect foresight unconstrained problem this requires

Assumption L.3. *(Return Impatience).* $\frac{\Ppatient}{\Rfree} < 1$. {ass:RIC}

Return impatience can be best understood as the tension between the income effect of capital income and substitution effect. As we show below, in the perfect foresight model, it is straightforward to derive the MPC out of overall (human plus nonhuman) wealth that would result in next period's wealth being identical to the current period's wealth. The answer turns out to be an MPC ('κ') of $\kappa = (1 - \Ppatient/\Rfree)$. The interesting point here is that κ depends both on our absolute patience factor $\Ppatient$ and on the return factor. This is the manifestation in this context of the interaction of the income effect (higher wealth yields higher interest income if $\Rfree > 1$) and the substitution effect (which we have already captured with $\Ppatient$).

See the discussion in section [wherever] for further elaboration.

Next, consider the weaker condition of a consumer whose absolute patience factor is suitably adjusted to take account of the probability of zero income is less than the market return.

Assumption L.4. *(Weak Return Impatience).* $\underbrace{\frac{(q\Rfree\beta)^{1/\gamma}}{\Rfree}}_{=\frac{q^{1/\gamma}\Ppatient}{\Rfree}} \leq 1$. {ass:WRIC}

This condition is 'weak' (relative to the plain return impatience) because the probability of the zero income events q is strictly less than 1. The role of q in this equation is related to the fact that a consumer with zero end-of-period assets today has a probability q of having no income and no assets to finance consumption (and $m_{t+1} = 0$ would yield negative infinite utility). In the case with no artificial constraint, our main results below show weak return impatience and finite value of autarky are sufficient to guarantee a 'sensible' (nondegenerate) solution. Moreover, weak return impatience is also necessary and cannot be relaxed further without an artificial liquidity constraint.

Review the remark under comment in latex. what do we mean by does not approach irrelevance? Should this even be a Remark?

Now that we have finished discussing the requirements for a nondegenerate solution, we turn to assumptions required for stability. We speak of a consumer whose absolute patience factor is less than the expected growth factor for their permanent income $\mathcal{G} = \mathbb{E}[\mathcal{G}\psi])$ as exhibiting 'growth impatience:'

Assumption S.1. *(Growth Impatience).* $\frac{\Ppatient}{\mathcal{G}} < 1$. {ass:GICRaw}

A final useful definition is 'strong growth impatience' which holds for a consumer for whom the expectation of the *ratio* of the absolute patience factor to the growth factor of permanent income is less than one,

Assumption S.2. *(Strong Growth Impatience).* $\mathbb{E}\left[\frac{\mathbf{\Phi}}{\mathcal{G}\psi}\right] = \frac{\mathbf{\Phi}}{\mathcal{G}}\mathbb{E}[\psi^{-1}] < 1.$ {ass:GICMod}

(The difference between growth impatience and strong growth impatience is that the first is the ratio of an expectation to an expectation, while the latter is the expectation of the ratio. With nondegenerate mean-one stochastic shocks to permanent income, the expectation of the ratio is strictly larger than the ratio of the expectations).

While neither growth impatience nor return impatience will by themselves be required for the existence of a limiting solution, the finite value of autarky condition stops individuals from becoming *both* growth and return patient.

{claim:noRICGIC

Claim 1. *If growth impatience fails ($\frac{\mathbf{\Phi}}{\mathcal{G}} > 1$) and return impatience fails ($\frac{\mathbf{\Phi}}{\mathsf{R}} > 1$), then finite value of autarky fails ($\beta\mathcal{G}^{1-\gamma}\mathbb{E}(\psi^{1-\gamma}) > 1$).*

Proof. Since $\frac{\mathbf{\Phi}}{\mathsf{R}} > 1$, we have:

$$\frac{\mathbf{\Phi}}{\mathsf{R}} = \frac{(\mathsf{R}\beta)^{\frac{1}{\gamma}}}{\mathsf{R}} > 1 \tag{7}$$

Next, multiplying both sides by $\mathsf{R}\mathcal{G}^{1-\gamma}$, we get

$$\beta\mathcal{G}^{1-\gamma}\mathsf{R}^{\frac{1}{\gamma}}\beta^{\frac{1-\gamma}{\gamma}} > \mathsf{R}\mathcal{G}^{1-\gamma} \Rightarrow \beta\mathcal{G}^{1-\gamma} > \left(\frac{\mathbf{\Phi}}{\mathcal{G}}\right)^{\gamma-1} \tag{8}$$

Finally, since $\gamma > 1$, applying $\frac{\mathbf{\Phi}}{\mathcal{G}} > 1$ gives us the result.

□

We discuss further intuition for these conditions below when they are used in the formal results. The relationship between the conditions and their implications for consumption behaviour will also be be discussed in detail in Section 5.

2.3 Perfect Foresight Benchmarks

{subsec:PFBdiscuss

To understand the economic implications of the patience conditions, we begin with the perfect foresight case.

When $q = 0$ and $\underline{\theta} = \overline{\theta} = \bar{\theta} = \underline{\psi} = \overline{\psi} = 1$, the process we have described above becomes a perfect foresight process. Below, when we say we assume perfect foresight, what we mean mathematically is:

{ass:pfincome}

Assumption I.2. *(Perfect Foresight Income Process).* $q = 0$ *and* $\underline{\theta} = \overline{\theta} = \bar{\theta} = \underline{\psi} = \overline{\psi} = 1$

Under perfect foresight, finite value of autarky reduces to a 'perfect foresight finite value of autarky' condition:

$$\beta\mathcal{G}^{1-\gamma} < 1. \tag{9}$$ {eq:PFFVAC}

2.3.1 Perfect Foresight without Liquidity Constraints

{subsubsec:PFUnco

Consider the familiar analytical solution to the perfect foresight model without liquidity constraints. In this case, the consumption Euler Equation always holds; with $\mathrm{u}'(\boldsymbol{c}) = \boldsymbol{c}^{-\gamma}$ and $\mathrm{u}'(\boldsymbol{c}_t) = \mathsf{R}\beta\mathrm{u}'(\boldsymbol{c}_{t+1})$ we have:

$$\boldsymbol{c}_{t+1}/\boldsymbol{c}_t = (\mathsf{R}\beta)^{1/\gamma}. \tag{10}$$ {eq:cGroFac}

Defining $\tilde{\mathcal{R}} :\!= \mathsf{R}/\mathcal{G}$, 'human wealth' is the present discounted value of income:

$$\begin{aligned} \boldsymbol{h}_t &= \boldsymbol{p}_t + \tilde{\mathcal{R}}^{-1}\boldsymbol{p}_t + \tilde{\mathcal{R}}^{-2}\boldsymbol{p}_t + \cdots + \tilde{\mathcal{R}}^{t-T}\boldsymbol{p}_t \\ &= \underbrace{\left(\frac{1-\tilde{\mathcal{R}}^{-(T-t+1)}}{1-\tilde{\mathcal{R}}^{-1}}\right)}_{\equiv h_t}\boldsymbol{p}_t. \end{aligned} \tag{11}$$ {eq:HDef}

For human wealth to have finite value, we must have:

{ass:FHWC}

Assumption I.3. *(Finite Human Wealth).*

$$\tilde{\mathcal{R}}^{-1} = \mathcal{G}/\mathsf{R} < 1. \tag{12}$$ {eq:FHWC2}

If $\tilde{\mathcal{R}}^{-1}$ is less than one, human wealth will be finite in the limit as $T \uparrow \infty$ because (noncapital) income growth is smaller than the interest rate at which that income is being discounted.

Under these conditions we can define a normalized finite-horizon perfect foresight consumption function (see Appendix C.1 for details) as follows:

$$\bar{\mathrm{c}}_{T-n}(m_{T-n}) = (\overbrace{m_{T-n}-1}^{\equiv b_{T-n}} + h_{T-n})\underline{\kappa}_{t-n}$$

where $\underline{\kappa}_t$ is the marginal propensity to consume (MPC) and satisfies:

$$\underline{\kappa}_{T-n}^{-1} = 1 + \left(\frac{\mathbf{\mathfrak{P}}}{\mathsf{R}}\right)\underline{\kappa}_{T-n+1}^{-1}. \tag{13}$$ {eq:MPCminInv}

Thus, for $\underline{\kappa}$ to be strictly positive as n goes to infinity, we must impose return impatience. The limiting consumption function then becomes:

$$\bar{\mathrm{c}}(m) = (m + h - 1)\underline{\kappa}, \tag{14}$$ {eq:cFuncPFUnc}

where (under return impatience)

$$\underline{\kappa} :\!= 1 - \frac{\mathbf{\mathfrak{P}}}{\mathsf{R}}. \tag{15}$$

In order to rule out the degenerate limiting solution in which $\bar{\mathrm{c}}(m) = \infty$, we also need to require (in the limit as the horizon extends to infinity) that human wealth remain

bounded (that is, we require a limiting 'finite human wealth'). Thus, while return impatience prevents a consumer from saving everything in the limit, 'finite value of human wealth' prevents infinite borrowing (against infinite human wealth) in the limit.

{prop:pfUCFHWC

Proposition 1. *Consider the normalized problem without liquidity constraints and with perfect foresight income (Assumption I.2). A nondegenerate limiting solution exists if and only if finite value of human wealth ($\tilde{\mathcal{R}}^{-1} < 1$) and return impatience (Assumption L.3) hold.*

The proof of the following claim is follows from straightforward algebra (see Appendix A.2).

{claim:PFConspC

Claim 2. *Assume finite limiting human wealth ($\tilde{\mathcal{R}}^{-1} < 1$). If growth impatience (Assumption S.1) holds, then finite value of autarky (Assumption L.1) holds. Moreover, if finite value of autarky (Assumption L.1) holds, then return impatience (Assumption L.3) holds.*

The claim implies that if we impose finite value of human wealth, then growth impatience is sufficient for nondegeneracy since finite value of autarky and return impatience follow. However, there are circumstances under which return impatience and finite value of human wealth can hold while the finite value of autarky fails. For example, if $\mathcal{G} = 0$, the problem is a standard 'cake-eating' problem with a nondegenerate solution under return impatience.

2.3.2 Perfect Foresight with Liquidity Constraints

Our ultimate interest is in the unconstrained problem with uncertainty. Here, we show that the perfect foresight constrained solution defines a useful limit for the unconstrained problem with uncertainty.

If a liquidity constraint requiring $b \geq 0$ is ever to be relevant, it must be relevant at the lowest possible level of market resources, $m_t = 1$, defined by the lower bound, $b_t = 0$ (if the constraint were relevant at any higher m, it would certainly be relevant here, because $\mathrm{u}' > 0$). The constraint is 'relevant' if it prevents the choice that would otherwise be optimal; at $m_t = 1$ it is relevant if the marginal utility from spending all of today's resources $c_t = m_t = 1$, exceeds the marginal utility from doing the same thing next period, $c_{t+1} = 1$; that is, if such choices would violate the Euler equation, Equation (4):

$$1^{-\gamma} > \mathsf{R}\beta\mathcal{G}^{-\gamma}1^{-\gamma}, \tag{16}$$

{eq:LiqConstrBinds

which is just a restatement of growth impatience. So, the constraint is relevant if and only if growth impatience holds.

{prop:PFCExist}

Proposition 2. *Consider the normalized problem with perfect foresight income (Assumption I.2) and assume $c_t \leq m_t$ for each t. If return impatience (Assumption L.3) holds, then a nondegenerate solution exists. Moreover, if return impatience does not*

hold, then a nondegenerate solution exists if and only if growth impatience (Assumption S.1) also holds.

Importantly, if return impatience fails ($\mathsf{R} < \mathbf{\Phi}$) and growth impatience holds ($\mathbf{\Phi} < \mathcal{G}$), then finite value of human wealth also fails ($\mathsf{R} < \mathcal{G}$). Despite the unboundedness of human wealth as the horizon extends arbitrarily, for any finite horizon the relevant liquidity constraint prevents borrowing. Similarly, when uncertainty is present, the natural borrowing constraint plays an analogous role in permitting a finite limiting solution despite unbounded limiting human wealth – we discuss the various parametric cases in Section 5.

2.4 Main Results for Problem with Uncertainty

{subsec:limSolExist

We are now ready to return to our primary interest, the model with permanent and transitory income shocks. Throughout this section, we assume the Friedman-Muth income process (Assumption I.1 holds).

2.4.1 Limiting MPCs

{subsubsec:cFuncBo

We first establish results regarding the shape of the consumption function.

{prop:cfuncprop}

Proposition 3. *For each t, c_t is increasing, twice continuously differentiable, strictly concave and $\mathrm{c}_t(0) = 0$.*

For a proof, see Appendix A.3.[22]

Next, we verify that the ratio of optimal consumption to market resources (c/m) is bounded by the minimal and maximal marginal propensities to consume (MPC). Recall that the MPCs answer the question 'if the consumer had an extra unit of resources, how much more spending would occur?' The minimal and maximal MPCs are the limits of the MPC as $m \to \infty$ and $m \to 0$, which we denote by $\underline{\kappa}_t$ and $\overline{\kappa}_t$ respectively. Since the consumer spends everything in the terminal period, $\underline{\kappa}_T = 1$ and $\overline{\kappa}_T = 1$. Furthermore, Proposition 3 will imply:[23]

$$\underline{\kappa}_t m_t \;\leq \mathrm{c}_t(m_t) \leq\; \overline{\kappa}_t m_t. \tag{17}$$

{eq:cBounds}

We define:

$$\underline{\kappa}\colon = \max\{0, 1 - \frac{\mathbf{\Phi}}{\mathsf{R}}\}, \qquad \overline{\kappa}\colon = 1 - q^{1/\gamma}\frac{\mathbf{\Phi}}{\mathsf{R}}, \tag{18}$$

{eq:MPCminmaxde

as the 'limiting minimal and maximal MPCs.' The following result verifies that the consumption share is bounded each period by the minimal and maximal MPCs, that the consumption function is asymptotically linear and that the MPCs converge to the limiting MPCs as the planning horizon recedes.[24]

[22] Carroll and Kimball [1996] proved concavity but not continuous differentiability.

[23] Note c'_t is positive, bounded above by 1 and decreasing, then apply L'Hôpital's Rule.

[24] Benhabib et al. [2015] show that the consumption function becomes linear as wealth approaches infinity in a model with capital income risk and liquidity constraints; Ma and Toda [2020] show that these results generalize to the limits derived here if capital income is added to the model.

Lemma 1. *(Limiting MPCs). If weak return impatience (Assumption L.4) holds, then:* {lemm:MPC}

(i) For each n:

$$\underline{\kappa}_{T-n}^{-1} = 1 + \left(\frac{\boldsymbol{\rho}}{\mathsf{R}}\right) \underline{\kappa}_{T-n+1}^{-1}. \tag{19}$$ {eq:MPCminInv}

Moreover, if return impatience (Assumption L.3) holds, then $\lim_{n\to\infty} \underline{\kappa}_{T-n} = \underline{\kappa} = 1 - \frac{\boldsymbol{\rho}}{\mathsf{R}}$ where $1 > \underline{\kappa} > 0$.

(ii) For each n:

$$\overline{\kappa}_{T-n}^{-1} = 1 + \left(q^{1/\gamma}\frac{\boldsymbol{\rho}}{\mathsf{R}}\right) \overline{\kappa}_{T-n+1}^{-1}. \tag{20}$$ {eq:MPCmaxInv}

Moreover, $\lim_{n\to\infty} \overline{\kappa}_{T-n} = \overline{\kappa}$, where $1 \geq \overline{\kappa} > 0$.

For the proof, see Appendix A.3.

The MPC bound as market resources approach infinity is easy to understand. Recall that $\bar{\mathrm{c}}$ from the perfect foresight case will be an upper bound in the problem with uncertainty; analogously, $\underline{\kappa}$ becomes the MPC's lower bound. As the *proportion* of consumption that will be financed out of human wealth approaches zero, the proportional difference between the solution to the model with uncertainty and the perfect foresight model shrinks to zero.

To understand the maximal limiting MPC, the essence of the argument is that as market resources approach zero, the overriding consideration that limits consumption is the (recursive) fear of the zero-income events — this is why the probability of the zero income event q appears in the expression for the maximal MPC. Weak return impatience is too weak to guarantee a lower bound on the share of consumption to market resources; it merely prevents the upper bound on the share of consumption to market resources from approaching zero. Weak return impatience thereby prevents a situation where *everyone* consumes an arbitrarily small share of current market resources as the planning horizon recedes. This insight plays a key role in the proof for the existence of a non-degenerate solution in what follows.

2.4.2 Existence of Limiting Nondegenerate Solution

{subsubsec:eventual

To address the challenges of unbounded state-spaces, Boyd [1990] provided a weighted contraction mapping theorem. We use this approach to first show that while the stationary operator $\mathbb{T}$ may be undefined, each period's Bellman operator will be a contraction. We then show the value function generated by the Bellman iteration given by Problem $(\mathcal{P}_N)$ generates a Cauchy sequence in a complete metric space; that is, the sequence of value functions converges to a nondegenerate solution in $\mathcal{C}(\mathbb{R}_{++}, \mathbb{R})$.

Let $\mathcal{C}(\mathbb{R}_{++}, \mathbb{R})$ be the space of continuous functions from $\mathbb{R}_{++}$ to $\mathbb{R}$.

Definition 2. *Fix* f *such that* $\mathrm{f} \in \mathcal{C}(\mathbb{R}_{++}, \mathbb{R})$ *and let* φ *be a function such that* $\varphi \in \mathcal{C}(\mathbb{R}_{++}, \mathbb{R})$ *and* $\varphi > 0$*. The function* f *will be* φ*-bounded if the* φ*-norm of* f*:*

$$\|\mathrm{f}\|_{\varphi} = \sup_{s \in \mathbb{R}_{++}} \left[\frac{|\mathrm{f}(s)|}{\varphi(s)} \right], \tag{21}$$ {eq:phinorm}

is finite. We will call $\mathcal{C}_{\varphi}(\mathbb{R}_{++}, \mathbb{R})$ *the subspace of functions in* $\mathcal{C}(\mathbb{R}_{++}, \mathbb{R})$ *that are* φ*-bounded.*

We define the weighting function as:

$$\varphi(x) = \zeta + x^{1-\gamma} \tag{22}$$

where $\zeta \in \mathbb{R}_{++}$ is a constant.

Next, for any lower bound $\underline{\nu}$ and upper-bound $\overline{\nu}$ on the share of consumption to market resources, define the Bellman operator $\mathbb{T}^{\underline{\nu},\overline{\nu}}$, with $\mathbb{T}^{\underline{\nu},\overline{\nu}} : \mathcal{C}_{\varphi}(S, Y) \to \mathcal{C}_{\varphi}(S, Y)$, as:

$$\mathbb{T}^{\underline{\nu},\overline{\nu}}\mathrm{f}(m) = \max_{c \in [\underline{\nu} m, \overline{\nu} m]} \left\{ \mathrm{u}(c) + \beta \mathbb{E} \tilde{\mathcal{G}}^{1-\gamma} \mathrm{f}(\tilde{\mathcal{R}}(m-c) + \boldsymbol{\xi}) \right\}, \; m \in \mathbb{R}_{++}, \mathrm{f} \in \mathcal{C}_{\varphi}(\mathbb{R}_{++}, Y). \tag{23}$$

Even if in the limit the minimal marginal propensity to consume approaches zero, the finite-horizon value functions defined by the normalized recursive problem, Problem $(\mathcal{P}_N)$, will satisfy $\mathrm{v}_t = \mathbb{T}^{\underline{\kappa}_t,\overline{\kappa}_t}\mathrm{v}_{t+1}$ since consumption shares are bounded by the minimal and maximal MPCs (Lemma 3 and Equation (17)). We now show this implies that the operator $\mathbb{T}^{\underline{\nu},\overline{\nu}}$ is a contraction on $\mathcal{C}_{\varphi}(\mathbb{R}_{++}, Y)$ for a suitably narrow interval $[\underline{\nu}, \overline{\nu}]$.

Theorem 1. *(Contraction Mapping Under Consumption Bounds). If weak return impatience (Assumption L.4) and finite value of autarky (Assumption L.1) hold, then there exists* k *large enough and* $\alpha \in (0,1)$ *such that for all* $\overline{\nu}$ *with* $\overline{\nu} \leq \overline{\kappa}_{T-k}$ *and* $\underline{\nu}$ *with* $\overline{\nu} > \underline{\nu} > 0$*, the Bellman operator* $\mathbb{T}^{\underline{\nu},\overline{\nu}}$ *is a contraction with modulus* α*.* {thm:cmap}

An implication of the theorem is that eventually, the maximal MPCs will be small enough such that each of the Bellman operators generating the sequence of value functions as the terminal time T recedes (that is, as the horizon expands) will be contraction maps. We can now relate the sequence of contraction maps to a limiting solution defined in Section 2.1.1.

Theorem 2. *(Existence of Nondegenerate Solution). If weak return impatience (Assumption L.4) and finite value of autarky (Assumption L.1) hold, then:* {thm:convgtobellm}

(i) There exists k *such that for all* t *with* $t < k$ *and* $\underline{\nu}$*, with* $\overline{\kappa}_t > \underline{\nu} > 0$*,* $\mathbb{T}^{\underline{\nu},\overline{\kappa}_t}$ *is a contraction with modulus* α*, where* $\alpha < 1$ *and the sequence* $\{\mathrm{v}_{T-n}\}_{n=0}^{\infty}$ *converges point-wise to* v*, with* $\mathrm{v} \in \mathcal{C}_{\varphi}(\mathbb{R}_{++}, \mathbb{R})$*.*

(ii) The function v *is a fixed point of* $\mathbb{T}$ *and there exists a measurable limiting policy function,* c*, such that* $\mathrm{c}\colon \mathbb{R}_+ \to \mathbb{R}$ *and:*

$$\mathbb{T}\mathrm{v}(m) = \mathrm{u}(\mathrm{c}(m)) + \beta \mathbb{E}\tilde{\mathcal{G}}^{1-\gamma}\mathrm{v}(\tilde{\mathcal{R}}(m - \mathrm{c}(m)) + \boldsymbol{\xi}), \qquad m \in \mathbb{R}_{++}. \tag{24}$$ {eq:stationarybelln

(iii) The sequence $\{\mathrm{c}_{T-n}\}_{n=0}^{\infty}$ *converges point-wise to* c *and* c *is a limiting nondegenerate solution.*

Proof. The first claim of Item (i) follows from Theorem 1, since $\underline{\nu} > 0$ and for each t, $\overline{\kappa}_t < \overline{\kappa}_k$ by Lemma 3. We now prove $\{\mathrm{v}_{T-n}\}_{n=0}^{\infty}$ converges point-wise to a limiting nondegenerate solution v. In the proof, to streamline the notation, we define an index t_n for the sequence t that depends on n. Specifically, $t_n = T - n$. Now, for all $n > k + 2$, $\mathrm{v}_{t_n} = \mathbb{T}^{\underline{\kappa}_{t_n}, \overline{\kappa}_{t_n}}\mathrm{v}_{t_{n-1}}$ holds by definition of the value function given by the Bellman Equation $(\mathcal{P}_N)$. Moreover, since $\overline{\kappa}_{t_{n-1}} \geq \overline{\kappa}_{t_n}$ by Lemma 3, we have:

$$\mathrm{v}_{t_n} = \mathbb{T}^{\underline{\kappa}_{t_n}, \overline{\kappa}_{t_{n-1}}}\mathrm{v}_{t_{n-1}}$$

and since $\underline{\kappa}_{t_n} \leq \underline{\kappa}_{t_{n-1}}$, we have:

$$\begin{aligned}\mathrm{v}_{t_{n-1}} &= \mathbb{T}^{\underline{\kappa}_{t_{n-1}}, \overline{\kappa}_{t_{n-1}}}\mathrm{v}_{n-2} \\ &= \mathbb{T}^{\underline{\kappa}_{t_n}, \overline{\kappa}_{t_{n-1}}}\mathrm{v}_{t_{n-2}}\end{aligned}$$

Next, take the φ-norm distance between v_{t_n} and $\mathrm{v}_{t_{n-1}}$, and note $\mathbb{T}^{\underline{\nu}, \overline{\kappa}_t}$ is a contraction. We have:

$$\|\mathrm{v}_{t_n} - \mathrm{v}_{t_{n-1}}\|_\varphi = \|\mathbb{T}^{\underline{\kappa}_{t_n}, \overline{\kappa}_{t_{n-1}}}\mathrm{v}_{t_{n-1}} - \mathbb{T}^{\underline{\kappa}_{t_n}, \overline{\kappa}_{t_{n-1}}}\mathrm{v}_{t_{n-2}}\|_\varphi \leq \alpha\|\mathrm{v}_{t_{n-1}} - \mathrm{v}_{t_{n-2}}\|_\varphi.$$

As such, $\|\mathrm{v}_{t_n} - \mathrm{v}_{t_{n-1}}\|_\varphi \leq \alpha\|\mathrm{v}_{t_{n-1}} - \mathrm{v}_{t_{n-2}}\|_\varphi$; because n is arbitrary and α holds for all n by Theorem 1, this is a sufficient condition for $\{\mathrm{v}_{T-n}\}_{n=k+2}^{\infty}$ to be Cauchy. Since $\mathcal{C}_\varphi(\mathbb{R}_{++}, Y)$ is a complete metric space, and $\mathrm{v}_{t_{n-2}} \in \mathcal{C}_\varphi(\mathbb{R}_{++}, Y)$ for each n, v_{t_n} converges to v, with $v \in \mathcal{C}_\varphi(\mathbb{R}_{++}, Y)$. (Proof continued in Appendix A.4.) □

Without the stronger assumption holding, $\underline{\kappa} = 0$ and $\mathbb{T}^{0,\kappa_k}$ will not be a well-defined operator from $\mathcal{C}_\varphi(\mathbb{R}_{++}, Y)$ to $\mathcal{C}_\varphi(\mathbb{R}_{++}, Y)$, even for k large enough (recall our discussion below Equation (5)). Nonetheless, the sequence of value functions produced by the composition of the per-period Bellman operators $\mathbb{T}^{\underline{\kappa}_t, \overline{\kappa}_t}$ will be a Cauchy sequence converging to the limiting solution. Due to weak return impatience, the upper bound on consumption converges to a strictly positive share of market resources, preventing consumption to limit to zero. The remainder of this proof for Item (ii) and Item (iii) in the appendix shows that the limiting value functions is a fixed point to the operator $\mathbb{T}$ and that the sequence of consumption functions converge.

Finite value of autarky is the second key assumption required to show existence of limiting solutions and guarantees the value is finite (in levels) for a consumer who spent exactly their permanent income every period (see Section 5.2). In the normalized problem $(\mathcal{P}_N)$, finite value of autarky ensures the expected value of the random discount

factor is less than one.[25] The intuition for the finite value of autarky condition is that, with an infinite-horizon, with any strictly positive initial amount of bank balances b_0, in the limit your value can always be made greater than you would get by consuming exactly the sustainable amount (say, by consuming $(\mathsf{r}/\mathsf{R})b_0 - \epsilon$ for some arbitrarily small $\epsilon > 0$).

Finally, we verify that the converged nondegenerate consumption functions satisfies the same consumption bounds as the per-period consumption functions.

{lemm:MPC}

Claim 3. *If weak return impatience (Assumption L.4) holds, then* $\underline{\kappa}m \leq \mathrm{c}(m) \leq \overline{\kappa}m$, *ii)* $\lim_{m\to\infty} \mathrm{c}(m)/m = \underline{\kappa}m$ *and iii)* $\lim_{m\to 0} \mathrm{c}(m)/m = \overline{\kappa}m$.

2.4.3 The Liquidity Constrained Solution as a Limit

{subsubsec:deatonIs

Recall the common assumption [Deaton, 1991, Aiyagari, 1994, Li and Stachurski, 2014, Ma et al., 2020] of a strictly positive minimum value of income and a non-trivial artificial liquidity constraint, namely $a_t \geq 0$. We will refer to the set-up from Section 2.1, with Assumption 2 modified so $q = 0$ as the "liquidity constrained problem." We now show a finite-horizon solution to the liquidity constrained problem is the limit of the problems as the probability q of the zero-income event approaches zero. Let $\mathrm{c}_t(\bullet; q)$ be the consumption function for a problem where Assumption I.1 holds for a given fixed q, with $q > 0$. Moreover, let $\grave{\mathrm{c}}_t$ be the limiting consumption function for the liquidity constrained problem (note that the liquidity constraint $c_t \leq m_t$, or $a_t \geq 0$, becomes relevant only when $q = 0$).

{prop:deatonIsLim

Proposition 4. *Assume the setting of Theorem 2. We have* $\lim_{q\to 0} \mathrm{c}_t(m; q) = \grave{\mathrm{c}}_t(m)$ *for each* t *and* $m \in \mathbb{R}_{++}$.

Intuitively, if we impose the artificial constraint without changing q and maintain $q > 0$, it would not affect behavior. This is because the possibility of earning zero income over the remaining horizon already prevents the consumer from ending the period with zero assets. For precautionary reasons, the consumer will save something. However, the *extent* to which the consumer feels the need to make this precautionary provision depends on the *probability* that it will turn out to matter. As $q \to 0$, the precautionary saving induced by the zero-income events approaches zero, and "zero" is the amount of precautionary saving that would be induced by a zero-probability event by the impatient liquidity constrained consumer. See Appendix A.6 for the formal proof.

3 Individual Buffer Stock Stability

{sec:individStability

In this section we analyse two notions of stability which will be useful for studying either an individual or a population of individuals who behave according to the converged

[25] Assumption 2.1 by Ma et al. [2020] specializes to finite value of autarky in our case. Assumption 2.2 by Ma et al. [2020] requires specializes $\beta\mathsf{R} < 1$. Our proofs do not require $\beta\mathsf{R} < 1$.

consumption rule. Consider a individual who at time t holds normalized and non-normalized market resources m_t and $\boldsymbol{m}_t$ and follows the converged decision function c. The time-t consumption for the consumer will be $c_t = \mathrm{c}(m_t)$ and the time $t+1$ market resources will be a random variable $m_{t+1} = \tilde{\mathcal{R}}_{t+1}(m_t - \mathrm{c}(m)) + \boldsymbol{\xi}_{t+1}$. At the individual level, we are interested in whether the current level of market resources is above or below a 'target' level such that the magnitude of the precautionary motive (which causes a consumer to save) exactly balances the impatience motive (which makes them want to dissave). At the individual 'target,' the expected market resources ratio in the next period, *conditioned on current market resources*, will be the same as the ratio in the current period. The intensifying strength of the precautionary motive with decreasing market resources can ensure stability of the target. Below the target, the urgency to save due to the precautionary motive leads to an expected rise in market resources. Conversely, above the target, impatience prevails, leading to an expected reduction of market resources. In this way, the 'target' essentially defines the desired 'buffer stock' of resources for the consumer.

To help motivate the theoretical results concerning existence of a target level of market resources, Figure 1 shows the expected growth factors for consumption, the level of market resources, and the market resources to permanent income ratio, $\mathbb{E}_t[\boldsymbol{c}_{t+1}/\boldsymbol{c}_t]$, $\mathbb{E}_t[\boldsymbol{m}_{t+1}/\boldsymbol{m}_t]$, and $\mathbb{E}_t[m_{t+1}/m_t]$. The figure is generated using parameters discussed in Section 5, Table 2. First, the figure shows how as $m_t \to \infty$ the expected consumption growth factor goes to Φ, indicated by the lower bound in Figure 1. Moreover, as m_t approaches zero the consumption growth factor approaches ∞. The following proposition establishes the asymptotic growth factors formally (See Appendix B for a proof).

{prop:convgGrowth

Proposition 5. *We have* $\lim_{m_t \to \infty} \mathbb{E}_t[\boldsymbol{c}_{t+1}/\boldsymbol{c}_t] = \boldsymbol{\Phi}$ *and* $\lim_{m_t \to 0} \mathbb{E}_t[\boldsymbol{c}_{t+1}/\boldsymbol{c}_t] = \infty$.

Next, consider the implications of Figure 1 for individual stability. The figure shows a value of the market resources ratio, $m_t = \check{m}$, at which point the expected growth factor of the level of market resources $\boldsymbol{m}$ matches the expected growth factor of permanent income $\mathcal{G}$. A distinct and larger target ratio, $\hat{m}$, also exists. At this ratio, $\mathbb{E}[m_{t+1}/m_t] = 1$, and the expected growth factor of consumption is less than $\mathcal{G}$. Importantly, at the individual level, this model does not have a single m at which $\boldsymbol{p}$, $\boldsymbol{m}$ and $\boldsymbol{c}$ are all expected to grow at the same rate. Yet, when we aggregate across individuals, balanced growth paths can exist, even if there does not exist a target ratio where $\mathbb{E}[m_{t+1}/m_t] = 1$. Before we discuss aggregates further, we'll first set the conditions required for the existence of individual targets.

3.1 Limits as m Approaches Zero

{subsec:LimitsAsmt

Equation (20) shows that the limiting value of $\overline{\kappa}$ is

$$\overline{\kappa} = 1 - \mathsf{R}^{-1}(q\mathsf{R}\beta)^{1/\gamma}.$$

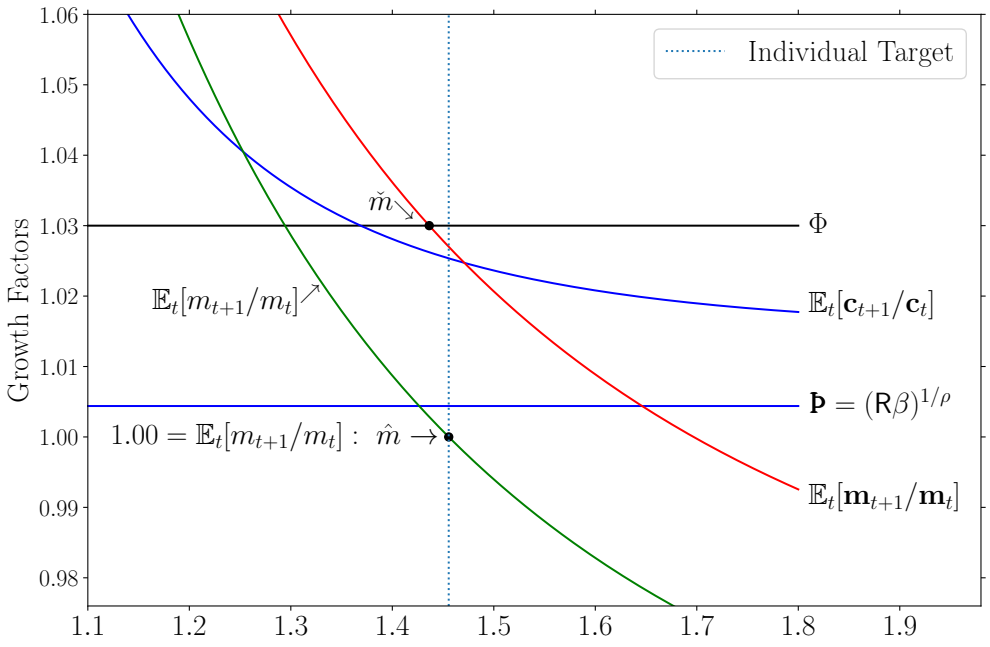

Figure 1 'Stable' (Target; Balanced Growth) m Values {fig:cNrmTargetFig

Defining $\mathrm{e}(m) = \mathrm{c}(m)/m$ we have

$$\lim_{m \downarrow 0} \mathrm{e}(m) = (1 - q^{1/\gamma}\frac{\mathbf{\dot{p}}}{\mathsf{R}}) = \overline{\kappa}.$$

Now using the continuous differentiability of the consumption function along with L'Hôpital's rule, we have

$$\lim_{m \downarrow 0} \mathrm{c}'(m) = \lim_{m \downarrow 0} \mathrm{e}(m) = \overline{\kappa}.$$

Figure 4 visually confirms that the numerical solution obtains this limit for the MPC as m approaches zero.

For consumption growth, as $m \downarrow 0$ we have

$$\begin{aligned}
\lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\mathrm{c}(m_{t+1})}{\mathrm{c}(m_t)} \right) \mathcal{G}_{t+1} \right] &> \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t) + \Xi_{t+1})}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&= q \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&\quad + (1-q) \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t) + \Theta_{t+1}/(1-q))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&> (1-q) \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\Theta_{t+1}/(1-q))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&= \infty
\end{aligned}$$

where the second-to-last line follows because $\lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right]$ is positive, and the last line follows because the minimum possible realization of Θ_{t+1} is $\underline{\Theta} > 0$ so the minimum possible value of expected next-period consumption is positive.[26]

3.2 Unique 'Stable' Points

{subsec:onetarget}

One kind of 'stable' point is a 'target' value $\hat{m}$ such that if $m_t = \hat{m}$, then $\mathbb{E}_t[m_{t+1}] = m_t$. Existence of such a target turns out to require the strong growth impatience condition.

Theorem 3. *(Individual Market-Resources-to-Permanent-Income Ratio Target). Consider the problem defined in Section 2.1. If weak return impatience (Assumption L.4), finite value of autarky (Assumption L.1) and strong growth impatience (Assumption S.2) hold, then there exists a unique market resources to permanent income ratio, $\hat{m}$, with $\hat{m} > 0$, such that:*

$$\mathbb{E}_t[m_{t+1}/m_t] = 1 \text{ if } m_t = \hat{m}. \tag{25}$$

Moreover, $\hat{m}$ is a point of 'stability' in the sense that:

$$\begin{aligned}
\forall m_t \in (0, \hat{m}),\ \mathbb{E}_t[m_{t+1}] > m_t \\
\forall m_t \in (\hat{m}, \infty),\ \mathbb{E}_t[m_{t+1}] < m_t.
\end{aligned} \tag{26}$$

Since $m_{t+1} = \tilde{\mathcal{R}}_{t+1}(m_t - \mathrm{c}(m_t))\tilde{\mathcal{R}}_{t+1} + \boldsymbol{\xi}_{t+1}$, the implicit equation for $\hat{m}$ becomes:

[26] None of the arguments in either of the two prior sections depended on the assumption that the consumption functions had converged. With more cumbersome notation, each derivation could have been replaced by the corresponding finite-horizon versions. This strongly suggests that it should be possible to extend the circumstances under which the problem can be shown to define a contraction mapping to the union of the parameter values under which {RIC,FHWC} hold and {FVAC,WRIC} hold. That extension is not necessary for our purposes here, so we leave it for future work.

$$\begin{aligned}\mathbb{E}_t[(\hat{m}-\mathrm{c}(\hat{m}))\tilde{\mathcal{R}}_{t+1}+\boldsymbol{\xi}_{t+1}] &= \hat{m} \\ (\hat{m}-\mathrm{c}(\hat{m}))\underbrace{\tilde{\mathcal{R}}\mathbb{E}_t[\psi^{-1}]}_{\equiv\bar{\tilde{\mathcal{R}}}}+1 &= \hat{m}. \end{aligned} \tag{27}$$ {eq:mTargImplicit}

The market-resources-to-permanent-income ratio target is the most restrictive among several competing definitions of stability. Our least restrictive definition of 'stability' derives from a traditional aggregate question in macro models: whether or not there is a 'balanced growth' equilibrium in which aggregate variables (income, consumption, market resources) all grow by the same factor $\mathcal{G}$. In particular, if growth impatience holds, the problem will exhibit a balanced-growth 'pseudo-steady-state' point, by which we mean that there is some $\check{m}$ such that if: $m_t > \check{m}$, then $\mathbb{E}_t[\boldsymbol{m}_{t+1}/\boldsymbol{m}_t] < \mathcal{G}$. Conversely if $m_t < \check{m}$ then $\mathbb{E}_t[\boldsymbol{m}_{t+1}/\boldsymbol{m}_t] > \mathcal{G}$. The target $\check{m}$ will be such that $\boldsymbol{m}$ growth matches $\mathcal{G}$, allowing us to write the implicit equation for $\check{m}$ as follows:

$$\begin{aligned}\mathbb{E}_t[\boldsymbol{m}_{t+1}]/\boldsymbol{m}_t &= \mathbb{E}_t[\boldsymbol{p}_{t+1}]/\boldsymbol{p}_t \\ \mathbb{E}_t[m_{t+1}\mathcal{G}\psi_{t+1}\boldsymbol{p}_t]/(m_t\boldsymbol{p}_t) &= \mathbb{E}_t[\boldsymbol{p}_t\mathcal{G}\psi_{t+1}]/\boldsymbol{p}_t \\ \mathbb{E}_t\left[\psi_{t+1}\underbrace{\left((m_t-\mathrm{c}(m_t)\mathsf{R}/(\mathcal{G}\psi_{t+1}))+\boldsymbol{\xi}_{t+1}\right)}_{m_{t+1}}\right]/m_t &= 1 \\ \mathbb{E}_t\left[(\check{m}-\mathrm{c}(\check{m}))\overbrace{\mathsf{R}/\mathcal{G}}^{\tilde{\mathcal{R}}}+\psi_{t+1}\boldsymbol{\xi}_{t+1}\right] &= \check{m} \\ (\check{m}-\mathrm{c}(\check{m}))\tilde{\mathcal{R}}+1 &= \check{m}.\end{aligned} \tag{28}$$ {eq:balgrostable}

The only difference between (28) and (27) is the substitution of $\tilde{\mathcal{R}}$ for $\bar{\tilde{\mathcal{R}}}$.[27,28]

Theorem 4. *(Individual Balanced-Growth 'Pseudo Steady State'). Consider the prob-*

[27] A third 'stable point' is the $\tilde{m}$ where $\mathbb{E}_t[\log \boldsymbol{m}_{t+1}] = \log \mathcal{G}\boldsymbol{m}_t$; this can be conveniently rewritten as $\mathbb{E}_t\left[\log\left((\tilde{m}-\mathrm{c}(\tilde{m}))\tilde{\mathcal{R}}+\psi_{t+1}\boldsymbol{\xi}_{t+1}\right)\right] = \log \tilde{m}_t$. Because the expectation of the log of a stochastic variable is less than the log of the expectation, if a solution for $\tilde{m}$ exists it will satisfy $\tilde{m} > \check{m}$; in turn, if $\hat{m}$ exists, $\hat{m} > \tilde{m}$. The target $\tilde{m}$ is guaranteed to exist when the log growth impatience condition is satisfied (see below). For our purposes, little would be gained by an analysis of this point parallel to those of the other points of stability; but to accommodate potential practical uses, the Econ-ARK toolkit computes the value of this point (when it exists) as `mBalLog`.

[28] Our choice to call to this the individual problem's 'individual balanced-growth pseudo-steady-state' $\check{m}$ is motivated by what happens in the case where all draws of all future shocks just happen to take on their expected value of 1.0. (They unexpectedly always take on their expected values). In that infinitely improbable case, the economy *would* exhibit balanced growth:

$$\mathbb{E}_t[m_{t+1}/m_t|\psi_{t+1}=\boldsymbol{\xi}_{t+1}=1] = \mathcal{G}\left(\check{m}-\mathrm{c}(\check{m})\tilde{\mathcal{R}}+1\right)/m = \mathcal{G}.$$

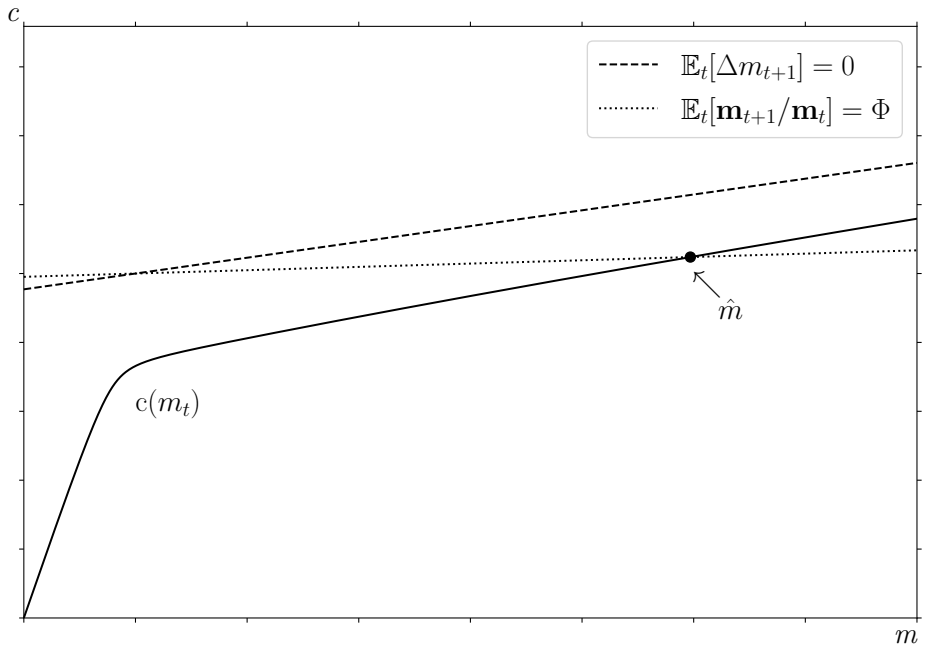

Figure 2 {FVAC,GIC,~~GIC-Mod~~}: No Target Exists But SS Does {fig:GICModFailsBu

lem defined in Section 2.1. If weak return impatience (Assumption L.4), finite value of autarky (Assumption L.1) and growth impatience (Assumption S.1) hold, then there exists a unique pseudo-steady-state market resources to permanent income ratio $\check{m} > 0$ *such that:*

$$\mathbb{E}_t[\psi_{t+1} m_{t+1}/m_t] = 1 \text{ if } m_t = \check{m}. \tag{29}$$

Moreover, $\check{m}$ *is a point of stability in the sense that:*

$$\begin{aligned} \forall m_t \in (0, \check{m}),\ \mathbb{E}_t[\boldsymbol{m}_{t+1}]/\boldsymbol{m}_t > \mathcal{G} \\ \forall m_t \in (\check{m}, \infty),\ \mathbb{E}_t[\boldsymbol{m}_{t+1}]/\boldsymbol{m}_t < \mathcal{G}. \end{aligned} \tag{30}$$

3.3 Example With Balanced-Growth $\check{m}$ But No Target $\hat{m}$

Because the equations defining target and pseudo-steady-state m, (27) and (28), differ only by substitution of $\tilde{\mathcal{R}}$ for $\tilde{\bar{\mathcal{R}}} = \tilde{\mathcal{R}}\mathbb{E}[\psi^{-1}]$, if there are no permanent shocks ($\psi \equiv 1$), the conditions are identical. For many parameterizations (e.g., under the baseline parameter values used for constructing figure 1), $\hat{m}$ and $\check{m}$ will not differ much.

An illuminating exception is exhibited in Figure 2, which modifies the baseline parameter values by quadrupling the variance of the permanent shocks, enough to cause failure of strong growth impatience; now there is no target level of market resources $\hat{m}$. Nonetheless, the pseudo-steady-state still exists because it turns off realizations of

the permanent shock. It is tempting to conclude that the reason target $\hat{m}$ does not exist is that the increase in the size of the shocks induces a precautionary motive that increases the consumer's effective patience. The interpretation is not correct because as market resources approach infinity, precautionary saving against noncapital income risk becomes negligible (as the proportion of consumption financed out of such income approaches zero). The correct explanation is more prosaic: The increase in uncertainty boosts the expected uncertainty-modified rate of return factor from $\tilde{\mathcal{R}}$ to $\bar{\tilde{\mathcal{R}}} > \tilde{\mathcal{R}}$ which reflects the fact that in the presence of uncertainty the expectation of the inverse of the growth factor increases: $\underline{\mathcal{G}} > \mathcal{G}$. That is, in the limit as $m \to \infty$ the increase in effective impatience reflected in $\frac{\mathbf{\Phi}}{\underline{\mathcal{G}}}\mathbb{E}[\psi^{-1}] < \frac{\mathbf{\Phi}}{\mathcal{G}}$ is entirely due to the certainty-equivalence growth adjustment, not to a (limiting) change in precaution. In fact, the next section will show that an aggregate balanced growth equilibrium will exist even when realizations of the permanent shock are not turned off: The required condition for aggregate balanced growth is the regular growth impatience, which ignores the magnitude of permanent shocks.[29]

Before we get to the formal arguments, the key insight can be understood by considering an economy that starts, at date t, with the entire population at $m_t = \check{m}$, but then evolves according to the model's assumed dynamics between t and $t+1$. Equation (28) will still hold, so for this first period, at least, the economy will exhibit balanced growth: the growth factor for aggregate $\boldsymbol{M}$ will match the growth factor for permanent income $\mathcal{G}$. It is true that there will be people for whom financial balances, b_{t+1}, where $b_{t+1} = k_{t+1}\mathsf{R}/(\mathcal{G}\psi_{t+1})$, are boosted by a small draw of ψ_{t+1}. However, their contribution to the *level* of the aggregate variable is given by $\boldsymbol{b}_{t+1} = b_{t+1}\boldsymbol{p}_t\psi_{t+1}$, so their b_{t+1} is reweighted by an amount that exactly unwinds that divisor-boosting. This means that it is possible for the consumption-to-permanent-income ratio for every consumer to be small enough that their market resources ratio is expected to rise, and yet for the economy as a whole to exhibit a balanced growth equilibrium with a finite aggregate balanced growth steady state $\check{M}$ (this is not numerically the same as the individual pseudo-steady-state ratio $\check{m}$ because the problem's nonlinearities have consequences when aggregated).[30]

4 Aggregate Invariant Relationships

In this section, we move from characterizing the individual decision rule to properties of a distribution of individuals following the converged nondegenerate consumption rule c. Assume a continuum of *ex ante* identical buffer-stock households, with constant total mass normalized to one and indexed by i. Szeidl [2013] proved that such a population,

[29] Szeidl [2013]'s impatience condition, discussed below, also tightens as uncertainty increases, but this is also not a consequence of a precaution-induced increase in patience – it represents an increase in the tightness of the requirements of the 'mixing condition' used in his proof.

[30] Still, the pseudo-steady-state can be calculated from the policy function without any simulation, and therefore serves as a low-cost starting point for the numerical simulation process; see Harmenberg-Aggregation for an example.

following the consumption rule c, will be characterized by invariant distributions of m, c, and a under the log growth impatience condition:[31]

$$\log \frac{\mathbf{\mathfrak{p}}}{\mathcal{G}} < \mathbb{E}[\log \psi] \tag{31}$$ {eq:GICSdl}

which is stronger than our growth impatience ($\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}} < 1$), but weaker than our strong growth impatience ($\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}\mathbb{E}[\psi^{-1}] < 1$).[32]

Harmenberg [2021] substitutes a clever change of probability-measure into Szeidl's proof, with the implication that under growth impatience, invariant permanent-income-weighted distributions of m and c exist. In particular, let $\mathcal{F}_{m_t,\boldsymbol{p}_t}$ be the joint CDF of normalized market resources and permanent income at time t.[33] The permanent-income-weighted CDF of m_t, $\tilde{\mathcal{F}}_{m_t}$, will be:

$$\tilde{\mathcal{F}}_{m_t}(x) = \mathcal{G}^{-t} \int_0^x \int_0^\infty \boldsymbol{p} \mathcal{F}_{m_t,\boldsymbol{p}_t}(dm, d\boldsymbol{p}) \tag{32}$$ {eq:HarmCDF}

Simply put, the permanent-income-weighted CDF shows how the total 'mass' of permanent income is distributed along normalized market resources.[34] The change of variables allows Harmenberg [2021] to prove a conjecture from an earlier draft of this paper (Carroll [2019, Submitted]) that under growth impatience, aggregate consumption grows at the same rate $\mathcal{G}$ as aggregate noncapital income in the long run (with the corollary that aggregate assets and market resources grow at that same rate). Harmenberg [2021] also shows how the reformulation can reduce costs of calculation by over a factor of 100.[35] The remainder of this section draws out the implications of these points for aggregate balanced growth factors.

[31] Szeidl [2013]'s equation (9), in our notation, is:

$$\begin{aligned} \mathbb{E} \log \mathsf{R}(1-\kappa) &< \mathbb{E} \log \mathcal{G}\psi \\ \mathbb{E} \log \mathsf{R}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} &< \mathbb{E} \log \mathcal{G}\psi \\ \log \frac{\mathbf{\mathfrak{p}}}{\mathcal{G}} &< \mathbb{E} \log \psi \end{aligned}$$

which, exponentiated, yields (31).

[32] Under our default (though not required) assumption that $\log \psi \sim \mathcal{N}(-\sigma_\psi^2/2, \sigma_\psi^2)$; strong growth impatience in this case, is $\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}} < \exp(-\sigma^2)$, so if strong growth impatience holds then Szeidl's condition will hold.

[33] In the notation in Harmenberg [2021], the permanent-income-weighted measures are denoted as $\tilde{\psi}^m$.

[34] The change of variables is analogous to weighting the mass of objects by coordinates and integrating to calculate the center of gravity. **?** also use a similar approach to compare the relative dependence on labour and capital income across the wealth distribution.

[35] The Harmenberg method is implemented in the Econ-ARK; see the last part of `test_Harmenbergs_method.sh`. Confirming the computational advantage of Harmenberg's method, this notebook finds that the Harmenberg method reduces the simulation size required for a given degree of accuracy by two orders of magnitude under the baseline parameter values defined above.

4.1 Aggregate Balanced Growth of Income, Consumption, and Wealth

{subsec:cGroEqPer

Define $\mathbb{M}$ to yield the expected value operator with respect to the empirical distribution of a variable across the population (as distinct from the operator $\mathbb{E}$ which represents beliefs about the future for a given individual).[36] Using boldface capitals for aggregates, the growth factor for aggregate noncapital income becomes:

$$\boldsymbol{Y}_{t+1}/\boldsymbol{Y}_{t} = \mathbb{M}\left[\boldsymbol{\xi}_{t+1}\mathcal{G}\psi_{t+1}\boldsymbol{p}_{t}\right]/\mathbb{M}\left[\boldsymbol{p}_{t}\boldsymbol{\xi}_{t}\right] = \mathcal{G}$$

because of the independence assumptions we have made about the shocks $\boldsymbol{\xi}$ and ψ.

Consider an economy that satisfies the Szeidl impatience condition (31) and has existed for long enough by date t that we can consider it as Szeidl-converged. In such an economy a microeconomist with a population-representative panel dataset could calculate the growth factor of consumption for each individual household, and take the average:

$$\begin{aligned}\mathbb{M}\left[\Delta \log \boldsymbol{c}_{t+1}\right] &= \mathbb{M}\left[\log c_{t+1}\boldsymbol{p}_{t+1} - \log c_{t}\boldsymbol{p}_{t}\right] \\ &= \mathbb{M}\left[\log \boldsymbol{p}_{t+1} - \log \boldsymbol{p}_{t}\right] + \mathbb{M}\left[\log c_{t+1} - \log c_{t}\right].\end{aligned} \tag{33}$$

{eq:MeanDeltaLogC

Because this economy is Szeidl-converged, distributions of c_t and c_{t+1} will be identical, so that the second term in (33) disappears; hence, mean cross-sectional growth factors of consumption and permanent income are the same:

$$\mathbb{M}\left[\Delta \log \boldsymbol{c}_{t+1}\right] = \mathbb{M}\left[\Delta \log \boldsymbol{p}_{t+1}\right] = \log \mathcal{G}. \tag{34}$$

{eq:MeanDeltaLogC

In a Harmenberg-invariant economy (and therefore also any Szeidl-invariant economy), a similar proposition holds in the cross-section as a direct implication of the fact that a constant proportion of total permanent income is accounted for by the successive sets of consumers with any particular m (recall Equation (32)). This fact is one way of interpreting Harmenberg's definition of the density of the permanent-income-weighted invariant distribution of m; call this density $\tilde{f}$. To understand $\tilde{f}$, we can see how total aggregate market resources held by people with given m will be:

$$\boldsymbol{M}_{t} = \boldsymbol{P}_{t}\tilde{f}(m)m \tag{35}$$

By implication of Theorem 7, $\boldsymbol{M}_t$ grows at a rate $\mathcal{G}$. We will now use this property of $\tilde{f}$ to show that aggregate consumption also grows at rate $\mathcal{G}$. Call $\boldsymbol{C}_t(m)$ the total amount of consumption at date t by persons with market resources m, and note that in the invariant economy this is given by the converged consumption function $\mathrm{c}(m)$ multiplied

[36] Formally, fix an individual i and let $\{\tilde{c}_t^i\}_{t=0}^{\infty}$ and $\{\tilde{m}_t^i\}_{t=0}^{\infty}$ be a stochastic recursive sequence generated by the converged consumption rule as follows, $\tilde{c}_t^i = \mathrm{c}(\tilde{m}_t^i)$ and $\tilde{m}_{t+1}^i = \tilde{\mathcal{R}}_{t+1}^i(\tilde{m}_t^i - \mathrm{c}(\tilde{m}_t^i)) + \boldsymbol{\xi}_{t+1}^i$, where the sequence of exogenous shocks are each defined on a *theoretical probability space* $(\Omega, \Sigma, \mathbb{P})$. Integration with respect to the measure $\mathbb{P}$ in the expected value operator $\mathbb{E}$ will be equivalent to *empirical* integration $\mathbb{M}$ with respect to a suitable measure of agents on a nonatomic agent space. In particular, for all j, $\mathbb{E}\mathrm{g}(\tilde{c}_t^j) = \int \tilde{c}_t \, d\mathbb{P} = \mathbb{M}\mathrm{g}(\tilde{c}_t) \colon = \int \mathrm{g}(\tilde{c}_t^i)\lambda(di)$, where λ is the measure of agents and for any measurable function g. For technical steps required to assert this claim, see **?**, which utilizes relatively recent results by **?** and also the detailed construction by **?**.

by the amount of permanent income accruing to such people $\tilde{f}(m)\boldsymbol{P}_t$. Since $\tilde{f}(m)$ is invariant and aggregate permanent income grows according to $\boldsymbol{P}_{t+1} = \mathcal{G}\boldsymbol{P}_t$, for any m, the following characterizes the growth of total consumption:

$$\begin{aligned}\log \boldsymbol{C}_{t+1}(m) - \log \boldsymbol{C}_t(m) &= \log \mathrm{c}(m)\tilde{f}(m)\boldsymbol{P}_{t+1} - \log \mathrm{c}(m)\tilde{f}(m)\boldsymbol{P}_t \\ &= \log \mathcal{G}.\end{aligned}$$

4.2 Aggregate Balanced Growth and Idiosyncratic Covariances

{subsec:Covariances

Harmenberg shows that the covariance between the individual consumption ratio c and the idiosyncratic component of permanent income $\boldsymbol{p}$ does not shrink to zero; thus, covariances are another potential measurement for construction of microfoundations.

Consider a date-t Harmenberg-converged economy, and define the mean value of the consumption ratio as $\bar{c}_{t+n} \equiv \mathbb{M}[c_{t+n}]$. Normalizing period-$t$ aggregate permanent income to $\boldsymbol{P}_t = 1$, total consumption at $t+1$ and $t+2$ are

$$\begin{aligned}\boldsymbol{C}_{t+1} &= \mathbb{M}[c_{t+1}\boldsymbol{p}_{t+1}] = \bar{c}_{t+1}\mathcal{G}^1 + \mathrm{cov}_{t+1}(c_{t+1}, \boldsymbol{p}_{t+1}) \\ \boldsymbol{C}_{t+2} &= \mathbb{M}[c_{t+2}\boldsymbol{p}_{t+2}] = \bar{c}_{t+2}\mathcal{G}^2 + \mathrm{cov}_{t+2}(c_{t+2}, \boldsymbol{p}_{t+2})\end{aligned} \tag{36}$$

{eq:atp2vsatp1}

and Harmenberg's proof that $\boldsymbol{C}_{t+2} - \mathcal{G}\boldsymbol{C}_{t+1} = 0$ allows us to obtain:

$$(\bar{c}_{t+2} - \bar{c}_{t+1})\,\mathcal{G}^2 = \mathcal{G}\mathrm{cov}_{t+1} - \mathrm{cov}_{t+2}. \tag{37}$$

{eq:cNrmvsCov}

In a Szeidl-invariant economy, $\bar{c}_{t+2} = \bar{c}_{t+1}$, so the economy exhibits balanced growth in the covariance:

$$\mathrm{cov}_{t+2} = \mathcal{G}\mathrm{cov}_{t+1}. \tag{38}$$

The more interesting case is when the economy is Harmenberg- but not Szeidl-invariant. In that case, if the cov and the $\bar{c}$ terms have constant growth factors Ω_{cov} and $\Omega_{\bar{c}}$,[37] an equation corresponding to (37) will hold in $t+n$:

$$\begin{aligned}(\overbrace{\Omega_{\bar{c}}^n \bar{c}_t}^{\bar{c}_{t+n}} - \Omega_{\bar{c}}^{n-1}\bar{c}_t)\mathcal{G}^n &= \left(\mathcal{G}\Omega_{\mathrm{cov}}^{n-1} - \Omega_{\mathrm{cov}}^n\right)\mathrm{cov}_t \\ (\Omega_{\bar{c}}\mathcal{G})^{n-1}(\Omega_{\bar{c}} - 1)\bar{c}_t\mathcal{G} &= \Omega_{\mathrm{cov}}^{n-1}(\mathcal{G} - \Omega_{\mathrm{cov}})\mathrm{cov}_t\end{aligned} \tag{39}$$

{eq:aNrmGrovsCov

so for the LHS and RHS to grow at the same rates we need

$$\Omega_{\mathrm{cov}} = \Omega_{\bar{c}}\mathcal{G}. \tag{40}$$

{eq:aNrmGrovscovC

This is intuitive: In the Szeidl-invariant economy, it just reproduces our result above that the covariance exhibits balanced growth because $\Omega_{\bar{c}} = 1$. The revised result just says that in the Harmenberg case where the mean value $\bar{c}$ of the consumption ratio c

[37] This 'if' is a conjecture, not something proven by Harmenberg (or anyone else). But see appendix F for an example of a Harmenberg-invariant economy in which simulations suggest this proposition holds.

can grow, the covariance must rise in proportion to any ongoing expansion of $\bar{c}$ (as well as in proportion to the growth in $\boldsymbol{p}$).

4.3 Implications for Microfoundations

{subsec:microfound

Thus we have microeconomic propositions, for both growth factors and for covariances of observable variables,[38] that can be tested in either cross-section or panel microdata to judge (and calibrate) the microfoundations that should hold for any macroeconomic analysis that requires balanced growth for its conclusions.

At first blush, these points are reassuring; one of the most persuasive arguments for the agenda of building microfoundations of macroeconomics is that newly available 'big data' allow us to measure cross-sectional covariances with great precision, so that we can use microeconomic natural experiments to disentangle questions that are hopelessly entangled in aggregate time-series data. Knowing that such covariances ought to be a stable feature of a stably growing economy is therefore encouraging.

But this discussion also highlights an uncomfortable point: In the model as specified, permanent income does not have a limiting distribution; it becomes ever more dispersed as the economy with infinite-horizon consumers continues to exist indefinitely.

A few microeconomic data sources attempt direct measurement of 'permanent income'; Carroll et al. [2017], for example, show that their assumptions about the magnitude of permanent shocks (and mortality; see below) yield a simulated distribution of permanent income that roughly matches answers in the U.S. Survey of Consumer Finances ('SCF') to a question designed to elicit a direct measure of respondents' permanent income. They use those results to calibrate a model to match empirical facts about the distribution of permanent income and wealth, showing that the model also does fits empirical facts about the marginal propensity to consume. The quantitative credibility of the argument depends on the model's match to the distribution of permanent income inequality, which would not be possible in a model without a nondegenerate steady-state distribution of permanent income.

For macroeconomists who want to build microfoundations by comparing the microeconomic implications of their models to micro data (directly – not in ratios to difficult-to-meaure 'permanent income'), it would be something of a challenge to determine how to construct empirical-data-comparable simulated results from a model with no limiting distribution of permanent income.

Death can solve this problem.

4.4 Mortality Yields Invariance

{sec:Mortality}

Most heterogeneous-agent models incorporate a constant positive probability of death, following Blanchard [1985] and Yaari [1965]. In the Blanchardian model, if the probability of death exceeds a threshold that depends on the size of the permanent shocks, Carroll et al. [2017] show that the limiting distribution of permanent income has a finite variance.

[38] Parallel results to those for consumption can be obtained for other measures like market assets.

Blanchard [1985] assumes a universal annuitization scheme in which estates of dying consumers are redistributed to survivors in proportion to survivors' wealth, giving the recipients a higher effective rate of return. This treatment has considerable analytical advantages, most notably that the effect of mortality on the time preference factor is the exact inverse of its effect on the (effective) interest factor. That is, if the 'pure' time preference factor is β and probability of remaining alive (not dead) is $\mathcal{L}$, then the assumption that no utility accrues after death makes the effective discount factor $\underline{\beta} = \beta\mathcal{L}$ while the enhancement to the rate of return from the annuity scheme yields an effective interest factor $\bar{\mathsf{R}} = \mathsf{R}/\mathcal{L}$ (recall that because of white-noise mortality, the average wealth of the two groups is identical). Combining these, the effective patience factor in the new economy $\underline{\beta}\bar{\mathsf{R}}$ is unchanged from its value in the infinite-horizon model:

$$\underline{\beta}\bar{\mathsf{R}} = (\beta\mathcal{L}\mathsf{R}/\mathcal{L})^{1/\gamma} = (\mathsf{R}\beta)^{1/\gamma} = \mathbb{Þ}. \tag{41}$$

The only adjustments this requires to the analysis above are therefore to the few elements that involve a role for the interest factor distinct from its contribution to $\mathbb{Þ}$ (principally, the RIC, which becomes $\mathbb{Þ}/\bar{\mathsf{R}}$).

Blanchard [1985]'s innovation was valuable not only for the insight it provided but also because when he wrote, the principal alternative, the Life Cycle model of Modigliani [1966], was computationally challenging given then-available technologies. Despite its (considerable) conceptual value, Blanchard's analytical solution is now rarely used because essentially all modern modeling incorporates uncertainty, constraints, and other features that rule out analytical solutions anyway.

The simplest alternative to Blanchard is to follow Modigliani in constructing a realistic description of income over the life cycle and assuming that any wealth remaining at death occurs accidentally (not implausible, given the robust finding that for the great majority of households, bequests amount to less than 2 percent of lifetime earnings, Hendricks [2001, 2016]).

Even if bequests are accidental, a macroeconomic model must make some assumption about how they are disposed of: As windfalls to heirs, estate tax proceeds, etc. We again consider the simplest choice, because it represents something of a polar alternative to Blanchard. Without a bequest motive, there are no behavioral effects of a 100 percent estate tax; we assume such a tax is imposed and that the revenues are effectively thrown in the ocean: The estate-related wealth effectively vanishes from the economy.

The chief appeal of this approach is the simplicity of the change it makes in the condition required for the economy to exhibit a balanced growth equilibrium (for consumers without a life cycle income profile). If $\mathcal{L}$ is the probability of remaining alive, the condition changes from the plain growth impatience to a looser mortality-adjusted version of growth impatience:

$$\mathcal{L}\mathbb{Þ}_{\mathcal{G}} < 1. \tag{42}$$ {eq:GICLivMod}

With no income growth, what is required to prohibit unbounded growth in aggregate wealth is the condition that prevents the per-capita wealth-to-permanent-income ratio

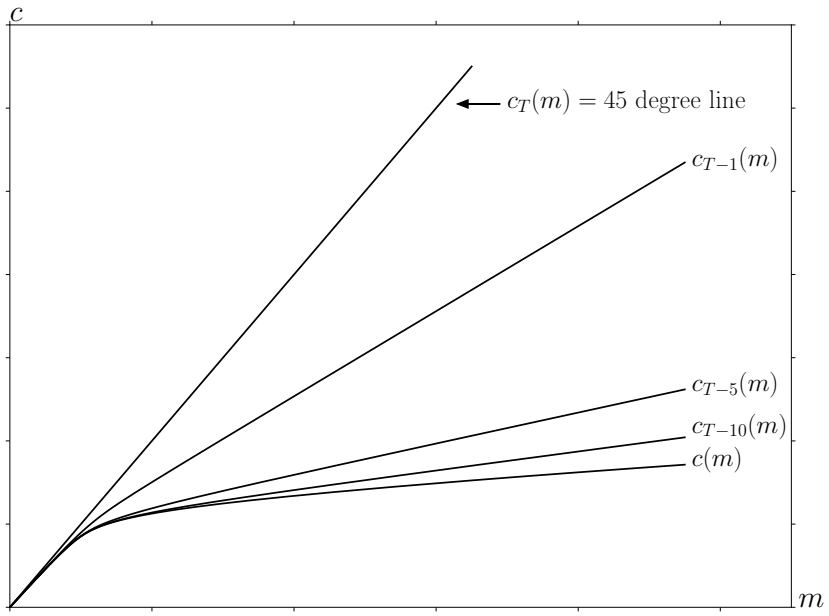

Figure 3 Convergence of the Consumption Rules {fig:cFuncsConverge

of surviving consumers from growing faster than the rate at which mortality diminishes their collective population. With income growth, the aggregate wealth-to-income ratio will head to infinity only if a cohort of consumers is patient enough to make the desired rate of growth of wealth fast enough to counteract combined erosive forces of mortality and productivity.

5 Patience and Limiting Consumer Behavior

{sec:GICdiscussion}

Having established our formal results, we are ready to describe how the various patience conditions determine the characteristics of the limiting consumption function. To fix ideas, we start with a quantitative example using the familiar benchmark case where return impatience, growth impatience and finite human wealth all hold, shown by Figure 3. The figure depicts the successive consumption rules that apply in the last period of life (c_T), the second-to-last period, and earlier periods under parameter values listed in Table 2. (The 45 degree line is $\mathrm{c}_T(m) = m$ because in the last period of life it is optimal to spend all remaining resources.)

Under the same parameter values, Figures 4–5 capture the theoretical bounds and MPCs of the converged consumption rule. In Figure 4, as m rises, the marginal propensity to consume approaches $\underline{\kappa} = (1 - \frac{\mathbf{\mathfrak{p}}}{\mathsf{R}})$ as $m \to \infty$, the same as the perfect foresight MPC. Moreover, as m approaches zero, the MPC approaches $\overline{\kappa} = (1 - q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}})$.

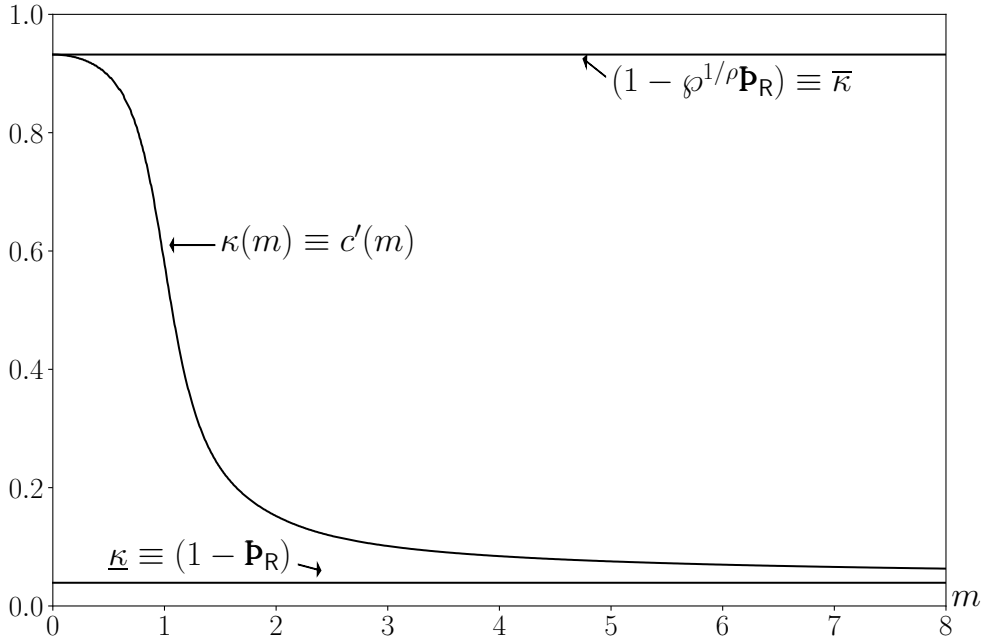

Figure 4 Limiting MPC's {fig:mpclimits}

While in the presence of a constraint neither return impatience nor growth impatience is individually necessary for nondegeneracy of $\mathrm{c}(m)$, a key conclusion of this section is that if both return impatience and growth impatience fail, the consumption function will be degenerate (limiting either to $\mathrm{c}(m) = 0$ or $c(m) = \infty$ as the horizon recedes). So, for a useful solution, at least one of these conditions must hold.[39] The case with growth impatience but return patience is particularly surprising, because it is not immediately clear what prevents our earlier conclusion (in other circumstances) that return patience leads $\mathrm{c}(m)$ to asymptote to zero. The trick is to note that if return patience holds, $\mathsf{R} < \mathbf{\Phi}$, while failure of growth impatience means $\mathbf{\Phi} < \mathcal{G}$; together these inequalities tell us that $\mathsf{R} < \mathcal{G}$ so (limiting) human wealth is infinite.[40] But, if at any m human wealth is unbounded, what prevents c from asymptoting to $\mathrm{c}(m) = \infty$ as the horizon gets arbitrarily long? This is where the natural borrowing constraint comes in. We will show that growth impatience is sufficient, at any fixed m, to guarantee an upper bound to $\mathrm{c}(m)$. The insight is best understood by first abstracting from uncertainty and studying the perfect foresight case (with and without constraints).

[39] Recall Claim 1 showing that a double-impatience failure implies autarky value is not finite; and see

[40] This logic holds even if both R and $\mathcal{G}$ are less than one – in this case, because the agent can borrow at a negative interest rate and always repay with income that shrinks more slowly than their debt.

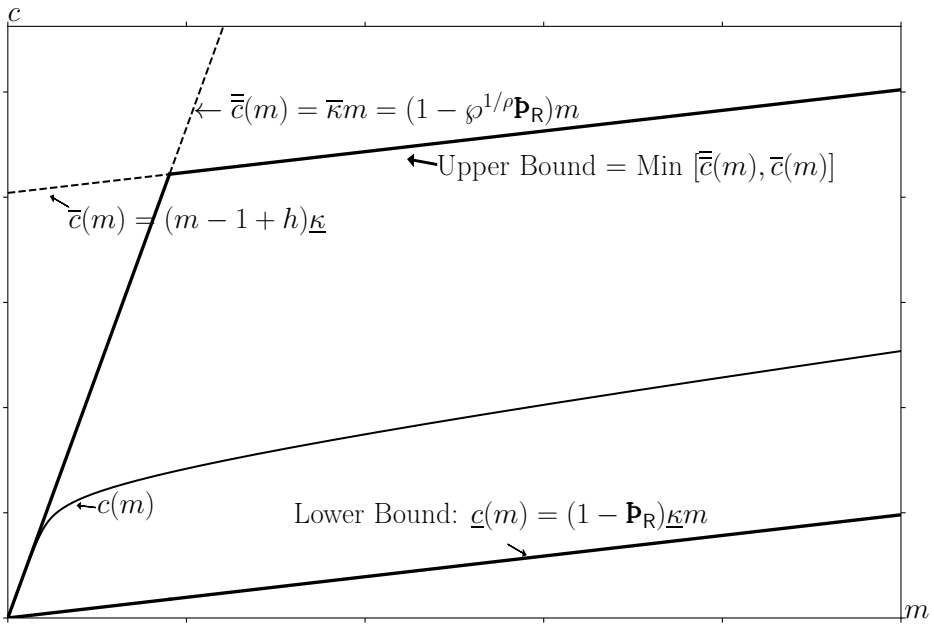

Figure 5 Upper and Lower Bounds on the Consumption Function {fig:cFuncBounds}

5.1 Model with Perfect Foresight {subsec:PFBdiscuss}

Claims 1-2 established the relationship between the finite value of autarky, return impatience and growth impatience in the context of a model with uncertainty. The easiest way to grasp the relations among these conditions is by studying Figure 6. Each node represents a quantity defined above. The arrow associated with each inequality imposes the condition, which is defined by the originating quantity being smaller than the arriving quantity. For example, one way we wrote the finite value of autarky (under perfect foresight) in Equation (9) is $\mathbb{P} < \mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$, so imposition of finite value of autarky is captured by the diagonal arrow connecting $\mathbb{P}$ and $\mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$. Traversing the boundary of the diagram clockwise starting at $\mathbb{P}$ involves imposing first growth impatience ($\mathbb{P} < \mathcal{G}$) then finite human wealth ($\mathcal{G} < \mathcal{G}(\mathsf{R}/\mathcal{G})^{1/\gamma} \longleftrightarrow \mathcal{G} < \mathsf{R}$), and the consequent arrival at the bottom right node tells us that these two conditions jointly imply perfect-foresight-finite-value-of-autarky. Reversal of a condition reverses the arrow's direction; so, for example, the bottom-most arrow going rightwards to $\mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$ implies finite human wealth fails; but we can cancel the cancellation and reverse the arrow. This would allow us to traverse the diagram clockwise from $\mathbb{P}$ through $\mathcal{G}$ to $\mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$ to R, revealing that imposition of growth impatience and finite human wealth (and, redundantly, finite human wealth again) let us conclude that return

impatience holds because the starting point is $\Þ$ and the endpoint is R (and we have traversed a chain of 'is greater than' relations).[41]

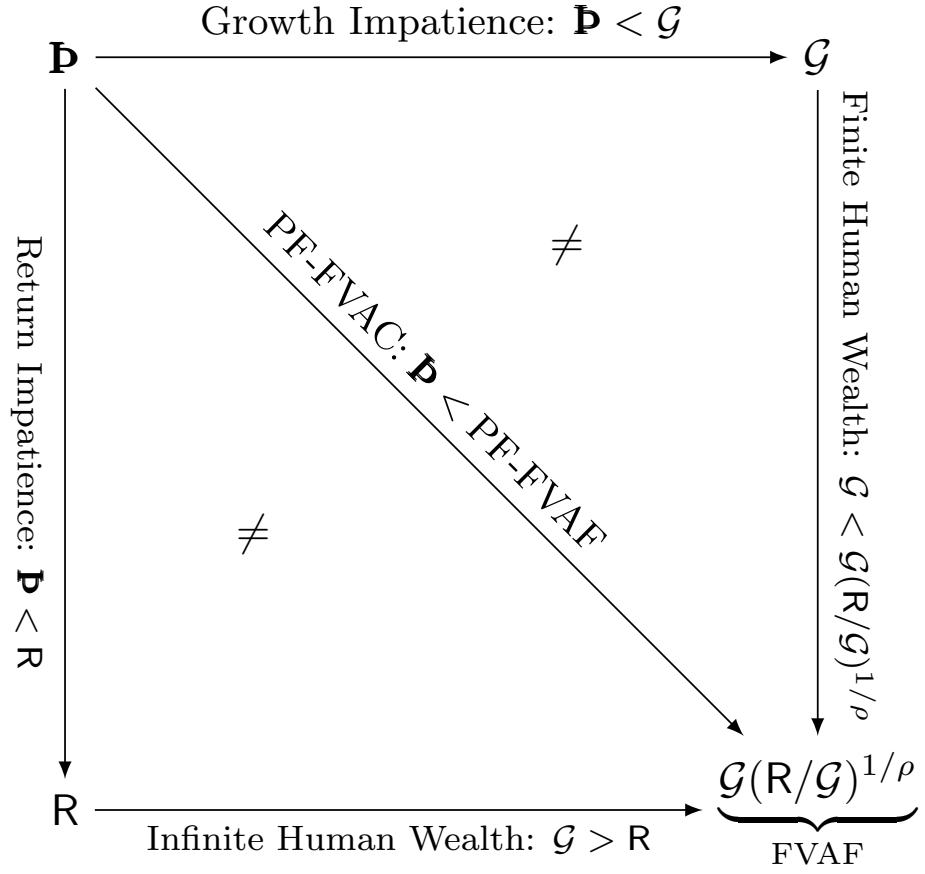

Figure 6 Perfect Foresight Relation of GIC, FHWC, RIC, and PFFVAC
An arrowhead points to the larger of the two quantities being compared; so, following the diagonal arrow imposes that absolute patience is smaller than the limit defined by the finite value of autarky factor, $\Þ < \mathcal{G}(\mathsf{R}/\mathcal{G})^{1/\gamma}$ (this is one way of writing the PF-FVAC, equation (9)). (The $\neq$ symbols indicate that the diagram is not commutative; that is, the different ways of reaching the conclusion that the PF-FVAC holds are not equivalent to each other).

In the unconstrained case, finite human wealth was necessary since, without constraints, only this condition could prevent infinite borrowing in the limit (Proposition 1). Looking at Figure 6, following the diagonal from $\Þ$ to the bottom-right corner corresponds to the direct of imposition of the finite value of autarky, which implies that the existence of a non-degenerate solution *requires* return impatience to hold. To see why, if return impatience failed, proceeding clockwise from the bottom left node of R would lead to $\mathsf{R} > \mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$, (equivalently $(\mathcal{G}/\mathsf{R})^{1-1/\gamma} < 1$) which corresponds to failure of finite human wealth (see also Case 3 in Section 5.2.1).

[41] Consult Appendix E for an exposition of diagrams of this type, which are a simple application of Category Theory (Riehl [2017]).

We can understand how failure of finite human wealth leads to infinite borrowing thinking about growth impatience. From Figure 6, let finite value of autarky hold (traverse the diagonal from $\mathbf{Þ}$) and then reverse the downward arrow from $\mathcal{G}$, signifying the failure of finite human wealth, so that as the horizon extends and income grows faster than the rate at which it is discounted, there is no upper bound to the present discounted value of future income (cf. Equation (16)). But the cancellation of finite human wealth also indirectly implies that growth impatience holds $\mathbf{Þ} > \$Rfree^{1/\gamma}\mathcal{G}^{1-\gamma} > \mathcal{G}$ which tells us that this is a consumer who wants to spend out of their human wealth. And therefore, at any fixed level of market resources, there is no upper bound to how much the consumer would choose to borrow as the horizon recedes.

Thus, in the perfect foresight unconstrained model, return impatience is the only condition at our disposal that can prevent consumption from limiting to zero as the planning horizon recedes. However, when we impose a liquidity constraint, the range of admissible parameters becomes more interesting.

5.1.1 Perfect Foresight Constrained Solution

{subsec:PFCon}

We now sketch the perfect foresight constrained solution and demonstrate that a solution can exist either under return impatience or without return impatience but with growth impatience (Proposition 2). Our discussion proceeds by examining implications of possible configurations of the patience conditions. (Tables 3 and 4 codify.)

Case 1: Growth impatience fails and return impatience holds. If growth impatience fails but return impatience holds, Appendix C shows that, for some $m_{\#}$, with $0 < m_{\#} < 1$, an unconstrained consumer behaving according to the perfect foresight solution (98) would choose $c < m$ for all $m > m_{\#}$. In this case the solution to the constrained consumer's problem is simple; for any $m \geq m_{\#}$ the constraint does not bind (and will never bind in the future). For such m the constrained consumption function is identical to the unconstrained one. If the consumer were somehow[42] to arrive at an $m_{\#}$ such that $m < m_{\#} < 1$ the constraint would bind and the consumer would consume $c = m$. Using $\grave{\mathrm{c}}$ for the perfect foresight consumption function in the presence of constraints (and analogously for all other functions):

$$\grave{\mathrm{c}}(m) = \begin{cases} m & \text{if } m < m_{\#} \\ \bar{\mathrm{c}}(m) & \text{if } m \geq m_{\#} \end{cases}$$

where $\bar{\mathrm{c}}(m)$ is the unconstrained perfect foresight solution.

Case 2: Growth impatience holds and return impatience holds. When return impatience and growth impatience both hold, Appendix C shows that the limiting constrained consumption function is piecewise linear, with $\grave{\mathrm{c}}(m) = m$ up to a first

[42]"Somehow" because $m < 1$ could only be obtained by entering the period with $b < 0$ which the constraint forbids.

'kink point' at $m^0_{\#} > 1$, and with discrete declines in the MPC at a set of kink points $\{m^1_{\#}, m^2_{\#}, \ldots\}$. As $m \to \infty$ the constrained consumption function $\grave{\mathrm{c}}(m)$ becomes arbitrarily close to the unconstrained $\bar{\mathrm{c}}(m)$, and the marginal propensity to consume, $\grave{\mathrm{c}}'(m)$, limits to $\underline{\kappa}$.[43] Similarly, the value function $\grave{\mathrm{v}}(m)$ is nondegenerate and limits to the value function of the unconstrained consumer.

This logic holds even when finite human wealth fails, because the constraint prevents the (limiting) consumer[44] from borrowing against unbounded human wealth to finance unbounded current consumption. Under these circumstances, the consumer who starts with any $b_t > 0$ will, over time, run those resources down so that after some finite number of periods τ the consumer will reach $b_{t+\tau} = 0$, and thereafter will set $\boldsymbol{c} = \boldsymbol{p}$ for eternity (which finite value of autarky says yields finite value). Using the same steps as for Equation (95), value of the interim program is also finite:

$$\boldsymbol{v}_{t+\tau} = \mathcal{G}^{\tau(1-\gamma)}\mathrm{u}(\boldsymbol{p}_t)\left(\frac{1-(\beta\mathcal{G}^{1-\gamma})^{T-(t+\tau)+1}}{1-\beta\mathcal{G}^{1-\gamma}}\right).$$

So, even when finite human wealth fails, the limiting consumer's value for any finite m will be the sum of two finite numbers: One due to the unconstrained choice made over the finite-horizon leading up to $b_{t+\tau} = 0$, and one reflecting the value of consuming $\boldsymbol{p}_{t+\tau}$ thereafter.

Case 3: Growth impatience holds and return impatience fails. The most peculiar possibility occurs only when return impatience fails. As noted above, this possibility is unavailable to us without a constraint. Without return impatience, finite human wealth must also fail (Appendix C), and the constrained consumption function is (surprisingly) nondegenerate. (See appendix Figure 8 for a numerical example). Even though human wealth is unbounded at any given level of m, since borrowing is ruled out, consumption cannot become unbounded at that m in the limit as the horizon recedes. However, the failure of return impatience does have some power: It means that as m rises without bound, the MPC approaches zero ($\lim_{m\to\infty} \grave{\mathrm{c}}'(m) = 0$). Nevertheless $\grave{\mathrm{c}}(m)$ is finite, strictly positive, and strictly increasing in m. This result reconciles the conflicting intuitions from the unconstrained case, where failure of return impatience would suggest a degenerate limit of $\grave{\mathrm{c}}(m) = 0$ while failure of finite human wealth would suggest a degenerate limit of $\grave{\mathrm{c}}(m) = \infty$.

5.2 Model with Uncertainty

{subsec:TheModelU

We now examine the case with uncertainty but without constraints, which we argued was a close parallel to the model with constraints but without uncertainty (recall Section 2.4.3).

[43] See Carroll et al. [2019] for details.

[44] That is, one obeying $\mathrm{c}(m) = \lim_{n\to\infty} \mathrm{c}_{t-n}(m)$.

Table 1 Microeconomic Model Calibration

{table:Parameters}

Calibrated Parameters			
Description	Parameter	Value	Source
Permanent Income Growth Factor	$\mathcal{G}$	1.03	PSID: Carroll (1992)
Interest Factor	R	1.04	Conventional
Time Preference Factor	β	0.96	Conventional
Coefficient of Relative Risk Aversion	γ	2	Conventional
Probability of Zero Income	q	0.005	PSID: Carroll (1992)
Std Dev of Log Permanent Shock	σ_ψ	0.1	PSID: Carroll (1992)
Std Dev of Log Transitory Shock	σ_θ	0.1	PSID: Carroll (1992)

Table 2 Model Characteristics Calculated from Parameters

{table:Calibration}

Description	Symbol and Formula			Approximate Calculated Value
Finite Human Wealth Factor	$\tilde{\mathcal{R}}^{-1}$	$\equiv$	$\mathcal{G}/\mathsf{R}$	0.990
PF Value of Autarky Factor	$\beth$	$\equiv$	$\beta\mathcal{G}^{1-\gamma}$	0.932
Growth Compensated Permanent Shock	$\underline{\psi}$	$\equiv$	$(\mathbb{E}[\psi^{-1}])^{-1}$	0.990
Uncertainty-Adjusted Growth	$\underline{\mathcal{G}}$	$\equiv$	$\mathcal{G}\underline{\psi}$	1.020
Utility Compensated Permanent Shock	$\underline{\underline{\psi}}$	$\equiv$	$(\mathbb{E}[\psi^{1-\gamma}])^{1/(1-\gamma)}$	0.990
Utility Compensated Growth	$\underline{\mathcal{G}}$	$\equiv$	$\mathcal{G}\underline{\underline{\psi}}$	1.020
Absolute Patience Factor	$\mathbf{\mathfrak{p}}$	$\equiv$	$(\mathsf{R}\beta)^{1/\gamma}$	0.999
Return Patience Factor	$\mathbf{\mathfrak{p}}_{\mathsf{R}}$	$\equiv$	$\mathbf{\mathfrak{p}}/\mathsf{R}$	0.961
Growth Patience Factor	$\mathbf{\mathfrak{p}}_{\mathcal{G}}$	$\equiv$	$\mathbf{\mathfrak{p}}/\mathcal{G}$	0.970
Modified Growth Patience Factor	$\mathbf{\mathfrak{p}}_{\mathcal{G}}\mathbb{E}[\psi^{-1}]$	$\equiv$	$\mathbf{\mathfrak{p}}/\underline{\mathcal{G}}$	0.980
Value of Autarky Factor	$\underline{\beth}$	$\equiv$	$\beta\mathcal{G}^{1-\gamma}\underline{\underline{\psi}}^{1-\gamma}$	0.941
Weak Return Impatience Factor	$q^{1/\gamma}\mathbf{\mathfrak{p}}$	$\equiv$	$(q\beta\mathsf{R})^{1/\gamma}$	0.071

Tables 1 and 2 present calibrations and values of model conditions in the case with uncertainty, where return impatience, growth impatience and finite value of autarky all hold. The full relationship among conditions is represented in Figure 7. Though the diagram looks complex, it is merely a modified version of the earlier simple diagram (Figure 6) with further (mostly intermediate) inequalities inserted. (Arrows with a "because" now label relations that always hold under the model's assumptions.)[45]

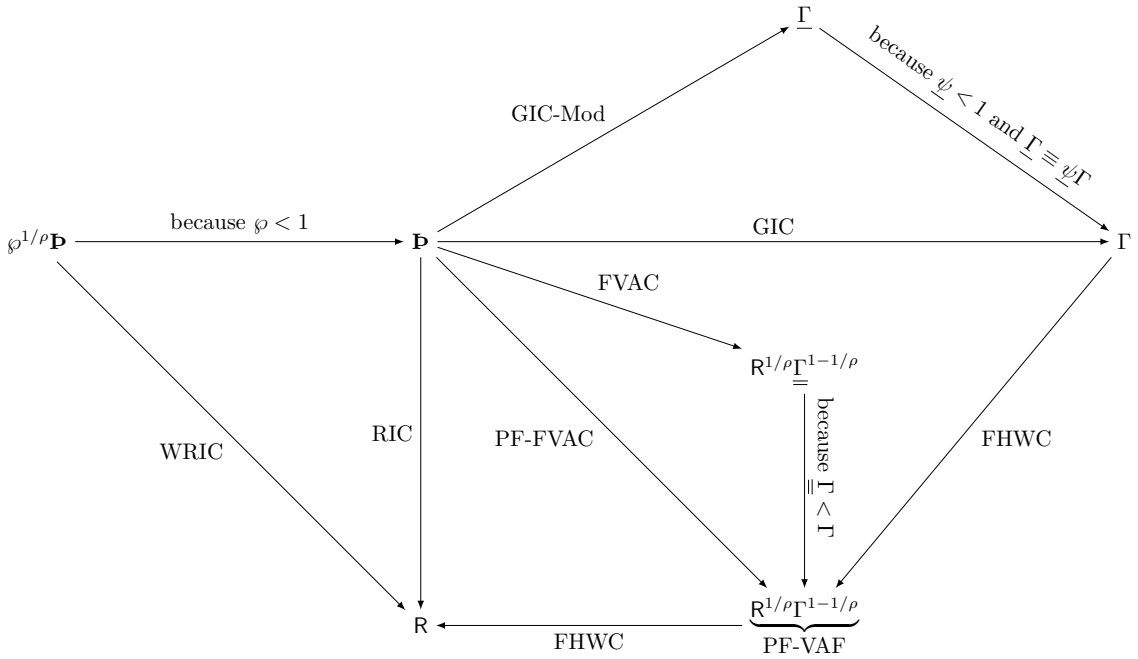

Figure 7 Relation of All Inequality Conditions
See Table 2 for Numerical Values of Nodes Under Baseline Parameters {fig:Inequalities}

Beyond finite value of autarky, the additional condition sufficient for contraction, weak return impatience, can be seen to be weak by asking 'under what circumstances would the finite value of autarky hold but the weak return impatience fail?' Algebraically, the requirement becomes:

$$\beta\mathcal{G}^{1-\gamma}\underline{\underline{\psi}}^{1-\gamma} < \ 1 \ < (q\beta)^{1/\gamma}/\mathsf{R}^{1-1/\gamma}. \tag{43}$$ {eq:WRICandFVAC

where $\underline{\underline{\psi}} := (\mathbb{E}[\psi^{1-\gamma}])^{1/(1-\gamma)} < 1$. If we require $\mathsf{R} \geq 1$, the weak return impatience is 'redundant' because now $\beta < 1 < \mathsf{R}^{\gamma-1}$, so that (with $\gamma > 1$ and $\beta < 1$) return impatience (and weak return impatience) must hold. But neither theory nor evidence demand that $\mathsf{R} \geq 1$. We can therefore approach the question of the relevance of weak return impatience by asking just how low R must be for the condition to be relevant. Suppose for illustration that $\gamma = 2$, $\underline{\underline{\psi}}^{1-\gamma} = 1.01$, $\mathcal{G}^{1-\gamma} = 1.01^{-1}$ and $q = 0.10$. In that

[45] Again, readers unfamiliar with such diagrams should see Appendix E for a more detailed exposition.

case (43) reduces to:

$$\beta < 1 < (0.1\beta/\mathsf{R})^{1/2},$$

but since $\beta < 1$ by assumption, the binding requirement becomes:

$$\mathsf{R} < \beta/10,$$

so that for example if $\beta = 0.96$ we would need $\mathsf{R} < 0.096$ (that is, a perpetual riskfree rate of return of worse than -90 percent a year) in order for weak return impatience to be nonredundant.

Perhaps the best way of thinking about this is to note that the space of parameter values for which the weak return impatience remains relevant shrinks out of existence as $q \to 0$, which Section 2.4.3 showed was the precise limiting condition under which behavior becomes arbitrarily close to the liquidity constrained solution (in the absence of other risks). On the other hand, when $q = 1$, the consumer has no noncapital income (so finite human wealth holds) and with $q = 1$ weak return impatience is identical to weak return impatience. However, weak return impatience is the only condition required for a solution to exist for a perfect foresight consumer with no noncapital income. Thus weak return impatience forms a sort of 'bridge' between the liquidity constrained and the unconstrained problems as q moves from 0 to 1.

5.2.1 Behavior Under Cases of Conditions

{subsubsec:casesUC

Case 1: Return impatience fails and growth impatience holds In the unconstrained perfect foresight problem (Section 5.1), return impatience was necessary for existence of a nondegenerate solution. It is surprising, therefore, that in the presence of uncertainty, the much weaker weak return impatience is sufficient for nondegeneracy (assuming that finite value of autarky holds). Given finite value of autarky, we can derive the features the problem must exhibit for return impatience to fail (that is, $\mathsf{R} < (\mathsf{R}\beta)^{1/\gamma}$) (given that growth impatience holds) as follows:

$$\begin{aligned} \mathsf{R} < (\mathsf{R}\beta)^{1/\gamma} \ &< \ (\mathsf{R}(\mathcal{G}\underline{\underline{\psi}})^{\gamma-1})^{1/\gamma} \\ \Rightarrow \mathsf{R} < (\mathsf{R}/\mathcal{G})^{1/\gamma}\mathcal{G}\underline{\underline{\psi}}^{1-1/\gamma} \qquad (44) \quad \{\text{eq:RICimplies}\} \\ \Rightarrow \mathsf{R}/\mathcal{G} < \underline{\underline{\psi}} \end{aligned}$$

but since $\underline{\underline{\psi}} < 1$ (for $\gamma > 1$ and nondegenerate ψ), this requires $\mathsf{R}/\mathcal{G} < 1$. Thus, given finite value of autarky, return impatience can fail only if human wealth is unbounded and growth impatience holds.[46]

[46]This algebraically complicated conclusion could be easily reached diagrammatically in Figure 7 by starting at the R node and imposing the failure of return impatience, which reverses the return impatience arrow and lets us traverse the diagram along any clockwise path to the perfect foresight

As in the perfect foresight constrained problem, unbounded limiting human wealth here does not lead to a degenerate limiting consumption function (finite human wealth is not required for Theorem 2). But, from equation (19) and the discussion surrounding it, an implication of the failure of return impatience is that $\lim_{m\to\infty} \mathrm{c}'(m) = 0$. Thus, interestingly, in this case (unavailable in the perfect foresight unconstrained) model the presence of uncertainty both permits unlimited human wealth (in the $n \to \infty$ limit) and at the same time prevents unlimited human wealth from resulting in (limiting) infinite consumption (at any finite m). Intuitively, the utility-imposed 'natural constraint' that arises from the possibility of a zero income event prevents infinite borrowing and at the same time allows infinite human wealth to prevent patience from resulting, as it does under other conditions, in the degenerate $\mathrm{c}(m) = 0$ as the planning horizon recedes. Thus, in presence of uncertainty of the kind we assume, pathological patience (which in the perfect foresight model results in a limiting consumption function of $\mathrm{c}(m) = 0$) plus unbounded human wealth (which the perfect foresight model prohibits because it leads to a limiting consumption function $\mathrm{c}(m) = \infty$ for any finite m) combine to yield a unique finite limiting (as $n \to \infty$) level of consumption and MPC for any finite value of m.

Note the close parallel to the conclusion in the perfect foresight liquidity constrained model in the case where return impatience fails (Case 3 in Section 5.1.1). There, too, the tension between infinite human wealth and pathological patience was resolved with a nondegenerate consumption function whose limiting MPC was zero.[47]

Case 2: Return impatience holds and growth impatience holds with finite human wealth This is the benchmark case we presented at the start of the Section. If return impatience and finite human wealth both hold, a perfect foresight solution exists (Section 5.1). As $m \to \infty$ the limiting c and v functions become arbitrarily close to those in the perfect foresight model, because human wealth pays for a vanishingly small portion of spending (Section **??**).

Case 3: Return impatience holds and growth impatience holds with infinite human wealth The more exotic case is where finite human wealth fails but both growth impatience and return impatience also hold. In the unconstrained perfect foresight model, this is the degenerate case with limiting $\bar{\mathrm{c}}(m) = \infty$. Here, infinite human wealth and finite value of autarky implies that (perfect foresight) finite value of autarky holds and that $\mathbb{P} < \mathcal{G}$. To see why, traverse Figure 7 clockwise from $\mathbb{P}$ by imposing finite value of autarky to reach the PF-FVAF node. Because the bottom arrow pointing to the right, connecting the R and perfect foresight finite value of autarky nodes imposes the failure of finite human wealth (and here we are assuming that condition holds), we can reverse the bottom arrow and traverse the resulting clockwise path from FVAC to

finite value of autarky node at which point we realize that we *cannot* impose finite human wealth because that would let us conclude $\mathsf{R} > \mathsf{R}$.

[47] Ma and Toda [2020] derive conditions under which the limiting MPC is zero in an even more general case where there is also capital income risk.

see that

$$\mathbb{Þ} < (\mathsf{R}/\mathcal{G})^{1/\gamma}\mathcal{G} \Rightarrow \mathbb{Þ} < \mathcal{G}$$

where the transition from the first to the second lines is justified because failure of finite human wealth implies $\Rightarrow (\mathsf{R}/\mathcal{G})^{1/\gamma} < 1$. So, under return impatience and finite human wealth, we must have growth impatience.

However, we are not entitled to conclude that strong growth impatience holds: $\mathbb{Þ} < \mathcal{G}$ does not imply $\mathbb{Þ} < \underline{\psi}\mathcal{G}$ where $\underline{\psi} < 1$.

We have now established the principal points of comparison between the perfect foresight solutions and the solutions under uncertainty; these are codified in the remaining parts of Tables 3 and 4.

Table 3 Definitions and Comparisons of

Perfect Foresight Versions	Uncertainty Versions
Finite Human Wealth	
$\mathcal{G}/\mathsf{R} < 1$ The growth factor for permanent income $\mathcal{G}$ must be smaller than the discounting factor R for human wealth to be finite.	$\mathcal{G}/\mathsf{R} < 1$ The model's risks are mean-preserving spreads, so the PDV of future income is unchanged by their introduction.
Absolute Impatience Condition	
$\mathbf{\mathfrak{P}} < 1$ The unconstrained consumer is sufficiently impatient that the level of consumption will be declining over time: $\boldsymbol{c}_{t+1} < \boldsymbol{c}_t$	$\mathbf{\mathfrak{P}} < 1$ *If wealth is large enough,* the *expectation* of consumption next period will be smaller than this period's consumption: $\lim_{m_t \to \infty} \mathbb{E}_t[\boldsymbol{c}_{t+1}] < \boldsymbol{c}_t$
Return Impatience	
Return Impatience Condition	Weak RIC (WRIC)
$\mathbf{\mathfrak{P}}/\mathsf{R} < 1$ The growth factor for consumption $\mathbf{\mathfrak{P}}$ must be smaller than the discounting factor R, so that the PDV of current and future consumption will be finite: $\mathrm{c}'(m) = 1 - \mathbf{\mathfrak{P}}/\mathsf{R} < 1$	$q^{1/\gamma}\mathbf{\mathfrak{P}}/\mathsf{R} < 1$ If the probability of the zero-income event is $q = 1$ then income is always zero and the condition becomes identical to the RIC. Otherwise, weaker. $\mathrm{c}'(m) < 1 - q^{1/\gamma}\mathbf{\mathfrak{P}}/\mathsf{R} < 1$
Growth Impatience	
GIC	GIC-Mod
$\mathbf{\mathfrak{P}}/\mathcal{G} < 1$ For an unconstrained PF consumer, the ratio of $\boldsymbol{c}$ to $\boldsymbol{p}$ will fall over time. For constrained, guarantees the constraint eventually binds. Guarantees $\lim_{m_t \uparrow \infty} \mathbb{E}_t[\psi_{t+1} m_{t+1}/m_t] = \frac{\mathbf{\mathfrak{P}}}{\mathcal{G}}$	$\mathbf{\mathfrak{P}}\mathbb{E}[\psi^{-1}]/\mathcal{G} < 1$ By Jensen's inequality stronger than GIC. Ensures consumers will not expect to accumulate m unboundedly. $\lim_{m_t \to \infty} \mathbb{E}_t[m_{t+1}/m_t] = \frac{\mathbf{\mathfrak{P}}}{\mathcal{G}}\mathbb{E}[\psi^{-1}]$
Finite Value of Autarky	
PF-FVAC	FVAC
$\beta\mathcal{G}^{1-\gamma} < 1$ equivalently $\mathbf{\mathfrak{P}} < \mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$ The discounted utility of constrained consumers who spend their permanent income each period should be finite.	$\beta\mathcal{G}^{1-\gamma}\mathbb{E}[\psi^{1-\gamma}] < 1$ By Jensen's inequality, stronger than the PF-FVAC because for $\gamma > 1$ and nondegenerate ψ, $\mathbb{E}[\psi^{1-\gamma}] > 1$.

Table 4 Sufficient Conditions for Nondegenerate[‡] Solution

Consumption Model(s)	Conditions	Comments
$\bar{c}(m)$: PF Unconstrained	RIC, FHWC°	RIC⇒ $\lvert v(m)\rvert < \infty$; FHWC⇒ $0 < \lvert v(m)\rvert$
$\underline{c}(m) = \underline{\kappa} m$		PF model with no human wealth ($h = 0$)
Section 2.3.1:		RIC prevents $\bar{c}(m) = \underline{c}(m) = 0$
Section 2.3.1:		FHWC prevents $\bar{c}(m) = \infty$
Eq (47):		PF-FVAC+FHWC ⇒ RIC
Eq (46):		GIC+FHWC ⇒ PF-FVAC
$\grave{c}(m)$: PF Constrained	~~GIC~~, RIC	FHWC holds ($\mathcal{G} < \bar{\mathbf{P}} < \mathsf{R} \Rightarrow \mathcal{G} < \mathsf{R}$)
Section 5.1.1:		$\grave{c}(m) = \bar{c}(m)$ for $m > m_{\#} < 1$
		(~~RIC~~ would yield $m_{\#} = 0$ so $\grave{c}(m) = 0$)
Appendix C:	GIC,RIC	$\lim_{m\to\infty} \grave{c}(m) = \bar{c}(m)$, $\lim_{m\to\infty} \grave{\boldsymbol{\kappa}}(m) = \underline{\kappa}$
		kinks where horizon to $b = 0$ changes*
Appendix C:	GIC,~~RIC~~	$\lim_{m\to\infty} \grave{\boldsymbol{\kappa}}(m) = 0$
		kinks where horizon to $b = 0$ changes*
$c(m)$: Friedman/Muth	Section 2.4.1 & 2.4.2 , Section 3.1	$\underline{c}(m) < c(m) < \bar{c}(m)$
		$\underline{v}(m) < v(m) < \bar{v}(m)$
Section **??**:	FVAC, WRIC	Sufficient for Contraction
Section **??**:		WRIC is weaker than RIC
Figure 7:		FVAC is stronger than PF-FVAC
Section **??**: Case 3		~~FHWC~~+RIC ⇒GIC, $\lim_{m\to\infty} \boldsymbol{\kappa}(m) = \underline{\kappa}$
Section **??**: Case 1		~~RIC~~ ⇒~~FHWC~~, $\lim_{m\to\infty} \boldsymbol{\kappa}(m) = 0$
Section 3.2:		"Buffer Stock Saving" Conditions
Theorem 6:		GIC ⇒ $\exists\ \check{m}$ s.t. $0 < \check{m} < \infty$
Theorem 7:		GIC-Mod ⇒ $\exists\ \hat{m}$ s.t. $0 < \hat{m} < \infty$

[‡]For feasible m satisfying $0 < m < \infty$, a nondegenerate limiting consumption function defines a unique optimal value of c satisfying $0 < c(m) < \infty$; a nondegenerate limiting value function defines a corresponding unique value of $-\infty < v(m) < 0$.

°RIC, FHWC are necessary as well as sufficient for the perfect foresight case. *That is, the first kink point in $c(m)$ is $m_{\#}$ s.t. for $m < m_{\#}$ the constraint will bind now, while for $m > m_{\#}$ the constraint will bind one period in the future. The second kink point corresponds to the m where the constraint will bind two periods in the future, etc.

**In the Friedman/Muth model, the RIC+FHWC are sufficient, but *not* necessary for nondegeneracy

6 Conclusions

Numerical solutions to optimal consumption problems, in both life cycle and infinite-horizon contexts, have become standard tools since the first reasonably realistic models were constructed in the late 1980s. One contribution of this paper is to show that finite-horizon ('life cycle') versions of the simplest such models, with assumptions about income shocks (transitory and permanent) dating back to Friedman [1957] and standard specifications of preferences — and without plausible (but computationally and mathematically inconvenient) complications like liquidity constraints — have attractive properties (like continuous differentiability of the consumption function, and analytical limiting MPC's as resources approach their minimum and maximum possible values).

The main focus of the paper, though, is on the limiting solution of the finite-horizon model as the time horizon approaches infinity. This simple model has other appealing features: A 'Finite Value of Autarky' condition guarantees convergence of the consumption function, under the mild additional requirement of a 'Weak Return Impatience Condition' that will never bind for plausible parameterizations, but provides intuition for the bridge between this model and models with explicit liquidity constraints. The paper also provides a roadmap for the model's relationships to the perfect foresight model without and with constraints. The constrained perfect foresight model provides an upper bound to the consumption function (and value function) for the model with uncertainty, which explains why the conditions for the model to have a nondegenerate solution closely parallel those required for the perfect foresight constrained model to have a nondegenerate solution.

The main use of infinite-horizon versions of such models is in heterogeneous-agent macroeconomics. The paper articulates intuitive 'Growth Impatience Conditions' under which populations of such agents, with Blanchardian (tighter) or Modiglianian (looser) mortality will exhibit balanced growth. Finally, the paper provides the analytical basis for many results about buffer-stock saving models that are so well understood that even without analytical foundations researchers uncontroversially use them as explanations of real-world phenomena like the cross-sectional pattern of consumption dynamics in the Great Recession.

Appendices

Below are the paper's appendices referenced in the text.

A c Functions Exist, are Concave, and Differentible

{sec:ApndxConcave

A.1 Appendix for Problem Formulation

A.1.1 Recovering the Non-Normalized Problem

{sec:recoverLevels}

Letting nonbold variables be the boldface counterpart normalized by $\boldsymbol{p}_t$ (as with $m = \boldsymbol{m}/\boldsymbol{p}$), consider the problem in the second-to-last period:

$$\begin{aligned}\mathbf{v}_{T-1}(\boldsymbol{m}_{T-1}, \boldsymbol{p}_{T-1}) &= \max_{0<c_{T-1}\leq m_{T-1}} \mathrm{u}(\boldsymbol{p}_{T-1}c_{T-1}) + \beta\mathbb{E}_{T-1}[\mathrm{u}(\boldsymbol{p}_T m_T)] \\ &= \boldsymbol{p}_{T-1}^{1-\gamma}\left\{\max_{0<c_{T-1}\leq m_{T-1}} \mathrm{u}(c_{T-1}) + \beta\mathbb{E}_{T-1}[\mathrm{u}(\tilde{\mathcal{G}}_T m_T)]\right\}. \end{aligned} \tag{45}$$ {eq:vBold}

Since $\mathrm{v}_T(m_T) = \mathrm{u}(m_T)$, defining $\mathrm{v}_{T-1}(m_{T-1})$ from Problem $(\mathcal{P}_N)$, we obtain:

$$\mathbf{v}_{T-1}(\boldsymbol{m}_{T-1}, \boldsymbol{p}_{T-1}) = \boldsymbol{p}_{T-1}^{1-\gamma}\mathrm{v}_{T-1}(\underbrace{\boldsymbol{m}_{T-1}/\boldsymbol{p}_{T-1}}_{=m_{T-1}}).$$

This logic induces to earlier periods; if we solve the normalized one-state-variable problem $(\mathcal{P}_N)$, we will have solutions to the original problem for any $t < T$ from:

$$\begin{aligned}\mathbf{v}_t(\boldsymbol{m}_t, \boldsymbol{p}_t) &= \boldsymbol{p}_t^{1-\gamma}\mathrm{v}_t(m_t), \\ \boldsymbol{c}_t(\boldsymbol{m}_t, \boldsymbol{p}_t) &= \boldsymbol{p}_t\mathrm{c}_t(m_t).\end{aligned}$$

A.1.2 Challenges with Standard Dynamic Programming Approaches

{subsubsec:challeng

A.1.3 Infinite Horizon Stochastic Dynamic Optimization Problem

{subsubsec:infiniteS

How does the limiting nondegenerate solution connect to the solution of an infinite horizon stochastic dynamic optimization problem [**???**]? The two problems are equivalent when the converged value function, v, is a fixed point of the stationary Bellman operator $\mathbb{T}$, and the nondegenerate consumption function is v-greedy, that is, Equation (24) holds. Given the particular approach taken by Theorem 2, and to aid the interpretation of our discussion on aggregate relationships, we state the standard result formally and present a proof Appendix **??**.

Let a sequence of shocks $\{\psi_k, \boldsymbol{\xi}_k\}_{k=0}^{\infty}$ be defined on a common probability space, $(\Omega, \Sigma, \mathbb{P})$, and fix the problem primitives defined in Section 2.1. Consider the value

function for stochastic infinite horizon sequence problem:

$$\tilde{\mathrm{v}}(m) = \max_{\{\tilde{c}_k\}_{k=0}^{\infty}} \mathbb{E} \sum_{k=0}^{\infty} \beta^k \Pi_{j=0}^{k} \tilde{\mathcal{G}}_j \mathrm{u}(\tilde{c}_k), \qquad m \in S \qquad (\mathcal{P}_\infty) \quad \{\text{eq:probsequence}\}$$

such that i) $\{\tilde{c}_k\}_{k=0}^{\infty}$ is a sequence of random variables defined on $(\Omega, \Sigma, \mathbb{P})$, progressively measurable with respect to the shocks $\{\psi_k, \boldsymbol{\xi}_k\}_{k=0}^{\infty}$, ii) the inter-temporal budget constraint holds almost everywhere: $\tilde{m}_{k+1} = \tilde{\mathcal{R}}_k(\tilde{m}_k - \tilde{c}_k) + \boldsymbol{\xi}_k$, iii) the cannot die in debt condition holds almost everywhere in the limit: $\lim_{k\to\infty} \tilde{m}_k \geq 0$ and iv) $\tilde{m}_0 = m$. The expectation $\mathbb{E}$ is taken with respect to $\mathbb{P}$.

Proposition 6. *Let the assumptions of Theorem 2 hold. If* v *and* c *are a limiting nondegenerate solution, them* $\mathrm{v} = \tilde{\mathrm{v}}$ *and the sequence* $\{\tilde{c}_k\}_{k=0}^{\infty}$ *generated by* $\tilde{c}_k = \mathrm{c}(\tilde{m}_k)$*, where* $\tilde{m}_{k+1} = \tilde{\mathcal{R}}_k(\tilde{m}_k - \mathrm{c}(\tilde{m}_k)) + \boldsymbol{\xi}_k$*, solves Problem* $(\mathcal{P}_\infty)$*.*

The proposition, an implication of the Bellman Principle of Optimality, says that an individual following the nondegenerate consumption rule has maximized the expected discounted sum of their future per-period utilities.

A.2 Perfect Foresight Benchmarks

{sec:PFBProofs}

PFBProofs

How do the finite value of human wealth, perfect foresight finite value of autarky and return impatience relate to each other? If the FHWC is satisfied, the PF-FVAC implies that the RIC is satisfied.[48] Likewise, if the FHWC and the GIC are both satisfied, PF-FVAC follows:

$$\begin{aligned} &\mathbf{\mathfrak{p}} < \mathcal{G} < \mathsf{R} \\ &\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} < \mathcal{G}/\mathsf{R} < (\mathcal{G}/\mathsf{R})^{1-1/\gamma} < 1 \end{aligned} \qquad (46) \quad \{\text{eq:GICandFHWCi}$$

(the last line holds because FHWC $\Rightarrow 0 \leq (\mathcal{G}/\mathsf{R}) < 1$ and $\gamma > 1 \Rightarrow 0 < 1 - 1/\gamma < 1$).

Divide both sides of the second inequality in (9) by R:

$$\mathbf{\mathfrak{p}}/\mathsf{R} < (\mathcal{G}/\mathsf{R})^{1-1/\gamma} \qquad (47) \quad \{\text{eq:FHWCandPFF}$$

and FHWC $\Rightarrow$ the RHS is < 1 because $(\mathcal{G}/\mathsf{R}) < 1$ (and the RHS is raised to a positive power (because $\gamma > 1$)).

[48] Divide both sides of the second inequality in (9) by R:

$$\mathbf{\mathfrak{p}}/\mathsf{R} < (\mathcal{G}/\mathsf{R})^{1-1/\gamma}$$

and FHWC $\Rightarrow$ the RHS is < 1 because $(\mathcal{G}/\mathsf{R}) < 1$ (and the RHS is raised to a positive power (because $\gamma > 1$)).

The first panel of Table 4 summarizes: The PF-Unconstrained model has a nondegenerate limiting solution if we impose the RIC and FHWC (these conditions are necessary as well as sufficient). Together the PF-FVAC and the FHWC imply the RIC. If we impose the GIC and the FHWC, both the PF-FVAC and the RIC follow, so GIC+FHWC are also sufficient. But there are circumstances under which the RIC and FHWC can hold while the PF-FVAC fails ('~~PF-FVAC~~'). For example, if $\mathcal{G} = 0$, the problem is a standard 'cake-eating' problem with a nondegenerate solution under the RIC (when the consumer has access to capital markets).

A.3 Properties of the Consumption Function and Limiting MPCs

{sec:MPCiterproofs}

We start by stating some properties of the value functions generated by Problem ($\mathcal{P}_N$).

{lemm:consC2}

Lemma 2. *If* v_t *is strictly negative, strictly increasing, strictly concave,* $\mathbf{C}^3$ *and satisfies* $\lim_{m\to 0} \mathrm{v}_t(m) = -\infty$*, then* c_t *is in* $\mathbf{C}^2$*.*

Proof. Now define an end-of-period value function $\mathfrak{v}_t(a)$ as:

$$\mathfrak{v}_t(a) = \beta \mathbb{E}_t \left[\tilde{\mathcal{G}}_{t+1}^{1-\gamma} \mathrm{v}_{t+1} \left(\tilde{\mathcal{R}}_{t+1} a + \boldsymbol{\xi}_{t+1} \right) \right]. \tag{48}$$

{eq:vfFrackdefn}

Since there is a positive probability that $\boldsymbol{\xi}_{t+1}$ will attain its minimum of zero and since $\tilde{\mathcal{R}}_{t+1} > 0$, it is clear that $\lim_{a\to 0} \mathfrak{v}_t(a) = -\infty$ and $\lim_{a\to 0} \mathfrak{v}'_t(a) = \infty$. So $\mathfrak{v}_t(a)$ is well-defined iff $a > 0$; it is similarly straightforward to show the other properties required for $\mathfrak{v}_t(a)$ to be satisfy the properties of the Proposition. (See Hiraguchi [2003].)

Next define $\underline{\mathrm{v}}_t(m, c)$ as:

$$\underline{\mathrm{v}}_t(m, c) = \mathrm{u}(c) + \mathfrak{v}_t(m - c). \tag{49}$$

Note that for fixed m, $c \mapsto \underline{\mathrm{v}}_t(m, c)$ is $\mathbf{C}^3$ on $(0, m)$ since $\mathfrak{v}_t$ and u are both $\mathbf{C}^3$. Next, observe that our problem's value function defined by Problem ($\mathcal{P}_N$) can be written as:

$$\mathrm{v}_t(m) = \max_c \ \underline{\mathrm{v}}_t(m, c), \tag{50}$$

where the function $\underline{\mathrm{v}}_t$ is well-defined if and only if $0 < c < m$. Furthermore, $\lim_{c\to 0} \underline{\mathrm{v}}_t(m, c) = \lim_{c\to m} \underline{\mathrm{v}}_t(m, c) = -\infty$, $\frac{\partial^2 \underline{\mathrm{v}}_t(m,c)}{\partial c^2} < 0$, $\lim_{c\to 0} \frac{\partial \underline{\mathrm{v}}_t(m,c)}{\partial c} = +\infty$, and $\lim_{c\to m} \frac{\partial \underline{\mathrm{v}}_t(m,c)}{\partial c} = -\infty$. It follows that the $\mathrm{c}_t(m)$ defined by:

$$\mathrm{c}_t(m) = \arg\max_{0<c<m} \underline{\mathrm{v}}_t(m, c) \tag{51}$$

exists and is unique and Problem ($\mathcal{P}_N$) has an interior solution. Moreover, by Berge's Maximum Theorem, c_t will be continuous on S. Next, note that c_t satisfies the first order condition:

$$\mathrm{u}'(\mathrm{c}_t(m)) = \mathfrak{v}'_t(m - \mathrm{c}_t(m)). \tag{52}$$

{eq:consumptionf}

By the Implicit Function Theorem, we then have that c_t is differentiable and:

$$\mathrm{c}_t'(m) = \frac{\mathfrak{v}_t''(\mathrm{a}_t(m))}{\mathrm{u}''(\mathrm{c}_t(m)) + \mathfrak{v}_t''(\mathrm{a}_t(m))}. \tag{53}$$ {eq:derivativeConsF}

Since both u and $\mathfrak{v}_t$ are three times continuously differentiable and c_t is continuous, the RHS of the above equation is continuous and we can conclude that c_t' is continuous and c_t is in $\mathbf{C}^1$.

Finally, $\mathrm{c}_t'(m)$ is differentiable because $\mathfrak{v}_t''$ is $\mathbf{C}^1$, $\mathrm{c}_t(m)$ is $\mathbf{C}^1$ and $\mathrm{u}''(\mathrm{c}_t(m))+\mathfrak{v}_t''(\mathrm{a}_t(m)) < 0$. The second derivative $\mathrm{c}_t''(m)$ will be given by:

$$\mathrm{c}_t''(m) = \frac{a_t'(m)\mathfrak{v}_t'''(a_t)\left[\mathrm{u}''(c_t) + \mathfrak{v}_t''(a_t)\right] - \mathfrak{v}_t''(a_t)\left[c_t'\mathrm{u}'''(c_t) + a_t'\mathfrak{v}_t'''(a_t)\right]}{\left[\mathrm{u}''(c_t) + \mathfrak{v}_t''(a_t)\right]^2}. \tag{54}$$

Since $\mathfrak{v}_t''(\mathrm{a}_t(m))$ is continuous, $\mathrm{c}_t''(m)$ is also continuous.

□

{prop:vfc3}

Proposition 7. *For each t, v_t is strictly negative, strictly increasing, strictly concave, $\mathbf{C}^3$ and satisfies $\lim_{m\to 0}\ \mathrm{v}_t(m) = -\infty$.*

Proof. We will say a function is 'nice' if it satisfies the properties stated by the Proposition. Assume that for some $t+1$, v_{t+1} is nice. Our objective is to show that this implies v_t is also nice; this is sufficient to establish that v_{t-n} is nice by induction for all $n > 0$ because $\mathrm{v}_T(m) = \mathrm{u}(m)$ and u, where $\mathrm{u}(m) = m^{1-\gamma}/(1-\gamma)$, is nice by inspection. By Lemma 2, if v_{t+1} is nice, c_t is in $\mathbf{C}^2$. Next, since both u and $\mathfrak{v}_t$ are strictly concave, both c_t and a_t, where $\mathrm{a}_t(m) = m - \mathrm{c}_t(m)$, are strictly increasing (Recall Equation (53)). This implies that $\mathrm{v}_t(m)$ is nice, since $\mathrm{v}_t(m) = \mathrm{u}(\mathrm{c}_t(m)) + \mathfrak{v}_t(\mathrm{a}_t(m))$. □

Proof for Proposition 3. By Proposition 7, each v_t is strictly negative, strictly increasing, strictly concave, $\mathbf{C}^3$ and satisfies $\lim_{m\to 0}\ \mathrm{v}_t(m) = -\infty$. As such, apply Lemma 2 to conclude the result.

□

Proof of Lemma 3 (Limiting MPCs). *Proof of (i): Minimal MPC*

Fix any t and for any m_t with $m_t > 0$, we can define $\mathrm{e}_t(m_t) = \mathrm{c}_t(m_t)/m_t$ and $\mathrm{a}_t(m_t) = m_t - \mathrm{c}_t(m_t)$. The Euler equation, Equation (4), can be rewritten as:

$$\mathrm{e}_t(m_t)^{-\gamma} = \beta \mathsf{R}\mathbb{E}_t\left[\mathrm{e}_{t+1}(m_{t+1})\left(\frac{\overbrace{\mathsf{R}\mathrm{a}_t(m) + \tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}^{=m_{t+1}\tilde{\mathcal{G}}_{t+1}}}{m_t}\right)\right]^{-\gamma} \tag{55}$$ {eq:eFuncEuler}

where $m_{t+1} = \tilde{\mathcal{R}}_{t+1}(m_t - \mathrm{c}_t(m_t)) + \boldsymbol{\xi}_{t+1}$. The minimal MPC's are obtained by letting where $m_t \to \infty$. Note that $\lim_{m_t\to\infty} m_{t+1} = \infty$ almost surely and thus $\lim_{m_t\to\infty} \mathrm{e}_{t+1}(m_{t+1}) =$

$\underline{\kappa}_{t+1}$ almost surely. Turning to the second term inside the marginal utility on the RHS, we can write:

$$\lim_{m_t\to\infty} \frac{\mathsf{Ra}_t(m_t)+\tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}{m_t} = \lim_{m_t\to\infty}\frac{\mathsf{Ra}_t(m_t)}{m_t} + \lim_{m_t\to\infty}\frac{\tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}{m_t} \tag{56}$$

$$= \mathsf{R}(1-\underline{\kappa}_t)+0, \tag{57}$$

since $\tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}$ is bounded. Thus, we can assert:

$$\lim_{m_t\to\infty}\left(\mathrm{e}_{t+1}(m_{t+1})\left(\frac{\mathsf{Ra}_t(m)+\tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}{m_t}\right)\right)^{-\gamma} = (\mathsf{R}\underline{\kappa}_{t+1}(1-\underline{\kappa}_t))^{-\gamma}, \tag{58}$$

almost surely. Next, the term inside the expectation operator at Equation (55) is bounded above by $\left(\mathsf{R}\underline{\kappa}_{t+1}(1-\overline{\kappa}_t)\right)^{-\gamma}$. Thus, by the Dominated Convergence Theorem, we have:

$$\lim_{m_t\to\infty}\beta\mathsf{R}\mathbb{E}_t\left(\mathrm{e}_{t+1}(m_{t+1})\left(\frac{\mathsf{Ra}_t(m_t)+\tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}{m_t}\right)\right)^{-\gamma} = \beta\mathsf{R}(R\underline{\kappa}_{t+1}(1-\underline{\kappa}_t))^{-\gamma}. \tag{59}$$ {eq:eFuncEulerMPC

Again applying L'Hôpital's rule to the LHS of Equation (55), letting $\lim_{m\to\infty}\mathrm{e}_t(m) = \underline{\kappa}_t$ and equating limits to the RHS, we arrive at:

$$\frac{\mathfrak{Dp}}{\mathsf{R}}\underline{\kappa}_t = (1-\underline{\kappa}_t)\underline{\kappa}_{t+1}$$

The minimal marginal propensity to consume satisfies the following recursive formula:

$$\underline{\kappa}_t^{-1} = 1+\underline{\kappa}_{t+1}^{-1}\frac{\mathfrak{Dp}}{\mathsf{R}}, \tag{60}$$ {eq:MPCminInvAp

which implies $(\{\underline{\kappa}_{T-n}^{-1}\})_{n=0}^{\infty}$ is an increasing convergent sequence. Define:

$$\underline{\kappa}^{-1} :\, = \lim_{n\to\infty}\kappa_{T-n}^{-1} \tag{61}$$

as the limiting (inverse) marginal MPC. If the RIC does *not* hold, then $\lim_{n\to\infty}\underline{\kappa}_{T-n}^{-1} = \infty$ and so the limiting MPC is $\underline{\kappa} = 0$. Otherwise if RIC holds, then $\underline{\kappa} > 0$.

Proof of (ii): Maximal MPC

The Euler Equation (4) can be rewritten as:

$$
\begin{aligned}
\mathrm{e}_t(m_t)^{-\gamma} &= \beta \mathsf{R} \mathbb{E}_t \left[\left(\mathrm{e}_{t+1}(m_{t+1}) \left(\frac{\overbrace{\mathsf{R}\mathrm{a}_t(m) + \tilde{\mathcal{G}}_{t+1}\boldsymbol{\xi}_{t+1}}^{=m_{t+1}\tilde{\mathcal{G}}_{t+1}}}{m_t} \right) \right)^{-\gamma} \right] \\
&= (1-q)\beta \mathsf{R} m_t^{\gamma} \mathbb{E}_t \left[\left(\mathrm{e}_{t+1}(m_{t+1}) m_{t+1} \tilde{\mathcal{G}}_{t+1} \right)^{-\gamma} \mid \boldsymbol{\xi}_{t+1} > 0 \right] \\
&\quad + q\beta \mathsf{R}^{1-\gamma} \mathbb{E}_t \left[\left(\mathrm{e}_{t+1}(\tilde{\mathcal{R}}_{t+1}\mathrm{a}_t(m)) \frac{m_t - \mathrm{c}_t(m)}{m_t} \right)^{-\gamma} \mid \boldsymbol{\xi}_{t+1} = 0 \right]
\end{aligned} \tag{62}
$$

{eq:eFuncEulerMPC}

Now consider the first conditional expectation in the second line of Equation (62). Recall that if $\boldsymbol{\xi}_{t+1} > 0$, then $\boldsymbol{\xi}_{t+1} = \theta_{t+1}/(1-q)$ by Assumption I.1. Since $\lim_{m_t \to 0} \mathrm{a}_t(m_t) = 0$, $\mathbb{E}_t[(\mathrm{e}_{t+1}(m_t')m_t'\mathcal{G}_{t+1})^{-\gamma} \mid \boldsymbol{\xi}_{t+1} > 0]$ is contained in the bounded interval $[(\mathrm{e}_{t+1}(\underline{\theta}/(1-q))\mathcal{G}\underline{\psi}\underline{\theta}/(1-q))^{-\gamma}, (\mathrm{e}_{t+1}(\bar{\theta}/(1-q))\mathcal{G}\bar{\psi}\bar{\theta}/(1-q))^{-\gamma}]$. As such, the first term after the second equality above converges to zero as m_t^{γ} converges to zero.

Turning to the second term after the second equality above, once again apply Dominated Convergence Theorem as noted above at Equation (59). As $m_t \to 0$, the expectation converges to $\overline{\kappa}_{t+1}^{-\gamma}(1-\overline{\kappa}_t)^{-\gamma}$.

Equating the limits on the LHS and RHS of Equation (62), we have $\overline{\kappa}_t^{-\gamma} = \beta q \mathsf{R}^{1-\gamma}\overline{\kappa}_{t+1}^{-\gamma}(1-\overline{\kappa}_t)^{-\gamma}$. Exponentiating by γ on both sides, we can conclude:

$$
\overline{\kappa}_t = q^{-1/\gamma}(\beta \mathsf{R})^{-1/\gamma} \mathsf{R}(1-\overline{\kappa}_t)\overline{\kappa}_{t+1}
$$

and,

$$
\underbrace{q^{1/\gamma} \overbrace{\mathsf{R}^{-1}(\beta \mathsf{R})^{1/\gamma}}^{\frac{\mathbf{Þ}}{\mathsf{R}}}}_{\equiv q^{1/\gamma}\frac{\mathbf{Þ}}{\mathsf{R}}} \overline{\kappa}_t = (1-\overline{\kappa}_t)\overline{\kappa}_{t+1} \tag{63}
$$

{eq:MPSminDef}

The equation above yields a useful recursive formula for the maximal marginal propensity to consume after some algebra:

$$
\begin{aligned}
(q^{1/\gamma}\frac{\mathbf{Þ}}{\mathsf{R}}\overline{\kappa}_t)^{-1} &= (1-\overline{\kappa}_t)^{-1}\overline{\kappa}_{t+1}^{-1} \\
\Rightarrow \overline{\kappa}_t^{-1}(1-\overline{\kappa}_t) &= q^{1/\gamma}\frac{\mathbf{Þ}}{\mathsf{R}}\overline{\kappa}_{t+1}^{-1} \\
\Rightarrow \overline{\kappa}_t^{-1} &= 1 + q^{1/\gamma}\frac{\mathbf{Þ}}{\mathsf{R}}\overline{\kappa}_{t+1}^{-1}
\end{aligned}
$$

As noted in the main text, we need the WRIC (**??**) for this to be a convergent sequence:

$$
0 \leq q^{1/\gamma}\frac{\mathbf{Þ}}{\mathsf{R}} < 1, \tag{64}
$$

{eq:WRICapndx}

Since $\overline{\kappa}_T = 1$, iterating (64) backward to infinity (because we are interested in the limiting consumption function) we obtain:

$$\lim_{n\to\infty} \overline{\kappa}_{T-n} = \overline{\kappa} \equiv 1 - q^{1/\gamma}\frac{\mathbf{\mathfrak{P}}}{\mathsf{R}} \tag{65}$$ {eq:MPCmaxDef}

□

A.4 Existence of Limiting Solutions

{sec:Tcontractionm

We state Boyd's contraction mapping Theorem (Boyd,1990) for completeness.

Theorem 5. *(Boyd's Contraction Mapping) Let $\mathbb{B} : \mathcal{C}_\varphi(S,Y) \to \mathcal{C}_\varphi(S,Y)$. If,*

1. *the operator $\mathbb{B}$ is non-decreasing, i.e. $x \leq y \Rightarrow \mathbb{B}x \leq \mathbb{B}y$,*
2. *we have $\mathbb{B}\mathbf{0}$ in $\mathcal{C}_\varphi(S,Y)$, where $\mathbf{0}$ is the null vector,*
3. *there exists some real $0 < \alpha < 1$ such that for all ζ with $\zeta > 0$, we have:*

$$\mathbb{B}(x + \zeta\varphi) \leq \mathbb{B}x + \zeta\alpha\varphi,$$

then $\mathbb{B}$ defines a contraction with a unique fixed point.

We must show that our operator $\mathbb{T}$ satisfies all of Boyd's conditions.

{claim: MPCMAX

Claim 4. *If WRIC (Assumption L.4) holds, then there exists k such that:*

$$q\beta(\mathsf{R}(1-\overline{\kappa}_k))^{1-\gamma} < 1 \tag{66}$$ {eq: MPCMAXKl

Proof. By straight-forward algebra, we have:

$$\begin{aligned} q\beta(\mathsf{R}(1-\overline{\kappa}))^{1-\gamma} &= q\beta\mathsf{R}^{1-\gamma}\left(q^{1/\gamma}\frac{(\mathsf{R}\beta)^{1/\gamma}}{\mathsf{R}}\right)^{1-\gamma} \\ &= q^{1/\gamma}\frac{(\mathsf{R}\beta)^{1/\gamma}}{\mathsf{R}} < 1 \end{aligned} \tag{67}$$

where the inequality holds by the WRIC (Assumption L.4). Finally, since the expression $q\beta(\mathsf{R}(1-\overline{\kappa}_k))^{1-\gamma}$ is continuous as a function of $\overline{\kappa}_k$, and we have $\bar{\kappa} > 0$ and $\overline{\kappa}_t \to \overline{\kappa}$, by the definition of continuity, there exists k such that Equation (66) holds. □

Proof of Theorem 2. Fix k such that Equation (66) holds. To show $\mathbb{T}^{\underline{b},\bar{b}}f$ satisfies the condition of Theorem 1, we first need to show $\mathbb{T}^{\underline{b},\bar{b}}f$ maps from $\mathcal{C}_\varphi(S,Y)$ to $\mathcal{C}_\varphi(S,Y)$. A preliminary requirement is therefore that $\mathbb{T}^{\underline{b},\bar{b}}\mathrm{f}$ be continuous for any $\varphi-$bounded f, $\mathbb{T}^{\underline{b},\bar{b}}\mathrm{f} \in \mathcal{C}(S,\mathbb{R})$. This is not difficult to show; see Hiraguchi [2003].

Proof of Condition 1

Consider condition (1). For this problem,

$$\{\mathbb{T}\mathrm{x}\}(m_t) \text{ is } \max_{c_t \in [\underline{\kappa} m_t, \overline{\kappa} m_t]} \left\{ \mathrm{u}(c_t) + \beta \mathbb{E}_t \left[\mathcal{G}_{t+1}^{1-\gamma} \mathrm{x}\left(m_{t+1}\right) \right] \right\}$$
$$\{\mathbb{T}\mathrm{y}\}(m_t) \text{ is } \max_{c_t \in [\underline{\kappa} m_t, \overline{\kappa} m_t]} \left\{ \mathrm{u}(c_t) + \beta \mathbb{E}_t \left[\mathcal{G}_{t+1}^{1-\gamma} \mathrm{y}\left(m_{t+1}\right) \right] \right\},$$

so $\mathrm{x}(\bullet) \leq \mathrm{y}(\bullet)$ implies $\{\mathbb{T}\mathrm{x}\}(m_t) \leq \{\mathbb{T}\mathrm{y}\}(m_t)$ by inspection.[49]

Condition (2) requires that $\{\mathbb{T}\mathbf{0}\} \in \mathcal{C}_\varphi(\mathcal{A}, \mathcal{B})$. By definition,

$$\{\mathbb{T}\mathbf{0}\}(m_t) = \max_{c_t \in [\underline{\kappa} m_t, \overline{\kappa} m_t]} \left\{ \left(\frac{c_t^{1-\gamma}}{1-\gamma} \right) + \beta 0 \right\}$$

the solution to which is patently $\mathrm{u}(\overline{\kappa} m_t)$. Thus, condition (2) will hold if $(\overline{\kappa} m_t)^{1-\gamma}$ is φ-bounded, which it is if we use the bounding function

$$\varphi(x) = \zeta + x^{1-\gamma} \tag{68}$$

defined in the main text.

Finally, we turn to condition (3), $\{\mathbb{T}(\mathrm{z} + \zeta\varphi)\}(m_t) \leq \{\mathbb{T}\mathrm{z}\}(m_t) + \zeta\alpha\varphi(m_t)$. The proof will be more compact if we define $\check{\mathrm{c}}$ and $\check{\mathrm{a}}$ as the consumption and assets functions[50] associated with $\mathbb{T}\mathrm{z}$ and $\hat{\mathrm{c}}$ and $\hat{\mathrm{a}}$ as the functions associated with $\mathbb{T}(\mathrm{z} + \zeta\varphi)$; using this notation, condition (3) can be rewritten

$$\mathrm{u}(\hat{\mathrm{c}}) + \beta\{\mathsf{E}(\mathrm{z} + \zeta\varphi)\}(\hat{\mathrm{a}}) \leq \mathrm{u}(\check{\mathrm{c}}) + \beta\{\mathsf{E}\mathrm{z}\}(\check{\mathrm{a}}) + \zeta\alpha\varphi.$$

where $\bar{\mathrm{m}}^{\mathtt{next}}(m) = \tilde{\mathcal{R}}(m - \bar{\mathrm{c}}(m)) + z^{\mathtt{next}}$ and $\hat{\mathrm{m}}^{\mathtt{next}}(m) = \tilde{\mathcal{R}}(m - \hat{\mathrm{c}}(m)) + z^{\mathtt{next}}$. If we now force the consumer facing f as the next period value function to consume the amount optimal for the consumer facing g, the value for the f consumer must be weakly lower. That is,

$$\mathrm{u}(\hat{\mathrm{c}}) + \beta\{\mathsf{E}\mathrm{z}\}(\hat{\mathrm{a}}) \leq \mathrm{u}(\check{\mathrm{c}}) + \beta\{\mathsf{E}\mathrm{z}\}(\check{\mathrm{a}}).$$

Thus, condition (3) will certainly hold under the stronger condition

$$\begin{aligned}
\mathrm{u} \circ \hat{\mathrm{c}} + \beta\mathbb{E}\tilde{\mathcal{G}}\mathrm{g} \circ \hat{m}^{\mathtt{next}} &\leq \mathrm{u} \circ \hat{\mathrm{c}} + \beta\mathbb{E}\mathcal{G}\mathrm{f} \circ \hat{m}^{\mathtt{next}} + \zeta\alpha\varphi \\
\beta\mathbb{E}\tilde{\mathcal{G}}(\mathrm{f} + \zeta\varphi)(\hat{m}^{\mathtt{next}}) &\leq \beta\mathbb{E}\tilde{\mathcal{G}}\mathrm{f} \circ \hat{m}^{\mathtt{next}} + \zeta\alpha\varphi \\
\beta\zeta\mathbb{E}\tilde{\mathcal{G}}\varphi \circ \hat{m}^{\mathtt{next}} &\leq \zeta\alpha\varphi \\
\beta\mathbb{E}\tilde{\mathcal{G}}\varphi \circ \hat{m}^{\mathtt{next}} &\leq \alpha\varphi
\end{aligned}$$

Recall by Claim 4, we have $q\beta(\mathsf{R}(1 - \overline{\kappa}_k))^{1-\gamma} < 1$. As such, use FVAC (Equation (9),

[49] For a fixed m_t, recall that m_{t+1} is just a function of c_t and the stochastic shocks.

[50] Section **??** proves existence of a continuously differentiable consumption function, which implies the existence of a corresponding continuously differentiable assets function.

which says $\beta\mathbb{E}\mathcal{G} < 1$) and fix α such that α satisfies $q\beta(\mathsf{R}(1-\overline{\kappa}_k))^{1-\gamma} < \alpha < 1$ and $\alpha > \beta\mathbb{E}\tilde{\mathcal{G}}$. Next, use $\varphi(m) = \bar{M} + m^{1-\gamma}$ and let $\hat{a}^{\texttt{next}} = m - \hat{\mathrm{c}}(m)$. The condition above will be satisfied if:

$$\beta\mathbb{E}[\tilde{\mathcal{G}}^{\texttt{next}}(\hat{a}^{\texttt{next}}\tilde{\mathcal{R}} + \boldsymbol{\xi})^{1-\gamma}] - \alpha m^{1-\gamma} < \alpha\bar{M}(1 - \alpha^{-1}\beta\mathbb{E}\tilde{\mathcal{G}})$$

where the last line follows because $0 < \alpha < 1$ by assumption.[51]

Using $\varphi(m) = \eta + m^{1-\gamma}$ and defining $\hat{a}_t = \hat{\mathrm{a}}(m_t)$, this condition is

$$\beta\mathbb{E}_t[\mathcal{G}_{t+1}^{1-\gamma}(\hat{a}_t\mathcal{R}_{t+1} + \Xi_{t+1})^{1-\gamma}] - m_t^{1-\gamma} < \eta(1 - \underbrace{\beta\mathbb{E}_t\mathcal{G}_{t+1}^{1-\gamma}}_{=\beth})$$

which by imposing PF-FVAC (equation (9), which says $\beth < 1$) can be rewritten as:

$$\bar{M} > \frac{\beta\mathbb{E}\left[\tilde{\mathcal{G}}(a^{\texttt{next}}\tilde{\mathcal{R}}^{\texttt{next}} + \boldsymbol{\xi}_{t+1})^{1-\gamma}\right] - \alpha m^{1-\gamma}}{\alpha(1 - \alpha^{-1}\beta\mathbb{E}\tilde{\mathcal{G}})}. \tag{69}$$ {eq:KeyCondition}

Since $\bar{M}$ is an arbitrary constant that we can pick, the proof reduces to showing the numerator of (69) is bounded from above:

$$\begin{aligned}
&(1-q)\beta\mathbb{E}_t\left[\tilde{\mathcal{G}}(\hat{a}^{\texttt{next}}\tilde{\mathcal{R}}^{\texttt{next}} + \theta^{\texttt{next}}/(1-q))^{1-\gamma}\right] \\
&\quad + q\beta\mathbb{E}_t\left[\tilde{\mathcal{G}}(\hat{a}^{\texttt{next}}\tilde{\mathcal{R}}^{\texttt{next}})^{1-\gamma}\right] - \alpha m^{1-\gamma} \\
&\leq (1-q)\beta\mathbb{E}_t\left[\tilde{\mathcal{G}}((1-\overline{\kappa}_k)m\tilde{\mathcal{R}}^{\texttt{next}} + \theta^{\texttt{next}}/(1-q))^{1-\gamma}\right] \\
&\quad + q\beta\mathsf{R}^{1-\gamma}((1-\overline{\kappa}_k)m)^{1-\gamma} - \alpha m^{1-\gamma} \\
&= (1-q)\beta\mathbb{E}_t\left[\tilde{\mathcal{G}}((1-\overline{\kappa}_k)m\tilde{\mathcal{R}}^{\texttt{next}} + \theta^{\texttt{next}}/(1-q))^{1-\gamma}\right] \\
&\quad + m^{1-\gamma}\left(\underbrace{q\beta(\mathsf{R}(1-\overline{\kappa}_k))^{1-\gamma}}_{<\alpha \text{ by construction}} - \alpha\right) \\
&< (1-q)\beta\mathbb{E}_t\left[\tilde{\mathcal{G}}(\underline{\theta}/(1-q))^{1-\gamma}\right] = \beta\mathbb{E}\tilde{\mathcal{G}}(1-q)^{\gamma}\underline{\theta}^{1-\gamma}.
\end{aligned} \tag{70}$$

We can thus conclude that equation (69) will certainly hold for any:

$$\bar{M} > \bar{\bar{M}} :\, = \frac{\beta\mathbb{E}\tilde{\mathcal{G}}(1-q)^{\gamma}\underline{\theta}^{1-\gamma}}{\alpha(1 - \alpha^{-1}\beta\mathbb{E}\tilde{\mathcal{G}})} \tag{71}$$

which is a positive finite number under our assumptions.

The proof that $\mathbb{T}$ defines a contraction mapping under the conditions (**??**) and (**??**) is now complete.

[51] The remainder of the proof could be reformulated using the second-to-last line at a small cost to intuition.

A.5 $\mathbb{T}$ and v

In defining our operator $\mathbb{T}$ we made the restriction $\underline{\underline{\kappa}} m_t \leq c_t \leq \overline{\kappa} m_t$. However, in the discussion of the consumption function bounds, we showed only (in (17)) that $\underline{\kappa}_t m_t \leq \mathrm{c}_t(m_t) \leq \overline{\kappa}_t m_t$. (The difference is in the presence or absence of time subscripts on the MPC's.) We have therefore not proven (yet) that the sequence of value functions (**??**) defines a contraction mapping.

Fortunately, the proof of that proposition is identical to the proof above, except that we must replace $\overline{\kappa}$ with $\overline{\kappa}_{T-1}$ and the WRIC must be replaced by a slightly stronger (but still quite weak) condition. The place where these conditions have force is in the step at (**??**). Consideration of the prior two equations reveals that a sufficient stronger condition is

$$\begin{aligned} q\beta(\mathsf{R}(1-\overline{\kappa}_{T-1}))^{1-\gamma} &< 1 \\ (q\beta)^{1/(1-\gamma)}(1-\overline{\kappa}_{T-1}) &> 1 \\ (q\beta)^{1/(1-\gamma)}(1-(1+q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}})^{-1}) &> 1 \end{aligned}$$

where we have used (20) for $\overline{\kappa}_{T-1}$ (and in the second step the reversal of the inequality occurs because we have assumed $\gamma > 1$ so that we are exponentiating both sides by the negative number $1-\gamma$). To see that this is a weak condition, note that for small values of q this expression can be further simplified using $(1+q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}})^{-1} \approx 1 - q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}$ so that it becomes

$$\begin{aligned} (q\beta)^{1/(1-\gamma)} q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} &> 1 \\ (q\beta) q^{(1-\gamma)/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}^{1-\gamma} &< 1 \\ \beta q^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}^{1-\gamma} &< 1. \end{aligned}$$

Calling the weak return patience factor $\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}^{\wp} = \wp^{1/\gamma}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}$ and recalling that the WRIC was $\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}^{\wp} < 1$, the expression on the LHS above is $\beta\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}^{-\gamma}$ times the WRPF. Since we usually assume β not far below 1 and parameter values such that $\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} \approx 1$, this condition is clearly not very different from the WRIC.

The upshot is that under these slightly stronger conditions the value functions for the original problem define a contraction mapping in $\varphi-$bounded space with a unique $\mathrm{v}(m)$. But since $\lim_{n\to\infty} \underline{\underline{\kappa}}_{T-n} = \underline{\underline{\kappa}}$ and $\lim_{n\to\infty} \overline{\kappa}_{T-n} = \overline{\kappa}$, it must be the case that the $\mathrm{v}(m)$ toward which these v_{T-n}'s are converging is the *same* $\mathrm{v}(m)$ that was the endpoint of the contraction defined by our operator $\mathbb{T}$. Thus, under our slightly stronger (but still quite weak) conditions, not only do the value functions defined by (**??**) converge, they converge to the same unique v defined by $\mathbb{T}$.[52]

[52]It seems likely that convergence of the value functions for the original problem could be proven

for any $c \in (0, \overline{\kappa} m]$, where $m_{t_{n(i)}+1} = \tilde{\mathcal{R}}(m - \mathrm{c}_{t_{n(i)}}(m)) + \boldsymbol{\xi}_{t_{n(i)}+1}$ and $\hat{m}^{\mathtt{next}} = \tilde{\mathcal{R}}(m - c) + \boldsymbol{\xi}_{t_{n(i)}+1}$. Allowing $n(i)$ to tend to infinity, the left-hand side converges to:

$$\mathrm{u}(\mathrm{c}(m)) + \beta \mathbb{E}\left[\tilde{\mathcal{G}}^{1-\gamma} \mathrm{v}(m^{\mathtt{next}})\right], \tag{72}$$

where $m^{\mathtt{next}} = \tilde{\mathcal{R}}(m - \mathrm{c}(m)) + \boldsymbol{\xi}$. Moreover, the right-hand side converges to:

$$\mathrm{u}(c) + \beta \mathbb{E}\left[\tilde{\mathcal{G}}^{1-\gamma} \mathrm{v}(\hat{m}^{\mathtt{next}})\right]. \tag{73}$$

Hence, as $n(i)$ tends to infinity, the following inequality is implied:

$$\mathrm{u}(\mathrm{c}(m)) + \beta \mathbb{E}\left[\tilde{\mathcal{G}}^{1-\gamma} \mathrm{v}(m^{\mathtt{next}})\right] \geq \mathrm{u}(c) + \beta \mathbb{E}\left[\tilde{\mathcal{G}}^{1-\gamma} \mathrm{v}(\hat{m}^{\mathtt{next}})\right]. \tag{74}$$

Since the c above was arbitrary, we have:

$$\mathrm{c}(m) \in \underset{c \in (0, \overline{\kappa} m]}{\arg\max} \left\{ \mathrm{u}(c) + \beta \mathbb{E}_t \left[\tilde{\mathcal{G}}_{t+1}^{1-\gamma} \mathrm{v}(\hat{m}^{\mathtt{next}})\right] \right\}. \tag{75}$$

Next, since $\mathrm{c}_{t_{n(i)}} \to \mathrm{c}$ point-wise, and $\mathrm{v}_{t_{n(i)}} \to \mathrm{v}$ point-wise, we have:

$$\mathrm{v}(m) = \lim_{i \to \infty} \mathrm{u}(\mathrm{c}_{t_{n(i)}}(m)) + \beta \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}_{t_{n(i)}+1}(m_{t_{n(i)}+1}) = \mathrm{u}(\mathrm{c}(m)) + \beta \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}(m^{\mathtt{next}}). \tag{76}$$ {eq:convgcvftni}

where $m_{t_n} = \tilde{\mathcal{R}}(m - \mathrm{c}_{t_n}(m))$ and $m^{\mathtt{next}} = \tilde{\mathcal{R}}(m - \mathrm{c}(m))$. The first equality stems form the fact that $\mathrm{v}_{t_n} \to \mathrm{v}$ point-wise, and because point-wise convergence implies point-wise convergence along a sub-sequence. To see why $\lim_{i \to \infty} \mathrm{u}(\mathrm{c}_{t_{n(i)}}(m)) = \mathrm{u}(\mathrm{c}(m))$, note the continuity of u and the convergence of $\mathrm{c}_{t_{n(i)}}$ to c point-wise. To see why $\lim_{i \to \infty} \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}_{t_{n(i)}+1}(m_{t_{n(i)}+1}) = \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}(m^{\mathtt{next}})$, note that $\mathrm{v}_{t_{n(i)}+1}$ converges in the φ-norm, hence converges uniformly over compact sets in S and apply Fact 3 from the standard mathematical results presented in Appendix D. This completes the proof of part (ii) of the Theorem.

Proof of (iii)

The limits at Equation (76) immediately imply:

$$\mathrm{v}(m) = \lim_{n \to \infty} \mathrm{u}(\mathrm{c}_{t_n}(m)) + \beta \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}_{t_n+1}(m_{t_n+1}) = \mathrm{u}(\mathrm{c}(m)) + \beta \mathbb{E} \tilde{\mathcal{G}} \mathrm{v}(m^{\mathtt{next}}), \tag{77}$$

since a real valued sequence can have at most one limit. Finally, applying Fact 8 from Appendix D, we get $\mathrm{c}_{t_n}(m)) \to \mathrm{c}(m)$, thus establishing that c_{t_n} converges point-wise to c.

□

even if only the WRIC were imposed; but that proof is not an essential part of the enterprise of this paper and is therefore left for future work.

A.6 The Liquidity Constrained Solution as a Limit

{sec:LiqConstrAsLi

Proof of Proposition 4. Formally, suppose we change the description of the problem by making the following two assumptions:

$$\begin{aligned} q &= 0 \\ c_t &\leq m_t, \end{aligned}$$

and we designate the solution to this consumer's problem $\grave{\mathrm{c}}_t(m)$. We will henceforth refer to this as the problem of the 'restrained' consumer (and, to avoid a common confusion, we will refer to the consumer as 'constrained' only in circumstances when the constraint is actually binding).

Redesignate the consumption function that emerges from our original problem for a given fixed q as $\mathrm{c}_t(m;q)$ where we separate the arguments by a semicolon to distinguish between m, which is a state variable, and q, which is not. The proposition we wish to demonstrate is

$$\lim_{q \downarrow 0} \mathrm{c}_t(m;q) = \grave{\mathrm{c}}_t(m). \tag{78}$$

{eq:RestrEqUnrestr

We will first examine the problem in period $T-1$, then argue that the desired result propagates to earlier periods. For simplicity, suppose that the interest, growth, and time-preference factors are $\beta = \mathsf{R} = \mathcal{G} = 1$, and there are no permanent shocks, $\psi = 1$; the results below are easily generalized to the full-fledged version of the problem.

The solution to the restrained consumer's optimization problem can be obtained as follows. Assuming that the consumer's behavior in period T is given by $\mathrm{c}_T(m)$ (in practice, this will be $\mathrm{c}_T(m) = m$), consider the unrestrained optimization problem

$$\grave{\mathrm{a}}^*_{T-1}(m) = \arg\max_a \left\{ \mathrm{u}(m-a) + \int_{\underline{\theta}}^{\bar{\theta}} \mathrm{v}_T(a+\theta) d\mathcal{F}_\theta \right\}. \tag{79}$$

{eq:vUnconstr}

As usual, the envelope theorem tells us that $\mathrm{v}'_T(m) = \mathrm{u}'(\mathrm{c}_T(m))$ so the expected marginal value of ending period $T-1$ with assets a can be defined as

$$\grave{\mathfrak{v}}'_{T-1}(a) \equiv \int_{\underline{\theta}}^{\bar{\theta}} \mathrm{u}'(\mathrm{c}_T(a+\theta)) d\mathcal{F}_\theta,$$

and the solution to (79) will satisfy

$$\mathrm{u}'(m-a) = \grave{\mathfrak{v}}'_{T-1}(a). \tag{80}$$

{eq:uPConstr}

$\grave{\mathrm{a}}^*_{T-1}(m)$ therefore answers the question "With what level of assets would the restrained consumer like to end period $T-1$ if the constraint $c_{T-1} \leq m_{T-1}$ did not exist?" (Note that the restrained consumer's income process remains different from the process for the unrestrained consumer so long as $q > 0$.) The restrained consumer's actual asset position will be

$$\grave{\mathrm{a}}_{T-1}(m) = \max[0, \grave{\mathrm{a}}^*_{T-1}(m)],$$

reflecting the inability of the restrained consumer to spend more than current resources, and note (as pointed out by Deaton [1991]) that

$$m^1_{\#} = \left(\grave{\mathfrak{v}}'_{T-1}(0)\right)^{-1/\gamma}$$

is the cusp value of m at which the constraint makes the transition between binding and non-binding in period $T-1$.

Analogously to (80), defining

$$\mathfrak{v}'_{T-1}(a;q) \equiv \left[qa^{-\gamma} + (1-q)\int_{\underline{\theta}}^{\bar{\theta}} (\mathrm{c}_T(a+\theta/(1-q)))^{-\gamma} d\mathcal{F}_\theta\right], \tag{81}$$ {eq:vFrakPrime}

the Euler equation for the original consumer's problem implies

$$(m-a)^{-\gamma} = \mathfrak{v}'_{T-1}(a;q) \tag{82}$$ {eq:uPUnconstr}

with solution $\mathrm{a}^*_{T-1}(m;q)$. Now note that for any fixed $a > 0$, $\lim_{q\downarrow 0} \mathfrak{v}'_{T-1}(a;q) = \grave{\mathfrak{v}}'_{T-1}(a)$. Since the LHS of (80) and (82) are identical, this means that $\lim_{q\downarrow 0} \mathrm{a}^*_{T-1}(m;q) = \grave{\mathrm{a}}^*_{T-1}(m)$. That is, for any fixed value of $m > m^1_{\#}$ such that the consumer subject to the restraint would voluntarily choose to end the period with positive assets, the level of end-of-period assets for the unrestrained consumer approaches the level for the restrained consumer as $q \downarrow 0$. With the same a and the same m, the consumers must have the same c, so the consumption functions are identical in the limit.

Now consider values $m \leq m^1_{\#}$ for which the restrained consumer is constrained. It is obvious that the baseline consumer will never choose $a \leq 0$ because the first term in (81) is $\lim_{a\downarrow 0} qa^{-\gamma} = \infty$, while $\lim_{a\downarrow 0} (m-a)^{-\gamma}$ is finite (the marginal value of end-of-period assets approaches infinity as assets approach zero, but the marginal utility of consumption has a finite limit for $m > 0$). The subtler question is whether it is possible to rule out strictly positive a for the unrestrained consumer.

The answer is yes. Suppose, for some $m < m^1_{\#}$, that the unrestrained consumer is considering ending the period with any positive amount of assets $a = \delta > 0$. For any such δ we have that $\lim_{q\downarrow 0} \mathfrak{v}'_{T-1}(a;q) = \grave{\mathfrak{v}}'_{T-1}(a)$. But by assumption we are considering a set of circumstances in which $\grave{\mathrm{a}}^*_{T-1}(m) < 0$, and we showed earlier that $\lim_{q\downarrow 0} \mathrm{a}^*_{T-1}(m;q) = \grave{\mathrm{a}}^*_{T-1}(m)$. So, having assumed $a = \delta > 0$, we have proven that the consumer would optimally choose $a < 0$, which is a contradiction. A similar argument holds for $m = m^1_{\#}$.

These arguments demonstrate that for any $m > 0$, $\lim_{q\downarrow 0} \mathrm{c}_{T-1}(m;q) = \grave{\mathrm{c}}_{T-1}(m)$ which is the period $T-1$ version of (78). But given equality of the period $T-1$ consumption functions, backwards recursion of the same arguments demonstrates that the limiting consumption functions in previous periods are also identical to the constrained function.

Note finally that another intuitive confirmation of the equivalence between the two problems is that our formula (65) for the maximal marginal propensity to consume satisfies

$$\lim_{q\downarrow 0} \overline{\kappa} \quad = 1,$$

which makes sense because the marginal propensity to consume for a constrained restrained consumer is 1 by our definitions of 'constrained' and 'restrained.'

□

B Proofs for Individual Stability (Section 3)

{sec:ApndxMTarget

Proof for Proposition 5. *Proof.* For consumption growth, as $m \downarrow 0$ we have

$$\begin{aligned}
\lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\mathrm{c}(m_{t+1})}{\mathrm{c}(m_t)} \right) \mathcal{G}_{t+1} \right] &> \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t) + \boldsymbol{\xi}_{t+1})}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&= q \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&\quad + (1-q) \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t) + \theta_{t+1}/(1-q))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&> (1-q) \lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\theta_{t+1}/(1-q))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right] \\
&= \infty
\end{aligned}$$

where the second-to-last line follows because $\lim_{m_t \downarrow 0} \mathbb{E}_t \left[\left(\frac{\underline{\mathrm{c}}(\mathcal{R}_{t+1}\mathrm{a}(m_t))}{\overline{\kappa} m_t} \right) \mathcal{G}_{t+1} \right]$ is positive, and the last line follows because the minimum possible realization of θ_{t+1} is $\underline{\theta} > 0$ so the minimum possible value of expected next-period consumption is positive.[53]

Next we establish the limit of the expected consumption growth factor as $m_t \uparrow \infty$:

$$\lim_{m_t \uparrow \infty} \mathbb{E}_t[\boldsymbol{c}_{t+1}/\boldsymbol{c}_t] = \lim_{m_t \uparrow \infty} \mathbb{E}_t[\mathcal{G}_{t+1} c_{t+1}/c_t].$$

But

$$\mathbb{E}_t[\mathcal{G}_{t+1}\underline{c}_{t+1}/\bar{c}_t] \leq \mathbb{E}_t[\mathcal{G}_{t+1} c_{t+1}/c_t] \leq \mathbb{E}_t[\mathcal{G}_{t+1}\bar{c}_{t+1}/\underline{c}_t]$$

and

$$\lim_{m_t \uparrow \infty} \mathcal{G}_{t+1}\underline{\mathrm{c}}(m_{t+1})/\bar{\mathrm{c}}(m_t) = \lim_{m_t \uparrow \infty} \mathcal{G}_{t+1}\bar{\mathrm{c}}(m_{t+1})/\underline{\mathrm{c}}(m_t) = \lim_{m_t \uparrow \infty} \mathcal{G}_{t+1} m_{t+1}/m_t,$$

[53]None of the arguments in either of the two prior sections depended on the assumption that the consumption functions had converged. With more cumbersome notation, each derivation could have been replaced by the corresponding finite-horizon versions. This strongly suggests that it should be possible to extend the circumstances under which the problem can be shown to define a contraction mapping to the union of the parameter values under which {RIC,FHWC} hold and {FVAC,WRIC} hold. That extension is not necessary for our purposes here, so we leave it for future work.

while (for convenience defining $\mathrm{a}(m_t) = m_t - \mathrm{c}(m_t)$),

$$\begin{aligned}\lim_{m_t\uparrow\infty} \mathcal{G}_{t+1} m_{t+1}/m_t &= \lim_{m_t\uparrow\infty} \left(\frac{\mathsf{R}\mathrm{a}(m_t) + \mathcal{G}_{t+1}\boldsymbol{\xi}_{t+1}}{m_t}\right) \\ &= (\mathsf{R}\beta)^{1/\gamma} = \mathbf{\Phi}\end{aligned} \tag{83}$$

{eq:xtp1toinfty}

because $\lim_{m_t\uparrow\infty} \mathrm{a}'(m) = \frac{\mathbf{\Phi}}{\mathsf{R}}$[54] and $\mathcal{G}_{t+1}\boldsymbol{\xi}_{t+1}/m_t \leq (\mathcal{G}\bar{\psi}\bar{\theta}/(1-q))/m_t$ which goes to zero as m_t goes to infinity.

Hence we have

$$\mathbf{\Phi} \leq \lim_{m_t\uparrow\infty} \mathbb{E}_t[\boldsymbol{c}_{t+1}/\boldsymbol{c}_t] \leq \mathbf{\Phi}$$

so as cash goes to infinity, consumption growth approaches its value $\mathbf{\Phi}$ in the perfect foresight model. □

□

This appendix proves Theorems 6-7 and:

{lemma:orderingP

Lemma 3. *If $\check{m}$ and $\hat{m}$ both exist, then $\check{m} \leq \hat{m}$.*

B.1 Proof of Theorem 6

Theorem 6. *(Individual Market-Resources-to-Permanent-Income Ratio Target). Consider the problem defined in Section 2.1. If weak return impatience (Assumption L.4), finite value of autarky (Assumption L.1) and strong growth impatience (Assumption S.2) hold, then there exists a unique market resources to permanent income ratio, $\hat{m}$, with $\hat{m} > 0$, such that:*

$$\mathbb{E}_t[m_{t+1}/m_t] = 1 \text{ if } m_t = \hat{m}. \tag{84}$$

Moreover, $\hat{m}$ is a point of 'stability' in the sense that:

$$\begin{aligned}\forall m_t \in (0, \hat{m}),\ \mathbb{E}_t[m_{t+1}] > m_t \\ \forall m_t \in (\hat{m}, \infty),\ \mathbb{E}_t[m_{t+1}] < m_t.\end{aligned} \tag{85}$$

The elements of the proof of Theorem 6 are:

- Existence and continuity of $\mathbb{E}_t[m_{t+1}/m_t]$
- Existence of a point where $\mathbb{E}_t[m_{t+1}/m_t] = 1$
- $\mathbb{E}_t[m_{t+1}] - m_t$ is monotonically decreasing

[54] $\lim_{m_t\uparrow\infty} \mathrm{a}(m_t)/m_t = 1 - \lim_{m_t\uparrow\infty} \mathrm{c}(m_t)/m_t = 1 - \lim_{m_t\uparrow\infty} \mathrm{c}'(m_t) = \frac{\mathbf{\Phi}}{\mathsf{R}}$.

B.2 Existence and Continuity of $\mathbb{E}_t[m_{t+1}/m_t]$

{subsubsec:RatExit

The consumption function exists because we have imposed sufficient conditions (the WRIC and FVAC; Theorem **??**).

Section **??** shows that for all t, $a_{t-1} = m_{t-1} - c_{t-1} > 0$. Since $m_t = a_{t-1}\tilde{\mathcal{R}}_t + \boldsymbol{\xi}_t$, even if $\boldsymbol{\xi}_t$ takes on its minimum value of 0, $a_{t-1}\tilde{\mathcal{R}}_t > 0$, since both a_{t-1} and $\tilde{\mathcal{R}}_t$ are strictly positive. With m_t and m_{t+1} both strictly positive, the ratio $\mathbb{E}_t[m_{t+1}/m_t]$ inherits continuity (and, for that matter, continuous differentiability) from the consumption function.

B.3 Existence of a point where $\mathbb{E}_t[m_{t+1}/m_t] = 1$.

This follows from:

1. Existence and continuity of $\mathbb{E}_t[m_{t+1}/m_t]$ (just proven)
2. Existence a point where $\mathbb{E}_t[m_{t+1}/m_t] < 1$
3. Existence a point where $\mathbb{E}_t[m_{t+1}/m_t] > 1$
4. The Intermediate Value Theorem

B.3.1 Existence of m where $\mathbb{E}_t[m_{t+1}/m_t] < 1$

If RIC holds. Logic exactly parallel to that of Section **??** leading to equation (83), but dropping the $\mathcal{G}_{t+1}$ from the RHS, establishes that

$$\begin{aligned}\lim_{m_t \uparrow \infty} \mathbb{E}_t[m_{t+1}/m_t] &= \lim_{m_t \uparrow \infty} \mathbb{E}_t\left[\frac{\mathcal{R}_{t+1}(m_t - \mathrm{c}(m_t)) + \boldsymbol{\xi}_{t+1}}{m_t}\right] \\ &= \mathbb{E}_t[(\mathsf{R}/\mathcal{G}_{t+1})\frac{\text{Þ}}{\mathsf{R}}] \\ &= \mathbb{E}_t[\text{Þ}/\mathcal{G}_{t+1}] \qquad (86) \quad \{\text{eq:emgro}\} \\ &< 1\end{aligned}$$

where the inequality reflects imposition of the GIC-Mod (**??**).

If RIC fails. When the RIC fails, the fact that $\lim_{m \uparrow \infty} \mathrm{c}'(m) = 0$ (see equation (19)) means that the limit of the RHS of (86) as $m \uparrow \infty$ is $\tilde{\bar{\mathcal{R}}} = \mathbb{E}_t[\tilde{\mathcal{R}}_{t+1}]$. In the next step of this proof, we will prove that the combination GIC-Mod and ~~RIC~~ implies $\tilde{\bar{\mathcal{R}}} < 1$.

So we have $\lim_{m \uparrow \infty} \mathbb{E}_t[m_{t+1}/m_t] < 1$ whether the RIC holds or fails.

B.3.2 Existence of m > 1 where $\mathbb{E}_t[m_{t+1}/m_t] > 1$

Paralleling the logic for c in Section 3.1: the ratio of $\mathbb{E}_t[m_{t+1}]$ to m_t is unbounded above as $m_t \downarrow 0$ because $\lim_{m_t \downarrow 0} \mathbb{E}_t[m_{t+1}] > 0$.

Intermediate Value Theorem. If $\mathbb{E}_t[m_{t+1}/m_t]$ is continuous, and takes on values above and below 1, there must be at least one point at which it is equal to one.

B.3.3 $\mathbb{E}_t[m_{t+1}] - m_t$ *is monotonically decreasing.*

Now define $\boldsymbol{\zeta}(m_t) \equiv \mathbb{E}_t[m_{t+1}] - m_t$ and note that

$$\begin{aligned} \boldsymbol{\zeta}(m_t) < 0 &\leftrightarrow \mathbb{E}_t[m_{t+1}/m_t] < 1 \\ \boldsymbol{\zeta}(m_t) = 0 &\leftrightarrow \mathbb{E}_t[m_{t+1}/m_t] = 1 \\ \boldsymbol{\zeta}(m_t) > 0 &\leftrightarrow \mathbb{E}_t[m_{t+1}/m_t] > 1, \end{aligned} \tag{87}$$ {eq:difNrmioEquiv}

so that $\boldsymbol{\zeta}(\hat{m}) = 0$. Our goal is to prove that $\boldsymbol{\zeta}(\bullet)$ is strictly decreasing on $(0, \infty)$ using the fact that

$$\begin{aligned} \boldsymbol{\zeta}'(m_t) \equiv \left(\frac{d}{dm_t}\right)\boldsymbol{\zeta}(m_t) &= \mathbb{E}_t\left[\left(\frac{d}{dm_t}\right)\left(\tilde{\mathcal{R}}_{t+1}(m_t - \mathrm{c}(m_t)) + \boldsymbol{\xi}_{t+1} - m_t\right)\right] \\ &= \bar{\tilde{\mathcal{R}}}\left(1 - \mathrm{c}'(m_t)\right) - 1. \end{aligned} \tag{88}$$ {eq:difFuncmNrmD

Now, we show that (given our other assumptions) $\boldsymbol{\zeta}'(m)$ is decreasing (but for different reasons) whether the RIC holds or fails.

If RIC holds. Equation (**??**) indicates that if the RIC holds, then $\underline{\kappa} > 0$. We show at the bottom of Section **??** that if the RIC holds then $0 < \underline{\kappa} < \mathrm{c}'(m_t) < 1$ so that

$$\begin{aligned} \bar{\tilde{\mathcal{R}}}\left(1 - \mathrm{c}'(m_t)\right) - 1 &< \bar{\tilde{\mathcal{R}}}(1 - \underbrace{(1 - \frac{\mathbf{\mathfrak{p}}}{\mathsf{R}})}_{\underline{\kappa}}) - 1 \\ &= \bar{\tilde{\mathcal{R}}}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} - 1 \\ &= \mathbb{E}_t\left[\frac{\mathsf{R}}{\mathcal{G}\psi}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}\right] - 1 \\ &= \underbrace{\mathbb{E}_t\left[\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}\psi}\right]}_{=\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}\mathbb{E}[\psi^{-1}]} - 1 \end{aligned}$$

which is negative because the GIC-Mod says $\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}\mathbb{E}[\psi^{-1}] < 1$.

If RIC fails. Under ~~RIC~~, recall that $\lim_{m\uparrow\infty} \mathrm{c}'(m) = 0$. Concavity of the consumption function means that c' is a decreasing function, so everywhere

$$\bar{\tilde{\mathcal{R}}}\left(1 - \mathrm{c}'(m_t)\right) < \bar{\tilde{\mathcal{R}}}$$

which means that $\boldsymbol{\zeta}'(m_t)$ from (88) is guaranteed to be negative if

$$\bar{\tilde{\mathcal{R}}} \equiv \mathbb{E}_t\left[\frac{\mathsf{R}}{\mathcal{G}\psi}\right] < 1. \tag{89}$$ {eq:RbarBelowOne}

But the combination of the GIC-Mod holding and the RIC failing can be written:

$$\overbrace{\mathbb{E}_t\left[\frac{\mathbb{Þ}}{\mathcal{G}\psi}\right]}^{\frac{\mathbb{Þ}}{\mathcal{G}}\mathbb{E}[\psi^{-1}]} < 1 < \overbrace{\frac{\mathbb{Þ}}{\mathsf{R}}}^{\frac{\mathbb{Þ}}{\mathsf{R}}},$$

and multiplying all three elements by $\mathsf{R}/\mathbb{Þ}$ gives

$$\mathbb{E}_t\left[\frac{\mathsf{R}}{\mathcal{G}\psi}\right] < \mathsf{R}/\mathbb{Þ} < 1$$

which satisfies our requirement in (89).

B.4 Proof of Theorem 7

Theorem 7. *(Individual Balanced-Growth 'Pseudo Steady State'). Consider the problem defined in Section 2.1. If weak return impatience (Assumption L.4), finite value of autarky (Assumption L.1) and growth impatience (Assumption S.1) hold, then there exists a unique pseudo-steady-state market resources to permanent income ratio $\check{m} > 0$ such that:*

$$\mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] = 1 \text{ if } m_t = \check{m}. \tag{90}$$

Moreover, $\check{m}$ is a point of stability in the sense that:

$$\begin{aligned} \forall m_t \in (0,\check{m}),\ \mathbb{E}_t[\boldsymbol{m}_{t+1}]/\boldsymbol{m}_t > \mathcal{G} \\ \forall m_t \in (\check{m},\infty),\ \mathbb{E}_t[\boldsymbol{m}_{t+1}]/\boldsymbol{m}_t < \mathcal{G}. \end{aligned} \tag{91}$$

The elements of the proof are:

- Existence and continuity of $\mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t]$
- Existence of a point where $\mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] = 1$
- $\mathbb{E}_t[\psi_{t+1}m_{t+1} - m_t]$ is monotonically decreasing

B.4.1 Existence and Continuity of the Ratio

Since by assumption $0 < \underline{\psi} \leq \psi_{t+1} \leq \bar{\psi} < \infty$, our proof in B.2 that demonstrated existence and continuity of $\mathbb{E}_t[m_{t+1}/m_t]$ implies existence and continuity of $\mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t]$.

B.4.2 Existence of a stable point

Since by assumption $0 < \underline{\psi} \leq \psi_{t+1} \leq \bar{\psi} < \infty$, our proof in Subsection B.2 that the ratio of $\mathbb{E}_t[m_{t+1}]$ to m_t is unbounded as $m_t \downarrow 0$ implies that the ratio $\mathbb{E}_t[\psi_{t+1}m_{t+1}]$ to m_t is unbounded as $m_t \downarrow 0$.

The limit of the expected ratio as m_t goes to infinity is most easily calculated by modifying the steps for the prior theorem explicitly:

$$\begin{aligned}
\lim_{m_t\uparrow\infty} \mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] &= \lim_{m_t\uparrow\infty} \mathbb{E}_t\left[\frac{\mathcal{G}_{t+1}\left((\mathsf{R}/\mathcal{G}_{t+1})\mathrm{a}(m_t)+\boldsymbol{\xi}_{t+1}\right)/\mathcal{G}}{m_t}\right] \\
&= \lim_{m_t\uparrow\infty} \mathbb{E}_t\left[\frac{(\mathsf{R}/\mathcal{G})\mathrm{a}(m_t)+\psi_{t+1}\boldsymbol{\xi}_{t+1}}{m_t}\right] \\
&= \lim_{m_t\uparrow\infty} \left[\frac{(\mathsf{R}/\mathcal{G})\mathrm{a}(m_t)+1}{m_t}\right] \\
&= (\mathsf{R}/\mathcal{G})\frac{\mathbb{Þ}}{\mathsf{R}} \qquad (92) \quad \{\text{eq:emgro2}\} \\
&= \frac{\mathbb{Þ}}{\mathcal{G}} \\
&< 1
\end{aligned}$$

where the last two lines are merely a restatement of the GIC (**??**).

The Intermediate Value Theorem says that if $\mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t]$ is continuous, and takes on values above and below 1, there must be at least one point at which it is equal to one.

B.4.3 $\mathbb{E}_t[\psi_{t+1}m_{t+1}] - m_t$ *is monotonically decreasing.*

Define $\boldsymbol{\zeta}(m_t) \equiv \mathbb{E}_t[\psi_{t+1}m_{t+1}] - m_t$ and note that

$$\begin{aligned}
\boldsymbol{\zeta}(m_t) < 0 &\leftrightarrow \mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] < 1 \\
\boldsymbol{\zeta}(m_t) = 0 &\leftrightarrow \mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] = 1 \qquad (93) \quad \{\text{eq:difLvlEquiv}\} \\
\boldsymbol{\zeta}(m_t) > 0 &\leftrightarrow \mathbb{E}_t[\psi_{t+1}m_{t+1}/m_t] > 1,
\end{aligned}$$

so that $\boldsymbol{\zeta}(\hat{m}) = 0$. Our goal is to prove that $\boldsymbol{\zeta}(\bullet)$ is strictly decreasing on $(0,\infty)$ using the fact that

$$\begin{aligned}
\boldsymbol{\zeta}'(m_t) \equiv \left(\frac{d}{dm_t}\right)\boldsymbol{\zeta}(m_t) &= \mathbb{E}_t\left[\left(\frac{d}{dm_t}\right)\left(\tilde{\mathcal{R}}(m_t - \mathrm{c}(m_t)) + \psi_{t+1}\boldsymbol{\xi}_{t+1} - m_t\right)\right] \qquad (94) \quad \{\text{eq:difFuncDecreas}\} \\
&= (\mathsf{R}/\mathcal{G})\left(1 - \mathrm{c}'(m_t)\right) - 1.
\end{aligned}$$

Now, we show that (given our other assumptions) $\boldsymbol{\zeta}'(m)$ is decreasing (but for different reasons) whether the RIC holds or fails (~~RIC~~).

If RIC holds. Equation (**??**) indicates that if the RIC holds, then $\underline{\kappa} > 0$. We show

at the bottom of Section **??** that if the RIC holds then $0 < \underline{\kappa} < \mathrm{c}'(m_t) < 1$ so that

$$\tilde{\mathcal{R}}\,(1-\mathrm{c}'(m_t)) - 1 < \tilde{\mathcal{R}}(1 - \underbrace{(1 - \frac{\mathbf{Þ}}{\mathsf{R}})}_{\underline{\kappa}}) - 1$$
$$= (\mathsf{R}/\mathcal{G})\frac{\mathbf{Þ}}{\mathsf{R}} - 1$$

which is negative because the GIC says $\frac{\mathbf{Þ}}{\mathcal{G}} < 1$.

C Perfect Foresight Liquidity Constrained Solution

{sec:ApndxLiqCons

We briefly interpret FVAC before turning to how all the conditions relate under uncertainty. Analogously to (95), the value for a consumer who spent exactly their permanent income every period would reflect the product of the expectation of the (independent) future shocks to permanent income:

$$= \mathrm{u}(\boldsymbol{p}_t)\left(\frac{1-(\beta\mathcal{G}^{1-\gamma}\mathbb{E}[\psi^{1-\gamma}])^{T-t+1}}{1-\beta\mathcal{G}^{1-\gamma}\mathbb{E}[\psi^{1-\gamma}]}\right),$$

The function $\mathbf{v}_t$ will be finite as T approaches ∞ if the FVAC holds. In the case without uncertainty, Because $\mathrm{u}(xy) = x^{1-\gamma}\mathrm{u}(y)$, the value the consumer would achieve is:

$$\begin{aligned}\mathbf{v}_t^{\text{autarky}} &= \mathrm{u}(\boldsymbol{p}_t) + \beta\mathrm{u}(\boldsymbol{p}_t\mathcal{G}) + \beta^2\mathrm{u}(\boldsymbol{p}_t\mathcal{G}^2) + \dots \\ &= \mathrm{u}(\boldsymbol{p}_t)\left(\frac{1-(\beta\mathcal{G}^{1-\gamma})^{T-t+1}}{1-\beta\mathcal{G}^{1-\gamma}}\right)\end{aligned}$$

which (for $\mathcal{G} > 0$) asymptotes to a finite number as n, with $n = T - t$, approaches $+\infty$.

C.1 Perfect Foresight Unconstrained Solution

{subsec:ApndxUCP

The first result relates to the perfect foresight case without liquidity constraints.

Proof. (Proof of Proposition 1) Consider a sequence of consumption $\{\boldsymbol{c}_{T-n}\}_{n=t}^{T}$ and a sequence of income $\{\boldsymbol{p}_{T-n}\}_{n=t}^{T}$ and let $\mathrm{PDV}_t^T(\boldsymbol{c})$ and $\mathrm{PDV}_t^T(\boldsymbol{p})$ denote the present discounted value of the consumption sequence and permanent income sequence respectively. The dynamic budget constraint, strictly positive marginal utility, and the can't-die-in-debt condition, Equation (1), imply an exactly-holding intertemporal budget constraint (IBC):

$$\mathrm{PDV}_t^T(\boldsymbol{c}) = \overbrace{\boldsymbol{m}_t - \boldsymbol{p}_t}^{\boldsymbol{b}_t} + \overbrace{\mathrm{PDV}_t^T(\boldsymbol{p})}^{\boldsymbol{h}_t}, \tag{95}$$

{eq:IBCFinite}

where $\boldsymbol{b}$ is beginning-of-period 'market' balances; with $\tilde{\mathcal{R}} \colon = \mathsf{R}/\mathcal{G}$ 'human wealth' can be written as:

$$\begin{aligned} \boldsymbol{h}_t &= \boldsymbol{p}_t + \tilde{\mathcal{R}}^{-1}\boldsymbol{p}_t + \tilde{\mathcal{R}}^{-2}\boldsymbol{p}_t + \cdots + \tilde{\mathcal{R}}^{t-T}\boldsymbol{p}_t \\ &= \underbrace{\left(\frac{1-\tilde{\mathcal{R}}^{-(T-t+1)}}{1-\tilde{\mathcal{R}}^{-1}}\right)}_{\equiv h_t}\boldsymbol{p}_t. \end{aligned} \tag{96}$$ {eq:HDef}

Let h denote the limiting value of normalized human wealth as the planning horizon recedes, we have $h \colon = \lim_{n\to\infty} h_{t_n}$.

Next, since consumption is growing by $\Bbbp$ but discounted by R:

$$\text{PDV}_t^T(\boldsymbol{c}) = \left(\frac{1-\frac{\Bbbp}{\mathsf{R}}^{T-t+1}}{1-\frac{\Bbbp}{\mathsf{R}}}\right)\boldsymbol{c}_t$$

from which the IBC (95) implies

$$\boldsymbol{c}_t = \overbrace{\left(\frac{1-\frac{\Bbbp}{\mathsf{R}}}{1-\frac{\Bbbp}{\mathsf{R}}^{T-t+1}}\right)}^{\equiv \underline{\kappa}_t}(\boldsymbol{b}_t + \boldsymbol{h}_t) \tag{97}$$ {eq:WDef}

defining a normalized finite-horizon perfect foresight consumption function:

$$\bar{\mathrm{c}}_{T-n}(m_{T-n}) = (\overbrace{m_{T-n}-1}^{\equiv b_{T-n}} + h_{T-n})\underline{\kappa}_{t-n}$$

where $\underline{\kappa}_t$ is the marginal propensity to consume (MPC). (The overbar signifies that $\bar{\mathrm{c}}$ will be an upper bound as we modify the problem to incorporate constraints and uncertainty; analogously, $\underline{\kappa}$ is the MPC's lower bound).

The horizon-exponentiated term in the denominator of (97) is why, for $\underline{\kappa}$ to be strictly positive as n goes to infinity, we must impose the RIC. The RIC thus implies that the consumer cannot be so pathologically patient as to wish, in the limit as the horizon approaches infinity, to spend nothing today out of an increase in current wealth (the RIC rules out the degenerate limiting solution $\bar{\mathrm{c}}(m) = 0$).

Given that the RIC holds, and (as before) defining limiting objects by the absence of a time subscript, the limiting upper bound consumption function will be

$$\bar{\mathrm{c}}(m) = (m + h - 1)\underline{\kappa}, \tag{98}$$ {eq:cFuncPFUnc}

and so in order to rule out the degenerate limiting solution $\bar{\mathrm{c}}(m) = \infty$ we need h to be finite; that is, we must impose the Finite Human Wealth Condition (FHWC), eq. (**??**). □

Under perfect foresight in the presence of a liquidity constraint requiring $b \geq 0$, this

appendix taxonomizes the varieties of the limiting consumption function $\grave{\mathrm{c}}(m)$ that arise under various parametric conditions.

Results are summarized in table 5.

C.2 If GIC Fails

A consumer is 'growth patient' if the perfect foresight growth impatience condition fails ($\cancel{\text{GIC}}$, $1 < \mathbf{\Phi}/\mathcal{G}$). Under $\cancel{\text{GIC}}$ the constraint does not bind at the lowest feasible value of $m_t = 1$ because $1 < (\mathsf{R}\beta)^{1/\gamma}/\mathcal{G}$ implies that spending everything today (setting $c_t = m_t = 1$) produces lower marginal utility than is obtainable by reallocating a marginal unit of resources to the next period at return R:[55]

$$\begin{aligned} 1 &< (\mathsf{R}\beta)^{1/\gamma}\mathcal{G}^{-1} \\ 1 &< \mathsf{R}\beta\mathcal{G}^{-\gamma} \\ \mathrm{u}'(1) &< \mathsf{R}\beta\mathrm{u}'(\mathcal{G}). \end{aligned}$$

Similar logic shows that under these circumstances the constraint will never bind at $m = 1$ for a constrained consumer with a finite horizon of n periods, so for $m \geq 1$ such a consumer's consumption function will be the same as for the unconstrained case examined in the main text.

RIC fails, FHWC holds. If the RIC fails ($1 < \frac{\mathbf{\Phi}}{\mathsf{R}}$) while the finite human wealth condition holds, the limiting value of this consumption function as $n \uparrow \infty$ is the degenerate function

$$\grave{\mathrm{c}}_{T-n}(m) = 0(b_t + h). \tag{99}$$

(that is, consumption is zero for any level of human or nonhuman wealth).

RIC fails, FHWC fails. $\cancel{\text{FHWC}}$ implies that human wealth limits to $h = \infty$ so the consumption function limits to either $\grave{\mathrm{c}}_{T-n}(m) = 0$ or $\grave{\mathrm{c}}_{T-n}(m) = \infty$ depending on the relative speeds with which the MPC approaches zero and human wealth approaches ∞.[56]

Thus, the requirement that the consumption function be nondegenerate implies that for a consumer satisfying $\cancel{\text{GIC}}$ we must impose the RIC (and the FHWC can be shown to be a consequence of $\cancel{\text{GIC}}$ and RIC). In this case, the consumer's optimal behavior is easy to describe. We can calculate the point at which the unconstrained consumer would choose $c = m$ from Equation (98):

$$\begin{aligned} m_{\#} &= (m_{\#} - 1 + h)\underline{\kappa} \\ m_{\#}(1 - \underline{\kappa}) &= (h - 1)\underline{\kappa} \\ m_{\#} &= (h - 1)\left(\frac{\underline{\kappa}}{1 - \underline{\kappa}}\right) \end{aligned} \tag{100}$$

[55] The point at which the constraint would bind (if that point could be attained) is the $m = c$ for which $\mathrm{u}'(c_{\#}) = \mathsf{R}\beta\mathrm{u}'(\mathcal{G})$ which is $c_{\#} = \mathcal{G}/(\mathsf{R}\beta)^{1/\gamma}$ and the consumption function will be defined by $\grave{\mathrm{c}}(m) = \min[m, c_{\#} + (m - c_{\#})\underline{\kappa}]$.

[56] The knife-edge case is where $\mathbf{\Phi} = \mathcal{G}$, in which case the two quantities counterbalance and the limiting function is $\grave{\mathrm{c}}(m) = \min[m, 1]$.

which (under these assumptions) satisfies $0 < m_{\#} < 1$.[57] For $m < m_{\#}$ the unconstrained consumer would choose to consume more than m; for such m, the constrained consumer is obliged to choose $\grave{\mathrm{c}}(m) = m$.[58] For any $m > m_{\#}$ the constraint will never bind and the consumer will choose to spend the same amount as the unconstrained consumer, $\bar{\mathrm{c}}(m)$.

(Stachurski and Toda [2019] obtain a similar lower bound on consumption and use it to study the tail behavior of the wealth distribution.)

C.3 If GIC Holds

Imposition of the GIC reverses the inequality in (99), and thus reverses the conclusion: A consumer who starts with $m_t = 1$ will desire to consume more than 1. Such a consumer will be constrained, not only in period t, but perpetually thereafter.

Now define $b^n_{\#}$ as the b_t such that an unconstrained consumer holding $b_t = b^n_{\#}$ would behave so as to arrive in period $t + n$ with $b_{t+n} = 0$ (with $b^0_{\#}$ trivially equal to 0); for example, a consumer with $b_{t-1} = b^1_{\#}$ was on the 'cusp' of being constrained in period $t - 1$: Had b_{t-1} been infinitesimally smaller, the constraint would have been binding (because the consumer would have desired, but been unable, to enter period t with negative, not zero, b). Given the GIC, the constraint certainly binds in period t (and thereafter) with resources of $m_t = m^0_{\#} = 1 + b^0_{\#} = 1$: The consumer cannot spend more (because constrained), and will not choose to spend less (because impatient), than $c_t = c^0_{\#} = 1$.

We can construct the entire 'prehistory' of this consumer leading up to t as follows. Maintaining the assumption that the constraint has never bound in the past, c must have been growing according to $\frac{\Phi}{\mathcal{G}}$, so consumption n periods in the past must have been

$$c^n_{\#} = \frac{\Phi}{\mathcal{G}}^{-n} c_t = \frac{\Phi}{\mathcal{G}}^{-n}. \tag{101}$$ {eq:cPreHist}

The PDV of consumption from $t - n$ until t can thus be computed as

$$\begin{aligned}
\mathbb{C}^t_{t-n} &= c_{t-n}(1 + \Phi/\mathsf{R} + \cdots + (\Phi/\mathsf{R})^n) \\
&= c^n_{\#}(1 + \frac{\Phi}{\mathsf{R}} + \cdots + \frac{\Phi}{\mathsf{R}}^n) \\
&= \frac{\Phi}{\mathcal{G}}^{-n} \left(\frac{1 - \frac{\Phi}{\mathsf{R}}^{n+1}}{1 - \frac{\Phi}{\mathsf{R}}} \right) \\
&= \left(\frac{\frac{\Phi}{\mathcal{G}}^{-n} - \frac{\Phi}{\mathsf{R}}}{1 - \frac{\Phi}{\mathsf{R}}} \right)
\end{aligned}$$

[57] Note that $0 < m_{\#}$ is implied by RIC and $m_{\#} < 1$ is implied by ~~GIC~~.

[58] As an illustration, consider a consumer for whom $\Phi = 1$, $\mathsf{R} = 1.01$ and $\mathcal{G} = 0.99$. This consumer will save the amount necessary to ensure that growth in market wealth exactly offsets the decline in human wealth represented by $\mathcal{G} < 1$; total wealth (and therefore total consumption) will remain constant, even as market wealth and human wealth trend in opposite directions.

and note that the consumer's human wealth between $t-n$ and t (the relevant time horizon, because from t onward the consumer will be constrained and unable to access post-t income) is

$$h^n_\# = 1 + \cdots + \tilde{\mathcal{R}}^{-n} \tag{102}$$

while the intertemporal budget constraint says

$$\mathbb{C}^t_{t-n} \quad = b^n_\# + h^n_\#$$

from which we can solve for the $b^n_\#$ such that the consumer with $b_{t-n} = b^n_\#$ would unconstrainedly plan (in period $t-n$) to arrive in period t with $b_t = 0$:

$$b^n_\# = \mathbb{C}^t_{t-n} - \overbrace{\left(\frac{1-\tilde{\mathcal{R}}^{-(n+1)}}{1-\tilde{\mathcal{R}}^{-1}}\right)}^{h^n_\#}. \tag{103}$$ {eq:bPound}

Defining $m^n_\# = b^n_\# + 1$, consider the function $\grave{\mathrm{c}}(m)$ defined by linearly connecting the points $\{m^n_\#, c^n_\#\}$ for integer values of $n \geq 0$ (and setting $\grave{\mathrm{c}}(m) = m$ for $m < 1$). This function will return, for any value of m, the optimal value of c for a liquidity constrained consumer with an infinite horizon. The function is piecewise linear with 'kink points' where the slope discretely changes; for infinitesimal ϵ the MPC of a consumer with assets $m = m^n_\# - \epsilon$ is discretely higher than for a consumer with assets $m = m^n_\# + \epsilon$ because the latter consumer will spread a marginal dollar over more periods before exhausting it.

In order for a unique consumption function to be defined by this sequence (103) for the entire domain of positive real values of b, we need $b^n_\#$ to become arbitrarily large with n. That is, we need

$$\lim_{n\to\infty} b^n_\# = \infty. \tag{104}$$ {eq:bToInfty}

C.3.1 If FHWC Holds

The FHWC requires $\tilde{\mathcal{R}}^{-1} < 1$, in which case the second term in (103) limits to a constant as $n \uparrow \infty$, and (104) reduces to a requirement that

$$\begin{aligned}
\lim_{n\to\infty} \left(\frac{\frac{\mathbf{\mathfrak{b}}}{\mathcal{G}}^{-n} - (\frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}/\frac{\mathbf{\mathfrak{b}}}{\mathcal{G}})^n \frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}}{1-\frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}}\right) &= \infty \\
\lim_{n\to\infty} \left(\frac{\frac{\mathbf{\mathfrak{b}}}{\mathcal{G}}^{-n} - \tilde{\mathcal{R}}^{-n}\frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}}{1-\frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}}\right) &= \infty \\
\lim_{n\to\infty} \left(\frac{\frac{\mathbf{\mathfrak{b}}}{\mathcal{G}}^{-n}}{1-\frac{\mathbf{\mathfrak{b}}}{\mathsf{R}}}\right) &= \infty.
\end{aligned}$$

Given the GIC $\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}^{-1} > 1$, this will hold iff the RIC holds, $\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} < 1$. But given that the FHWC $\mathsf{R} > \mathcal{G}$ holds, the GIC is stronger (harder to satisfy) than the RIC; thus, the FHWC and the GIC together imply the RIC, and so a well-defined solution exists. Furthermore, in the limit as n approaches infinity, the difference between the limiting constrained consumption function and the unconstrained consumption function becomes vanishingly small, because the date at which the constraint binds becomes arbitrarily distant, so the effect of that constraint on current behavior shrinks to nothing. That is,

$$\lim_{m\to\infty} \grave{\mathrm{c}}(m) - \bar{\mathrm{c}}(m) = 0. \tag{105}$$

C.3.2 If FHWC Fails

If the FHWC fails, matters are a bit more complex.

Given failure of FHWC, (104) requires

$$\lim_{n\to\infty} \left(\frac{\tilde{\mathcal{R}}^{-n}\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} - \frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}^{-n}}{\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} - 1} \right) + \left(\frac{1 - \tilde{\mathcal{R}}^{-(n+1)}}{\tilde{\mathcal{R}}^{-1} - 1} \right) = \infty$$
$$\lim_{n\to\infty} \left(\frac{\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}}{\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} - 1} - \frac{\tilde{\mathcal{R}}^{-1}}{\tilde{\mathcal{R}}^{-1} - 1} \right) \tilde{\mathcal{R}}^{-n} - \left(\frac{\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}^{-n}}{\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}} - 1} \right) = \infty$$

If RIC Holds. When the RIC holds, rearranging (106) gives

$$\lim_{n\to\infty} \left(\frac{\frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}^{-n}}{1 - \frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}} \right) - \tilde{\mathcal{R}}^{-n} \left(\frac{\frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}}{1 - \frac{\mathbf{\mathfrak{p}}}{\mathsf{R}}} + \frac{\tilde{\mathcal{R}}^{-1}}{\tilde{\mathcal{R}}^{-1} - 1} \right) = \infty$$

and for this to be true we need

$$\begin{aligned} \frac{\mathbf{\mathfrak{p}}}{\mathcal{G}}^{-1} &> \tilde{\mathcal{R}}^{-1} \\ \mathcal{G}/\mathbf{\mathfrak{p}} &> \mathcal{G}/\mathsf{R} \\ 1 &> \mathbf{\mathfrak{p}}/\mathsf{R} \end{aligned}$$

which is merely the RIC again. So the problem has a solution if the RIC holds. Indeed, we can even calculate the limiting MPC from

$$\lim_{n\to\infty} \kappa^n_{\#} = \lim_{n\to\infty} \left(\frac{c^n_{\#}}{b^n_{\#}} \right) \tag{106}$$ {eq:MPCConstrLim

which with a bit of algebra[59] can be shown to asymptote to the MPC in the perfect foresight model:[60]

$$\lim_{m\to\infty} \grave{\boldsymbol{\kappa}}(m) = 1 - \frac{\Thorn}{\mathsf{R}}. \tag{108}$$

If RIC Fails. Consider now the ~~RIC~~ case, $\frac{\Thorn}{\mathsf{R}} > 1$. We can rearrange (106)as

$$\lim_{n\to\infty} \left(\frac{\frac{\Thorn}{\mathsf{R}}(\tilde{\mathcal{R}}^{-1}-1)}{(\tilde{\mathcal{R}}^{-1}-1)(\frac{\Thorn}{\mathsf{R}}-1)} - \frac{\tilde{\mathcal{R}}^{-1}(\frac{\Thorn}{\mathsf{R}}-1)}{(\tilde{\mathcal{R}}^{-1}-1)(\frac{\Thorn}{\mathsf{R}}-1)} \right) \tilde{\mathcal{R}}^{-n} - \left(\frac{\frac{\Thorn}{\mathcal{G}}^{-n}}{\frac{\Thorn}{\mathsf{R}}-1} \right) = \infty. \tag{109}$$

which makes clear that with ~~FHWC~~ $\Rightarrow \tilde{\mathcal{R}}^{-1} > 1$ and ~~RIC~~ $\Rightarrow \frac{\Thorn}{\mathsf{R}} > 1$ the numerators and denominators of both terms multiplying $\tilde{\mathcal{R}}^{-n}$ can be seen transparently to be positive. So, the terms multiplying $\tilde{\mathcal{R}}^{-n}$ in (106) will be positive if

$$\begin{aligned} \frac{\Thorn}{\mathsf{R}}\tilde{\mathcal{R}}^{-1} - \frac{\Thorn}{\mathsf{R}} &> \tilde{\mathcal{R}}^{-1}\frac{\Thorn}{\mathsf{R}} - \tilde{\mathcal{R}}^{-1} \\ \tilde{\mathcal{R}}^{-1} &> \frac{\Thorn}{\mathsf{R}} \\ \mathcal{G} &> \Thorn \end{aligned}$$

which is merely the GIC which we are maintaining. So the first term's limit is $+\infty$. The combined limit will be $+\infty$ if the term involving $\tilde{\mathcal{R}}^{-n}$ goes to $+\infty$ faster than the term involving $-\frac{\Thorn}{\mathcal{G}}^{-n}$ goes to $-\infty$; that is, if

$$\begin{aligned} \tilde{\mathcal{R}}^{-1} &> \frac{\Thorn}{\mathcal{G}}^{-1} \\ \mathcal{G}/\mathsf{R} &> \mathcal{G}/\Thorn \\ \Thorn/\mathsf{R} &> 1 \end{aligned}$$

which merely confirms the starting assumption that the RIC fails.

What is happening here is that the $c_{\#}^{n}$ term is increasing backward in time at rate dominated in the limit by $\mathcal{G}/\Thorn$ while the $b_{\#}$ term is increasing at a rate dominated by $\mathcal{G}/\mathsf{R}$ term and

$$\mathcal{G}/\mathsf{R} > \mathcal{G}/\Thorn \tag{110}$$

because ~~RIC~~ $\Rightarrow \Thorn > \mathsf{R}$.

---

[59] Calculate the limit of

$$\left(\frac{\frac{\Thorn}{\mathcal{G}}^{-n}}{\frac{\Thorn}{\mathcal{G}}^{-n}/(1-\frac{\Thorn}{\mathsf{R}}) - (1-\tilde{\mathcal{R}}^{-1}\tilde{\mathcal{R}}^{-n})/(1-\tilde{\mathcal{R}}^{-1})} \right) = \left(\frac{1}{1/(1-\frac{\Thorn}{\mathsf{R}}) + \tilde{\mathcal{R}}^{-n}\tilde{\mathcal{R}}^{-1}/(1-\tilde{\mathcal{R}}^{-1})} \right) \tag{107}$$

[60] For an example of this configuration of parameters, see the notebook `doApndxLiqConstr.nb` in the Mathematica software archive.

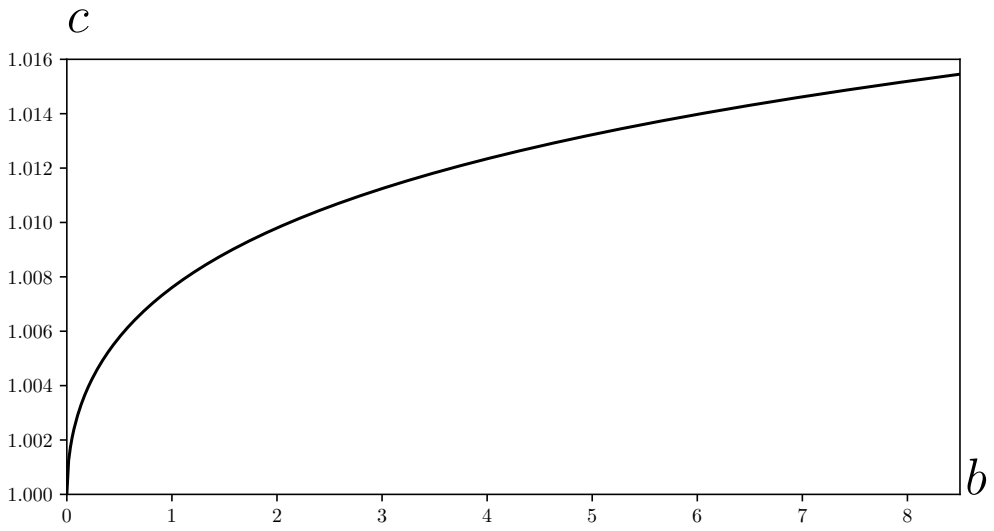

Figure 8 Appendix: Nondegenerate c Function with ~~FHWC~~ and ~~RIC~~ {fig:PFGICHoldsFH

Consequently, while $\lim_{n\uparrow\infty} b^n_\# = \infty$, the limit of the *ratio* $c^n_\#/b^n_\#$ in (106) is zero. Thus, surprisingly, the problem has a well defined solution with infinite human wealth if the RIC fails. It remains true that ~~RIC~~ implies a limiting MPC of zero,

$$\lim_{m\to\infty} \grave{\boldsymbol{\kappa}}(m) = 0, \tag{111}$$

but that limit is approached gradually, starting from a positive value, and consequently the consumption function is *not* the degenerate $\grave{\mathrm{c}}(m) = 0$. (Figure 8 presents an example for $\gamma = 2$, $\mathsf{R} = 0.98$, $\beta = 1.00$, $\mathcal{G} = 0.99$; note that the horizontal axis is bank balances $b = m-1$; the part of the consumption function below the depicted points is uninteresting — $c = m$ — so not worth plotting).

We can summarize as follows. Given that the GIC holds, the interesting question is whether the FHWC holds. If so, the RIC automatically holds, and the solution limits into the solution to the unconstrained problem as $m \uparrow \infty$. But even if the FHWC fails, the problem has a well-defined and nondegenerate solution, whether or not the RIC holds.

Although these results were derived for the perfect foresight case, we know from work elsewhere in this paper and in other places that the perfect foresight case is an upper bound for the case with uncertainty. If the upper bound of the MPC in the perfect foresight case is zero, it is not possible for the upper bound in the model with uncertainty to be greater than zero, because for any $\kappa > 0$ the level of consumption in the model with uncertainty would eventually exceed the level of consumption in the absence of uncertainty.

Ma and Toda [2020] characterize the limits of the MPC in a more general framework that allows for capital and labor income risks in a Markovian setting with liquidity constraints, and find that in that much more general framework the limiting MPC is also zero.

D Supporting Standard Results in Real Analysis

{fact:realanalysis}

Proposition 8. *Let* $\mathrm{f} : \mathbb{R}_{++} \to \mathbb{R}_{+}$ *be a continuous function. Consider sequences* x^n *in* $\mathbb{R}_{++}$ *and* $\mathrm{f}^n(x^n)$ *in* $\mathbb{R}_{+}$*. If* $\mathrm{f}^n(x^n) \to \mathrm{f}(x)$ *as* $n \to \infty$*, then* $x^n \to x$ *as* $n \to \infty$*.*

Proof. Given that f is continuous at x (with $x \in \mathbb{R}_{++}$), for every $\epsilon > 0$, there exists a $\delta > 0$ such that for all y in $\mathbb{R}_{++}$ with $|y - x| < \delta$, we have $|\mathrm{f}(y) - \mathrm{f}(x)| < \epsilon$.

Given $\mathrm{f}^n(x^n) \to \mathrm{f}(x)$, for the above ϵ, there exists an N such that for all $n > N$, $|\mathrm{f}^n(x^n) - \mathrm{f}(x)| < \epsilon$.

Assume for the sake of contradiction that x^n doesn't converge to x. This implies that there exists a $\delta > 0$ such that for infinitely many terms of the sequence x^n, $|x^n - x| \geq \delta$.

By the continuity of f at x, if $|x^n - x| \geq \delta$ for infinitely many n, then $|\mathrm{f}^n(x^n) - \mathrm{f}(x)| \geq \epsilon$ for those n, contradicting our assumption that $\mathrm{f}^n(x^n) \to \mathrm{f}(x)$.

Therefore, our assumption for contradiction is false, and it follows that $x^n \to x$ as $n \to \infty$. □

Fact 1. *Let* $\mathrm{g} : X \to \mathbb{R}_{+}$ *be a continuous function, where* $X \subseteq \mathbb{R}^n$ *is an open convex set. Define the weighted supremum norm* $\|\cdot\|_{\mathrm{g}}$ *of a real-valued function* $\mathrm{f} : X \to \mathbb{R}$ *by*

$$\|\mathrm{f}\|_{\mathrm{g}} := \sup_{x \in X} \frac{|\mathrm{f}(x)|}{\mathrm{g}(x)}. \tag{112}$$

If $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}^\star\|_{\mathrm{g}} = 0$*,* f_n *converges to* $\mathrm{f}^\star$ *uniformly on compact sets.*

Proof. Let $\tilde{X}$ be an arbitrary compact subset of X. Since $\tilde{X}$ is compact, there exists a positive lower bound for g on this subset, denoted as

$$\bar{\mathrm{g}} = \min_{x \in \tilde{X}} \mathrm{g}(x) > 0. \tag{113}$$

Hence, on $\tilde{X}$, if $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}^\star\|_{\mathrm{g}} = 0$, then $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}^\star\|_{\infty} = 0$ on $\tilde{X}$, where $\|\cdot\|_{\infty}$ denotes the supremum norm.

Now, let K be a compact subset of X. Given the continuity of g, there exists a positive maximum value for g on K, denoted as M_K. Then, we have

$$\sup_{x \in K} |\mathrm{f}_n(x) - \mathrm{f}(x)| \leq M_K \sup_{x \in K} \frac{|\mathrm{f}_n(x) - \mathrm{f}(x)|}{\mathrm{g}(x)} \leq M_K \sup_{x \in X} \frac{|\mathrm{f}_n(x) - \mathrm{f}(x)|}{\mathrm{g}(x)}. \tag{114}$$

Thus, $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}\|_{=}0$ implies that f_n converges uniformly to f on the compact set K. It's also worth noting that the convexity and openness of X aren't strictly necessary for this argument. □

{fact:compactnt}

Fact 2. *Let* $\{\mathrm{f}_n\}$ *be a sequence of continuous functions defined on a subset of the real line and converging uniformly to a function*f *on compact sets. If*$\{x_n\}$ *is a convergent sequence of real numbers with limit* x*, then* $\mathrm{f}_n(x_n)$ *converges to* $\mathrm{f}(x)$*.*

Proof. Let $\tilde{X}$ be an arbitrary compact subset of X. Since $\tilde{X}$ is compact, there exists a positive lower bound for g on this subset, denoted as

$$\bar{\mathrm{g}} = \min_{x \in \tilde{X}} \mathrm{g}(x) > 0. \tag{115}$$

Hence, on $\tilde{X}$, if $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}^\star\|_{\mathrm{g}} = 0$, then $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}^\star\|_\infty = 0$ on $\tilde{X}$, where $\|\cdot\|_\infty$ denotes the supremum norm.

Now, let K be a compact subset of X. Given the continuity of g, there exists a positive maximum value for g on K, denoted as M_K. Then, we have

$$\sup_{x\in K} |\mathrm{f}_n(x) - \mathrm{f}(x)| \leq M_K \sup_{x\in K} \frac{|\mathrm{f}_n(x) - \mathrm{f}(x)|}{\mathrm{g}(x)} \leq M_K \sup_{x\in X} \frac{|\mathrm{f}_n(x) - \mathrm{f}(x)|}{\mathrm{g}(x)}. \tag{116}$$

Thus, $\lim_{n\to\infty} \|\mathrm{f}_n - \mathrm{f}\|_{=}0$ implies that f_n converges uniformly to f on the compact set K. It's also worth noting that the convexity and openness of X aren't strictly necessary for this argument. □

{fact:compactnt}

Fact 3. *Let* $\{\mathrm{f}_n\}$ *be a sequence of continuous functions defined on a subset of the real line and converging uniformly to a function* f *on compact sets. If* $\{x_n\}$ *is a convergent sequence of real numbers with limit* x*, then* $\mathrm{f}_n(x_n)$ *converges to* $\mathrm{f}(x)$*.*

Proof. Since x_n converges to x, the sequence $\{x_n\}$ is bounded. Therefore, there exists a compact set K (specifically, a closed interval in the real line) that contains all the x_n as well as x.

Given the uniform convergence of f_n to f on K, for every $\epsilon > 0$, there exists an N such that for all $n \geq N$ and for all $y \in K$, we have

$$|\mathrm{f}_n(y) - \mathrm{f}(y)| < \epsilon.$$

In particular, for $y = x_n$, we have

$$|\mathrm{f}_n(x_n) - \mathrm{f}(x_n)| < \epsilon$$

for all $n \geq N$.

Now, each f_n being continuous and x_n converging to x implies that $\mathrm{f}(x_n)$ converges to $\mathrm{f}(x)$. Thus, for sufficiently large n, $\mathrm{f}(x_n)$ can be made arbitrarily close to $\mathrm{f}(x)$.

Combining the two inequalities and taking n large enough, we deduce that $|\mathrm{f}_n(x_n) - \mathrm{f}(x)|$ can be made smaller than any given ϵ. Hence, $\mathrm{f}_n(x_n)$ converges to $\mathrm{f}(x)$. □

E Relational Diagrams for the Inequality Conditions

{sec:ApndxConditio

This appendix explains in detail the paper's 'inequalities' diagrams (Figures 6, 7).

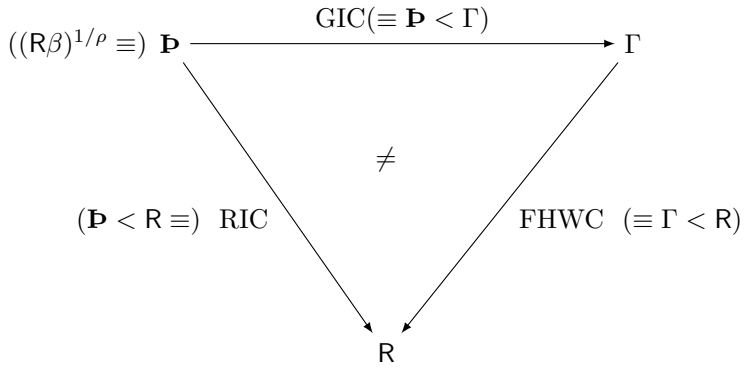

Figure 9 Appendix: Inequality Conditions for Perfect Foresight Model
(Start at a node and follow arrows)

{fig:InequalityPFGI

E.1 The Unconstrained Perfect Foresight Model

A simple illustration is presented in Figure 9, whose three nodes represent values of the absolute patience factor $\mathbf{\Þ}$, the permanent-income growth factor $\mathcal{G}$, and the riskfree interest factor R. The arrows represent imposition of the labeled inequality condition (like, the uppermost arrow, pointing from $\mathbf{\Þ}$ to $\mathcal{G}$, reflects imposition of the GIC condition (clicking GIC should take you to its definition; definitions of other conditions are also linked below)).[61] Annotations inside parenthetical expressions containing $\equiv$ are there to make the diagram readable for someone who may not immediately remember terms and definitions from the main text. (Such a reader might also want to be reminded that R, β, and Γ are all in $\mathbb{R}_{++}$, and that $\gamma > 1$).

Navigation of the diagram is simple: Start at any node, and deduce a chain of inequalities by following any arrow that exits that node, and any arrows that exit from successive nodes. Traversal must stop upon arrival at a node with no exiting arrows. So, for example, we can start at the $\mathbf{\Þ}$ node and impose the GIC and then the FHWC, and see that imposition of these conditions allows us to conclude that $\mathbf{\Þ} < \mathsf{R}$.

One could also impose $\mathbf{\Þ} < \mathsf{R}$ directly (without imposing GIC and FHWC) by following the downward-sloping diagonal arrow exiting $\mathbf{\Þ}$. Although alternate routes from one node to another all justify the same core conclusion ($\mathbf{\Þ} < \mathsf{R}$, in this case), $\neq$ symbol in the center is meant to convey that these routes are not identical in other respects. This notational convention is used in category theory diagrams,[62] to indicate that the diagram is not commutative.[63]

Negation of a condition is indicated by the reversal of the corresponding arrow. For example, negation of the RIC, ~~RIC~~ $\equiv \mathbf{\Þ} > \mathsf{R}$, would be represented by moving the

[61] For convenience, the equivalent ($\equiv$) mathematical statement of each condition is expressed nearby in parentheses.

[62] For a popular introduction to category theory, see Riehl [2017].

[63] But the rest of our notation does not necessarily abide by the other conventions of category theory diagrams.

arrowhead from the bottom right to the top left of the line segment connecting $\mathbb{Þ}$ and R.

If we were to start at R and then impose ~~FHWC~~, that would reverse the arrow connecting R and $\mathcal{G}$, but the $\mathcal{G}$ node would then have no exiting arrows so no further deductions could be made. However, if we *also* reversed GIC (that is, if we imposed ~~GIC~~), that would take us to the $\mathbb{Þ}$ node, and we could deduce $\mathsf{R} > \mathbb{Þ}$. However, we would have to stop traversing the diagram at this point, because the arrow exiting from the $\mathbb{Þ}$ node points back to our starting point, which (if valid) would lead us to the conclusion that $\mathsf{R} > \mathsf{R}$. Thus, the reversal of the two earlier conditions (imposition of ~~FHWC~~ and ~~GIC~~) requires us also to reverse the final condition, giving us ~~RIC~~.[64]

Under these conventions, Figure 6 in the main text presents a modified version of the diagram extended to incorporate the PF-FVAC (reproduced here for convenient reference).

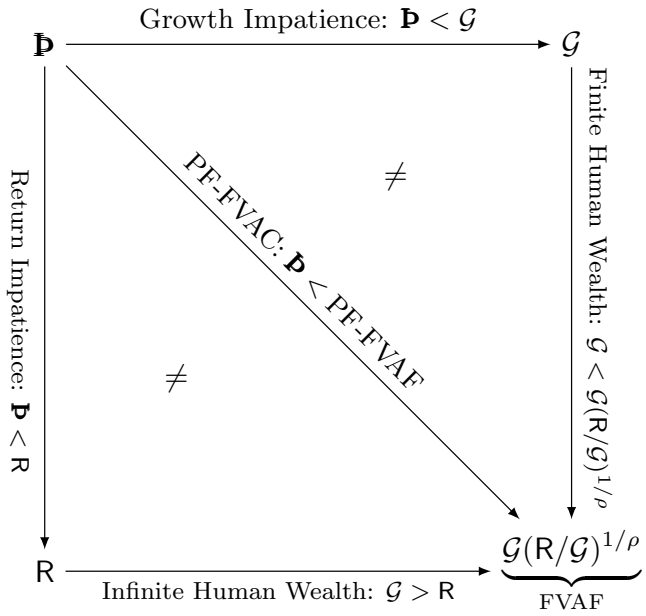

Figure 10 Appendix: Relation of GIC, FHWC, RIC, and PFVAC {fig:RelatePFGICFH
An arrowhead points to the larger of the two quantities being compared. For example, the diagonal arrow indicates that $\mathbb{Þ} < \mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma}$, which is an alternative way of writing the PF-FVAC, (9)

This diagram can be interpreted, for example, as saying that, starting at the $\mathbb{Þ}$ node, it is possible to derive the PF-FVAC[65] by imposing both the GIC and the FHWC; or by imposing RIC and ~~FHWC~~. Or, starting at the $\mathcal{G}$ node, we can follow the imposition

[64]The corresponding algebra is

$$\begin{aligned} \text{~~FHWC~~}: &\quad \mathsf{R} < \mathcal{G} \\ \text{~~GIC~~}: &\quad \mathcal{G} < \mathbb{Þ} \\ \Rightarrow \text{~~RIC~~}: &\quad \mathsf{R} < \mathbb{Þ}, \end{aligned}$$

.

[65]in the form $\mathbb{Þ} < (\mathsf{R}/\mathcal{G})^{1/\gamma}\mathcal{G}$

of the FHWC (twice — reversing the arrow labeled ~~FHWC~~) and then ~~RIC~~ to reach the conclusion that $\Ppatience < \mathcal{G}$. Algebraically,

$$\begin{aligned} \text{FHWC}: \quad & \mathcal{G} < \mathsf{R} \\ \text{~~RIC~~}: \quad & \mathsf{R} < \Ppatience \\ & \mathcal{G} < \Ppatience \end{aligned} \tag{117}$$ {eq:cnclGICRaw}

which leads to the negation of both of the conditions leading into $\Ppatience$. ~~GIC~~ is obtained directly as the last line in (117) and ~~PF-FVAC~~ follows if we start by multiplying the Return Patience Factor (RPF=$\Ppatience/\mathsf{R}$) by the FHWF (=$\mathcal{G}/\mathsf{R}$) raised to the power $1/\gamma - 1$, which is negative since we imposed $\gamma > 1$. FHWC implies FHWF < 1 so when FHWF is raised to a negative power the result is greater than one. Multiplying the RPF (which exceeds 1 because ~~RIC~~) by another number greater than one yields a product that must be greater than one:

$$\begin{aligned} 1 &< \overbrace{\left(\frac{(\mathsf{R}\beta)^{1/\gamma}}{\mathsf{R}}\right)}^{>1 \text{ from } \text{~~RIC~~}} \overbrace{(\mathcal{G}/\mathsf{R})^{1/\gamma - 1}}^{>1 \text{ from FHWC}} \\ 1 &< \left(\frac{(\mathsf{R}\beta)^{1/\gamma}}{(\mathsf{R}/\mathcal{G})^{1/\gamma}\mathsf{R}\mathcal{G}/\mathsf{R}}\right) \\ \mathsf{R}^{1/\gamma}\mathcal{G}^{1-1/\gamma} = (\mathsf{R}/\mathcal{G})^{1/\gamma}\mathcal{G} &< \Ppatience \end{aligned}$$

which is one way of writing ~~PF-FVAC~~.

The complexity of this algebraic calculation illustrates the usefulness of the diagram, in which one merely needs to follow arrows to reach the same result.

After the warmup of constructing these conditions for the perfect foresight case, we can represent the relationships between all the conditions in both the perfect foresight case and the case with uncertainty as shown in Figure 7 in the paper (reproduced here).

Finally, the next diagram substitutes the values of the various objects in the diagram under the baseline parameter values and verifies that all of the asserted inequality conditions hold true.

F Apparent Balanced Growth in $\bar{c}$ and $\mathbf{cov}(c, \boldsymbol{p})$

{sec:ApndxBalance

Section 4.2 demonstrates some propositions under the assumption that, when an economy satisfies the GIC, there will be constant growth factors $\Omega_{\bar{c}}$ and Ω_{cov} respectively for $\bar{c}$ (the average value of the consumption ratio) and $\text{cov}(c, \boldsymbol{p})$. In the case of a Szeidl-invariant economy, the main text shows that these are $\Omega_{\bar{c}} = 1$ and $\Omega_{\text{cov}} = \mathcal{G}$. If the economy is Harmenberg- but not Szeidl-invariant, no proof is offered that these growth factors will be constant.

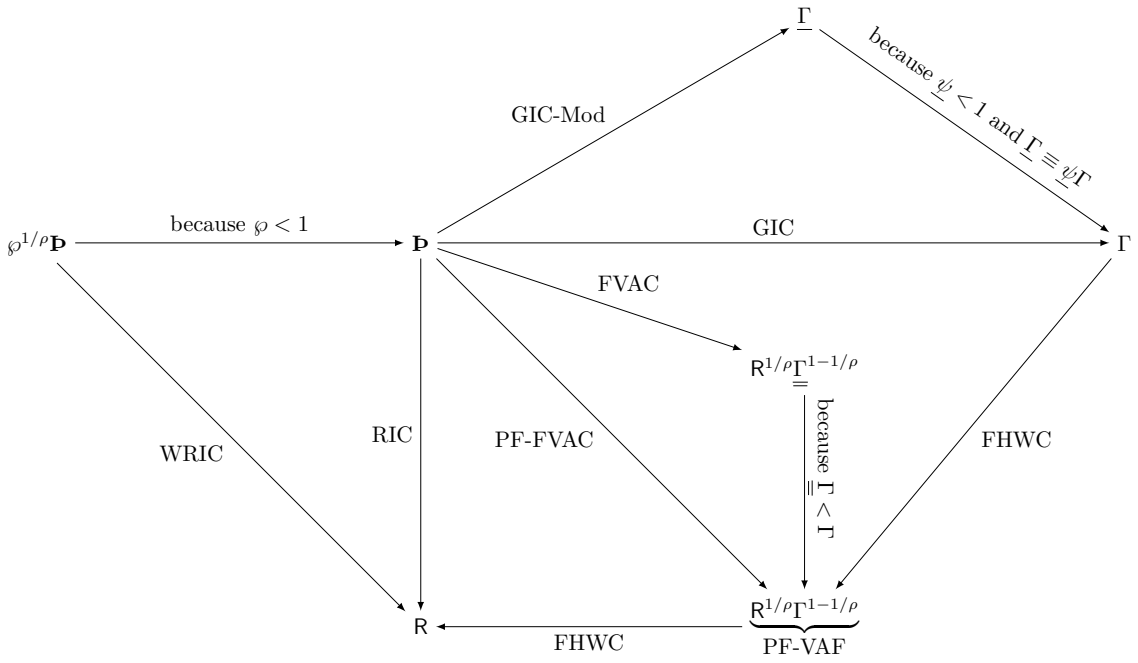

Figure 11 Appendix: Relation of All Inequality Conditions {fig:InequalitiesApp

F.1 $\log c$ and $\log(\text{cov}(c, \boldsymbol{p}))$ Grow Linearly

Figures 13 and 14 plot the results of simulations of an economy that satisfies Harmenberg- but not Szeidl-invariance with a population of 4 million agents over the last 1000 periods (of a 2000 period simulation).[66] The first figure shows that $\log \bar{c}$ increases apparently linearly. The second figure shows that $\log(-\text{cov}(c, \boldsymbol{p}))$ also increases apparently linearly. (These results are produced by the notebook `ApndxBalancedGrowthcNrmAndCov.ipynb`).

[66] For an exposition of our implementation of Harmenberg's method, see this supplemental appendix.

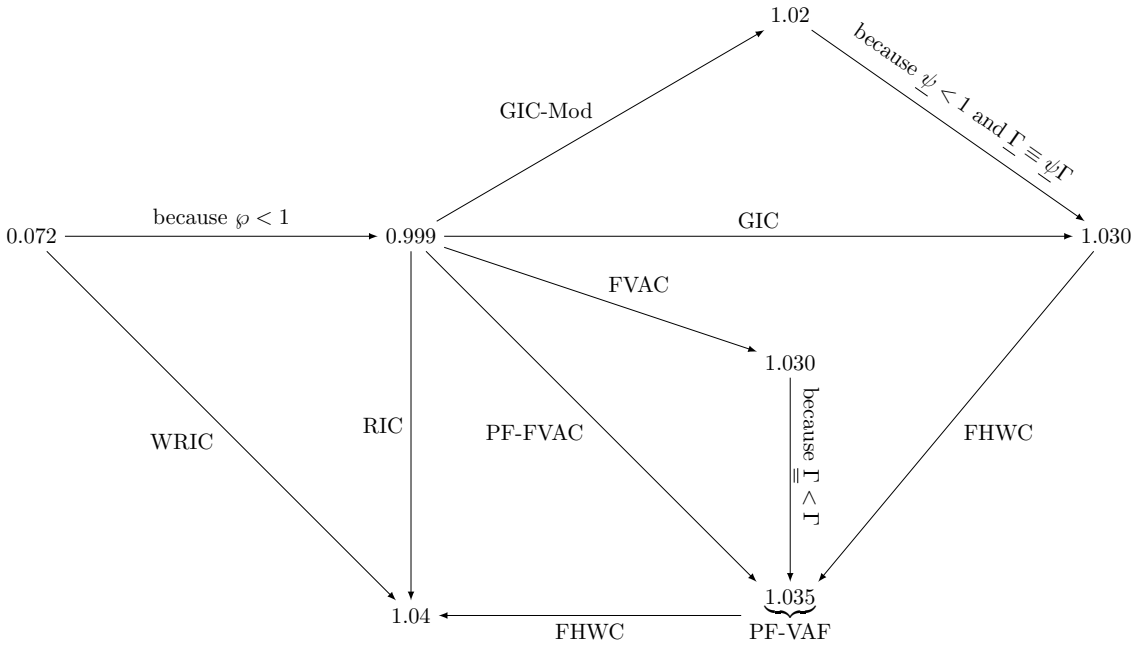

Figure 12 Appendix: Numerical Relation of All Inequality Conditions {fig:InequalitiesApp

References

John M. Abowd and David Card. On the covariance structure of earnings and hours changes. Econometrica, 57:411–445, 1989.

S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. Quarterly Journal of Economics, 109:659–684, 1994.

Fernando Alvarez and Nancy L Stokey. Dynamic programming with homogeneous functions. Journal of economic theory, 82(1):167–189, 1998.

Jess Benhabib, Alberto Bisin, and Shenghao Zhu. The wealth distribution in bewley economies with capital income risk. Journal of Economic Theory, 159:489–515, 2015. URL `https://www.sciencedirect.com/science/article/pii/S0022053115001362`. Available at `https://www.nber.org/papers/w20157.pdf`.

Truman Bewley. The permanent income hypothesis: A theoretical formulation. Journal of Economic Theory, 16:252–292, 1977.

Olivier J. Blanchard. Debt, deficits, and finite horizons. Journal of Political Economy, 93(2):223–247, April 1985.

John H. Boyd. Recursive utility and the ramsey problem. Journal of Economic Theory, 50(2):326–345, April 1990.

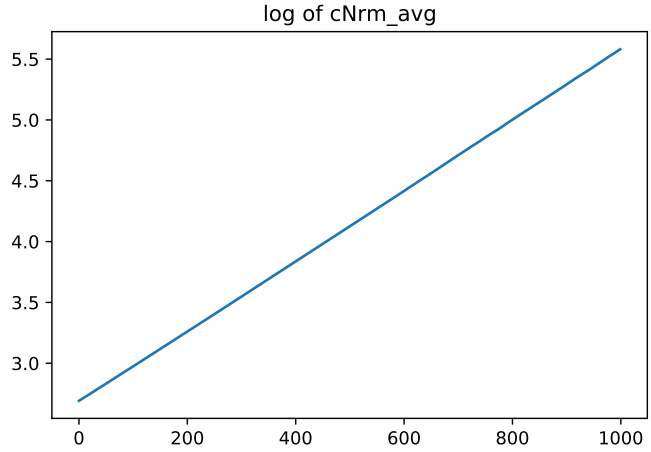

Figure 13 Appendix: log $\mathfrak{c}$ Appears to Grow Linearly {fig:logcNrm}

Marco Cagetti. Wealth accumulation over the life cycle and precautionary savings. Journal of Business and Economic Statistics, 21(3):339–353, 2003.

Christopher Carroll, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. The distribution of wealth and the marginal propensity to consume. Quantitative Economics, 8(3):977–1020, 2017. doi: https://doi.org/10.3982/QE694. URL `https://onlinelibrary.wiley.com/doi/abs/10.3982/QE694`.

Christopher D. Carroll. The buffer-stock theory of saving: Some macroeconomic evidence. Brookings Papers on Economic Activity, 1992(2):61–156, 1992. URL `https://www.econ2.jhu.edu/people/ccarroll/BufferStockBPEA.pdf`. `https://www.econ2.jhu.edu/people/ccarroll/BufferStockBPEA.pdf`.

Christopher D. Carroll. Buffer stock saving and the life cycle/permanent income hypothesis. Quarterly Journal of Economics, CXII(1):1–56, 1997.

Christopher D. Carroll. Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order approximation). Advances in Macroeconomics, 1(1):Article 6, 2001. doi: 10.2202/1534-6013.1003.

Christopher D. Carroll. Theoretical foundations of buffer stock saving. Quantitative Economics, 2019, Submitted. doi: https://zenodo.org/badge/latestdoi/304124725. URL `https://econ-ark.github.io/BufferStockTheory`.

Christopher D. Carroll and Miles S. Kimball. On the Concavity of the Consumption Function. Econometrica, 64(4):981–992, 1996. URL `https://www.econ2.jhu.edu/people/ccarroll/concavity.pdf`. `https://www.econ2.jhu.edu/people/ccarroll/concavity.pdf`.

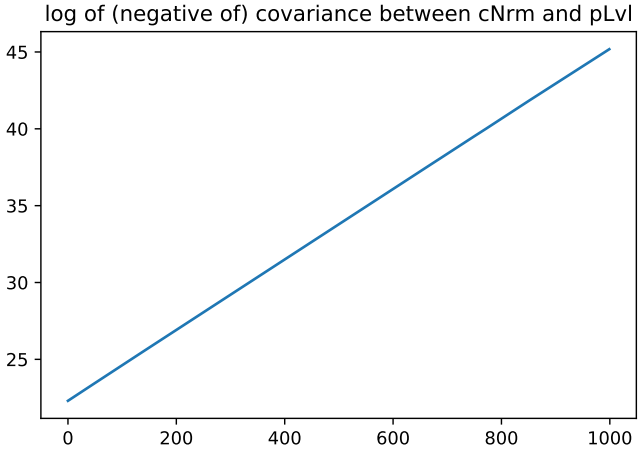

Figure 14 Appendix: log $(-\text{cov}(c, \boldsymbol{p}))$ Appears to Grow Linearly {fig:logcov}

Christopher D. Carroll and Miles S. Kimball. Precautionary saving and precautionary wealth. Palgrave Dictionary of Economics and Finance, 2nd Ed., 2007. URL `https://www.econ2.jhu.edu/people/ccarroll/papers/PalgravePrecautionary.pdf`. `https://www.econ2.jhu.edu/people/ccarroll/papers/PalgravePrecautionary.pdf`.

Christopher D. Carroll and Andrew A. Samwick. The Nature of Precautionary Wealth. Journal of Monetary Economics, 40(1):41–71, 1997. URL `https://www.econ2.jhu.edu/people/ccarroll/papers/nature.pdf`.

Christopher D. Carroll, Martin Holm, and Miles S. Kimball. Liquidity constraints and precautionary saving. Manuscript, Johns Hopkins University, 2019. URL `https://www.econ2.jhu.edu/people/ccarroll/papers/LiqConstr`. `https://www.econ2.jhu.edu/people/ccarroll/papers/LiqConstr`.

Gary Chamberlain and Charles A. Wilson. Optimal intertemporal consumption under uncertainty. Review of Economic Dynamics, 3(3):365–395, 2000.

Christopher D. Carroll, Alexander M. Kaufman, Jacqueline L. Kazil, Nathan M. Palmer, and Matthew N. White. The Econ-ARK and HARK: Open Source Tools for Computational Economics. In Fatih Akici, David Lippa, Dillon Niederhut, and M Pacer, editors, Proceedings of the 17th Python in Science Conference, pages 25 – 30, 2018. doi: 10.25080/Majora-4af1f417-004.

Richard H. Clarida. Consumption, liquidity constraints, and asset accumulation in the face of random fluctuations in income. International Economic Review, XXVIII:339–351, 1987.

Moira Daly, Dmytro Hryshko, and Iourii Manovskii. Improving the measurement of earnings dynamics. Technical report, National Bureau of Economic Research, 2016.

Angus S. Deaton. Saving and liquidity constraints. Econometrica, 59: 1221–1248, 1991. URL http://www.jstor.org/stable/2938366. https://www.jstor.org/stable/2938366.

Jorge Durán. Discounting long run average growth in stochastic dynamic programs. Economic Theory, 22(2):395–413, 2003.

Irving Fisher. The Theory of Interest. MacMillan, New York, 1930.

Milton A. Friedman. A Theory of the Consumption Function. Princeton University Press, 1957.

Pierre-Olivier Gourinchas and Jonathan Parker. Consumption over the life cycle. Econometrica, 70(1):47–89, 2002.

Karl Harmenberg. Aggregating heterogeneous-agent models with permanent income shocks. Journal of Economic Dynamics and Control, 129:104185, 2021. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2021.104185. URL https://www.sciencedirect.com/science/article/pii/S0165188921001202.

Lutz Hendricks. Bequests and Retirement Wealth in the United States. University of Arizona, 2001. URL https://lhendricks.org/Research/wps/bequdata_paper.pdf.

Lutz Hendricks. Wealth distribution and bequests. Lecture Notes, Economics 821, University of North Carolina, 2016. URL https://lhendricks.org/econ821/wealth/wealth_bequ_sl.pdf.

Ryoji Hiraguchi. On the convergence of consumption rules. Manuscript, Johns Hopkins University, 2003.

Dmytro Hryshko and Iourii Manovskii. How much consumption insurance in the us? Manuscript, University of Alerta, August 2020. URL http://www.artsrn.ualberta.ca/econweb/hryshko/Papers/HryshkoManovskii-how-much-aug25-202

Tullio Jappelli and Luigi Pistaferri. Intertemporal choice and consumption mobility. Econometric Society World Congress 2000 Contributed Paper Number 0118, August 2000. URL http://fmwww.bc.edu/RePEc/es2000/0118.pdf.

Miles S. Kimball. Standard risk aversion. Econometrica, 61(3):589–611, May 1993.

Dirk Krueger, Kurt Mitman, and Fabrizio Perri. Macroeconomics and household heterogeneity. Handbook of Macroeconomics, 2:843–921, 2016. doi: 10.3386/w22319. URL http://www.nber.org/papers/w22319.

Huiyu Li and John Stachurski. Solving the income fluctuation problem with unbounded rewards. Journal of Economic Dynamics and Control, 45:353–365, 2014.

Qingyin Ma and Alexis Akira Toda. A theory of the saving rate of the rich, 2020.

Qingyin Ma, John Stachurski, and Alexis Akira Toda. The income fluctuation problem and the evolution of wealth. Journal of Economic Theory, 187, 2020.

Thomas MaCurdy. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. Journal of Econometrics, 18(1):83–114, 1982.

V Filipe Martins-da Rocha and Yiannis Vailakis. Existence and uniqueness of a fixed point for local contractions. Econometrica, 78(3):1127–1141, 2010.

Janusz Matkowski and Andrzej S. Nowak. On discounted dynamic programming with unbounded returns. Economic Theory, 46:455–474, 2011.

Franco Modigliani. The life cycle hypothesis, the demand for wealth, and the supply of capital. Social Research, 33:160–217, 1966.

John F. Muth. Optimal properties of exponentially weighted forecasts. Journal of the American Statistical Association, 55(290):299–306, 1960.

Emily Riehl. Category theory in context. Courier Dover Publications, 2017.

Juan Pablo Rincón-Zapatero and Carlos Rodríguez-Palmero. Existence and uniqueness of solutions to the bellman equation in the unbounded case. Econometrica, 71(5): 1519–1555, 2003.

Jack Schechtman and Vera Escudero. Some results on 'an income fluctuation problem'. Journal of Economic Theory, 16:151–166, 1977.

José Scheinkman and Laurence Weiss. Borrowing constraints and aggregate economic activity. Econometrica, 54(1):23–46, 1986.

Stephanie Schmitt-Grohé and Martın Uribe. Closing small open economy models. Journal of international Economics, 61(1):163–185, 2003.

John Stachurski and Alexis Akira Toda. An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity. Journal of Economic Theory, 182:1–24, July 2019. doi: 10.1016/j.jet.2019.04.001.

Nancy L. Stokey, Robert E. Lucas, and Edward C. Prescott. Recursive Methods in Economic Dynamics. Harvard University Press, 1989.

Adam Szeidl. Stable invariant distribution in buffer-stock saving and stochastic growth models. Manuscript, Central European University, 2013. Available at `http://www.personal.ceu.hu/staff/Adam_Szeidl/papers/invariant_revision.pdf`.

Patrick Toche. A tractable model of precautionary saving in continuous time. Economics Letters, 87(2):267–272, 2005.

Menahem E Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. The Review of Economic Studies, 32(2):137–150, 1965.

Jiaxiong Yao. The theoretical foundations of buffer stock saving: A note. Manuscript, Johns Hopkins University, 2012.

Stephen P. Zeldes. Optimal consumption with stochastic income: Deviations from certainty equivalence. Quarterly Journal of Economics, 104(2):275–298, May 1989.

Table 5 Appendix: Perfect Foresight Liquidity Constrained Taxonomy

For constrained $\grave{c}$ and unconstrained $\bar{c}$ consumption functions

Main Condition Subcondition	Math	Outcome, Comments or Results
~~GIC~~	$1 < \Phi/\mathcal{G}$	Constraint never binds for $m \geq 1$
and RIC	$\Phi/\mathsf{R} < 1$	FHWC holds ($\mathsf{R} > \mathcal{G}$); $\grave{c}(m) = \bar{c}(m)$ for $m \geq 1$
and ~~RIC~~	$1 < \Phi/\mathsf{R}$	$\grave{c}(m)$ is degenerate: $\grave{c}(m) = 0$
GIC	$\Phi/\mathcal{G} < 1$	Constraint binds in finite time $\forall\ m$
and RIC	$\Phi/\mathsf{R} < 1$	FHWC may or may not hold $\lim_{m\uparrow\infty} \bar{c}(m) - \grave{c}(m) = 0$ $\lim_{m\uparrow\infty} \grave{\kappa}(m) = \underline{\kappa}$
and ~~RIC~~	$1 < \Phi/\mathsf{R}$	~~FHWC~~ $\lim_{m\uparrow\infty} \grave{\kappa}(m) = 0$

Conditions are applied from left to right; for example, the second row indicates conclusions in the case where ~~GIC~~ and RIC both hold, while the third row indicates that when the GIC and the RIC both fail, the consumption function is degenerate; the next row indicates that whenever the GICholds, the constraint will bind in finite time.