



Università
Ca' Foscari
Venezia

Bachelor's Degree
programme in
Economics, Markets
and Finance

Supervisor

Ch. Prof. Pietro Enrico Dino Dindo

Graduand

Mattia Cogo
Matriculation Number: 902265


Academic Year

2025/2026

Final Thesis

When nothing is worth something

Intrinsically worthless assets in OLG frameworks



When nothing is worth something

Intrinsically worthless assets in OLG frameworks

Mattia Cogo

Abstract

This thesis examines the role of money and other intrinsically valueless assets as stores of value within the framework of overlapping generations. Starting from Samuelson’s seminal contribution, which demonstrates how fiat money can improve efficiency by enabling intertemporal transfers. The analysis turns to Diamond’s extension, where public debt serves as a substitute for money. Finally, Tirole’s work on asset bubbles is considered, showing how bubbles despite lacking fundamental value can persist and even enhance welfare under certain dynamic conditions. By comparing these three classical contributions, the thesis aims towards a unifying theme: the ability of “social contrivances” such as money, debt, and bubbles to serve as intertemporal stores of value. The conclusion links these theoretical insights to present world, emphasizing that such contrivances remain central in understanding modern assets and financial exchanges.

Dedication

*A chi ha sostenuto i miei passi e a chi, dubitandone,
mi ha spinto a correre più forte.*

(To those who supported my steps, and to those who,
by doubting them, pushed me to run faster.)

Acknowledgements

I would like to express my deepest gratitude to Professor Dindo for guiding my work with extreme competence, patience and valuable suggestions throughout the entire writing this thesis.

Furthermore, I am profoundly grateful to Professor Pellizzari for his availability, patience, care, interest, and invaluable suggestions throughout the design and analysis of the simulation presented in this work.

I would like to thank Ca' Foscari University and Goethe-Universität for providing me with a stimulating academic environment and the resources necessary for my studies.

I am infinitely grateful to my parents, my sister Mariavittoria, Tommaso, and my entire family for supporting, encouraging and spurring me on during my studies: without their support, I would not be here. Thank you for your interest in my career, my achievements and my future plans.

I would like to thank my colleagues, especially Alex, with whom I shared most of my journey and my stay in Frankfurt. I thank him for being a valuable source of discussion and comparison throughout almost my entire university career.

I would like to thank my all my friends, especially Alessandro, Simone, Noemi, and Paolyne, for their support and encouragement during the most intense moments.

Contents

Introduction	1
1 From Barter to Institution: The Social Foundations of Money	2
1.1 The Problem of Value Without Worth	2
1.2 Money as a Social Contrivance	3
1.3 From Barter to Monetary Exchange	4
2 The Rational Basis of Fiat Money	8
2.1 Why Samuelson Is Central to This Thesis	8
2.2 The Model Setup	8
2.3 The Stationary Population Case	11
2.4 The Growing Population Case	12
2.5 The Biological Interest Rate	12
2.6 The Infinity Paradox: A Failure of Coordination	14
2.7 Why Pareto Efficiency Holds but Society Is Worse Off?	16
2.8 The Social Contrivance of Money	17
2.9 Critiques of Samuelson	21
3 Claims on the Future: Debt as a Store of Value	24
3.1 The Importance of Diamond's Extension	24
3.2 Diamond's OLG Model	24
3.3 Competitive Equilibrium and the Possibility of Failure	27
3.4 Debt as Competing Assets and Its Dual Nature	31
3.4.1 External Debt - The Benchmark	32
3.4.2 Internal Debt - The Portfolio Substitution Channel	39
3.4.3 Internal Versus External Debt	45
3.5 Remarks	47
3.6 Safe Assets and the Demand for Debt	48
3.6.1 What Makes Debt Safe	49
3.6.2 Collateral, Opacity, and the Production of Safe Debt	50
3.6.3 The Convenience Yield	51
3.6.4 Regulation and the Rise of Shadow Banking	51
3.6.5 Implications for Diamond's Framework	53
4 When Bubbles Enhance Welfare	55
4.1 The Role of Tirole	55
4.2 The Bubbly Economy	56
4.3 The Existence Condition	57
4.4 The Return to the Golden Rule	60

4.5	Three Faces of One Phenomenon, and What Sustains Them	63
5	Computational Evidence from an OLG Simulation	66
5.1	Stationary Economies	70
5.2	Growing Economies	74
5.3	Remarks	79
6	Is Bitcoin Samuelson Money?	81
6.1	Bitcoin as Samuelson Money	81
6.1.1	The Informal Arguments Translated	83
6.2	Medium of Exchange	84
6.2.1	The El Salvador Natural Experiment	85
6.3	Store of Value	86
6.4	Unit of Account	88
	Conclusions	91
	Bibliography	97
	Appendix A	99
	Declaration	110

Introduction

When analysing how agents allocate resources over time or exchange them with one another, it is reasonable to assume that such exchanges occur through monetary transactions. Money serves as the natural medium that enables goods, services, credit, and savings to circulate efficiently among individuals and across generations. However, this obvious naturality conceals a deeper economic and sociological dimension. Money, being an intrinsically valueless asset, must be regarded as a social artefact whose value arises from collective belief and institutional trust, and must be analysed by economic means.

This thesis evaluates how societies coordinate to sustain belief in intrinsically valueless assets and why it is rational to sustain such belief. First, borrowing a sociological lens, this study briefly defines money as a "social contrivance" and explains how society moved from barter to the use of money.

The core of the work examines how rational coordination around intrinsically valueless assets can arise and persist in overlapping generations models, analysing three key contributions to economic theory: first, Samuelson's "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money", *Journal of Political Economy*, 1958, which demonstrates how fiat money can improve efficiency by enabling intertemporal transfers; second, Diamond's "National Debt in a Neoclassical Growth Model", *The American Economic Review*, 1965, which extends Samuelson's findings by allowing national debt to serve as a tool for intertemporal exchange that can sustain equilibria analogous to those in Samuelson's monetary model; and third, Tirole's "Asset Bubbles and Overlapping Generations", *Econometrica*, 1985, which shows that pure bubbles can similarly serve as stores of value in overlapping-generations economies, generalising Samuelson's framework to a broader class of speculative assets. The thesis then turns to a computational simulation of Samuelson's findings.

These classical models provide a theoretical foundation for understanding how belief in valueless assets can coordinate society and expectations, showing that Overlapping Generations models remain vital for understanding contemporary financial phenomena such as cryptocurrencies, including whether Bitcoin can be considered money as defined by the OLG framework.

Chapter 1 - From Barter to Institution: The Social Foundations of Money

1.1 The Problem of Value Without Worth

Societies assign value and purchasing power to assets that, in themselves, have neither intrinsic nor use value. Money, intended as fiat currency, possesses no inherent worth, yet it serves as the foundation of economic coordination, reducing friction in exchanges. This apparent contradiction that something devoid of material value can serve as the most essential medium of exchange raises a deeper question about the nature of social belief and coordination.

This chapter wants to introduce the central problem of collective belief in valueless assets, establishing the conceptual foundation of this thesis and providing the historical and sociological background for the analytical models examined later, Samuelson (1958), P. A. Diamond (1965), and Tirole (1985); these works, although grounded in formal economic modelling, share a common intuition, that the persistence of intrinsically valueless assets depends on a form of collective trust and coordination intertwined with pure individual optimisation.

Why do societies coordinate to accept and sustain belief in assets that lack intrinsic value, and why is it rational for individuals to participate in this collective fiction? The first part of the question pertains to the sociological dimension; the second concerns economics and will be developed in the core of this work.

This chapter will argue that money, and even debt and speculative bubbles, can all be understood as "social contrivances", far-fetched constructs that derive their function not from material backing but from collective acceptance and coordination. Through this lens, money becomes not merely an economic instrument but a social institution, one that embodies trust, predictability, and shared meaning.

The structure of the chapter follows this reasoning. The first section examines the sociological foundations of money. The second section traces the historical and theoretical shift from barter to transactions, highlighting how inefficiencies in barter gave rise to the need for a universally accepted medium of exchange. Together, these sections lay the groundwork for the formal economic analysis that follows in later chapters, where overlapping-generations models will illustrate how such beliefs can emerge and persist in equilibrium.

1.2 Money as a Social Contrivance

Economists traditionally define money through its functional triad: medium of exchange, store of value, and unit of account. A sociological lens exposes the underlying collective processes through which money becomes meaningful and operative. Money is not only an economic tool, but a social fact in Durkheim's sense: an institution that exists outside the individual, constraining and enabling behaviours through shared norms and collective belief; its validity is rooted not in material substance, but in the web of trust, obligation, and reciprocity that binds individuals into society.

In *The Philosophy of Money* (Simmel, 1900), Georg Simmel advanced the first systematic sociology of money. Simmel argued that money must be understood not through its material form but through its relational function, as the purest expression of social interaction and exchange; its worth does not come from intrinsic properties, but from its function as a mediator in social interactions (Deflem, 2003). For Simmel, the value of money is entirely social; money condenses the mutual dependencies of modern life. Simmel saw in money both alienation and liberation: while it depersonalises relations and commodifies life, it also frees individuals from the constraints of personal dependence, allowing them to participate in a wider network of exchanges (Deflem, 2003).

Émile Durkheim viewed money as part of the moral order of society; in his view, the power of money stems not from its economic function but from collective representations that endow it with legitimacy. Like other social facts, its force lies in the shared belief that compels individuals to treat it as real. For Durkheim, the force of money does not derive from its intrinsic value, but from the regulation to which it is submitted, and the nature of that regulation (Deflem, 2003). Money expresses the moral and normative structures that hold society together; it functions because society collectively agrees it should; without such normative backing, monetary systems dissolve into anomie.

Thus, money is a social institution: its stability depends on the density of social regulation and trust that sustains it. This insight anticipates the later notion that banking crises are not merely economic breakdowns, but crises of collective confidence (D. W. Diamond and Dybvig, 1983).

Max Weber viewed money as part of the broader process of rationalisation that defines modernity. In his analysis of capitalism, Weber embodies the principle of calculability in money: the transformation of value into measurable quantities that allow rational coordination across complex systems. Money thus supports the rise of

bureaucracy, contracts, and market institutions governed by impersonal rules rather than informal procedures. Money is the most abstract and impersonal element in human life, but it depends on shared faith in its validity (Deflem, 2003). The rational money economy therefore rests on deeply social foundations: belief in the credibility of promises, institutions, and legal enforcement.

Across these perspectives, a common line emerges: money's value is sustained by collective belief and institutional recognition. Money exists not only because agents trust others will accept it, but also because institutions enforce that trust. This belief is not irrational but socially rational, an equilibrium of expectations sustained by shared norms and credible institutions. In this sense, money can be described as a collective fiction with real effects: a symbol that coordinates human behaviour precisely because everyone believes in it, or to put it more simply, a social contrivance.

1.3 From Barter to Monetary Exchange

"In barter exchange there must be what Edgeworth (1881) called a double coincidence of wants. That is, not only does A have something that B wants, but it is also true that B has something that A wants, and both want what the other has more than they want what they themselves have, which they are willing to give up in exchange." (Coleman, 1990)

Barter represents the most primitive and limited form of exchange. Its inefficiencies are well known: transactions require a double coincidence of needs and wants, goods are usually indivisible, and exchanges cannot occur across time. These limitations make repeated trades across complex systems of agents improbable, and the trades that nevertheless occur are subject to friction and limitations. Without a common standard of value or a means of storing purchasing power, trade remains restricted to sporadic, bilateral interactions, severely limiting specialisation and the extent of the market. The emergence of a generally accepted medium of exchange thus arises as a social and economic necessity.

Historically, however, this transition did not occur through a simple technological improvement or an immediate transition from barter to coins. Anthropological evidence (Graeber, 2011) shows that early economies often revolved around webs of credit and social debt. Exchanges were centred on long-term relationships in which credit preceded money. Money, firstly in its token form, in Graeber's view, did not arise to replace barter but to record and standardise social trust, enabling exchanges among strangers where personal connections could not guarantee reciprocity.

Commodity monies, whether precious metals, shells, salt or other durable goods, emerged as a symbolic medium that condensed social trust into tangible form (Graeber, 2011). Commodity monies did not derive their legitimacy from intrinsic value but from collective belief in their acceptability, and some use value, as in the case of salt or precious metals.

Following Gary Gorton (2017), a safe asset must be accepted at face value and must be information insensitive, meaning that no questions are asked about the value of such an asset. Over time, this NQA property of coins has been doubted due to clipping, debasement, and counterfeiting. These problems with coinage paved the way for institutions to create money-like liabilities. The scarcity of coins and historical doubts about their credibility necessitated the creation of privately or publicly issued debt to serve as a store of value and a medium of exchange. Starting in the Middle Ages, the bill of exchange became the dominant form of this debt. A bill of exchange was a document by which the issuer authorised an agent to make a payment to a third party on a specified date. These bills could circulate, and if a holder required immediate funds, they could "discount" the bill with another party, receiving its discounted value early. To ensure the safety of the bills, a system of joint liability emerged: as a bill circulated, each party that endorsed it assumed liability for the final payment in the event of the originator's default. Over time, these bills evolved from simple payment instructions into substitutes for money, backed by a developing legal infrastructure and social cohesion that made them credible "safe assets".

By the 17th century in London, a new form of debt emerged when goldsmiths began issuing receipts for gold deposits in their safes. These receipts, circulated as early banknotes, derived their value from their convertibility into gold reserves and the credibility of the issuing goldsmith.

The problem with all these instruments was ensuring the holder's on-demand convertibility. To address this, debt was backed by collateral, such as precious metals for goldsmiths, real property for bills, or state bonds and loans for later banks. Short maturities and convertibility options allowed creditors to monitor whether issuers held sufficient backing. Additional safety mechanisms included unlimited liability, used by Scottish banks to assure creditors, and monitoring systems such as the Suffolk Bank system in the U.S., which prevented the depreciation of notes by threatening to redeem them if backing was insufficient. Section 6 of Chapter 3 provides a further discussion on the moneyness of public versus private debt.

Over time, as societies developed institutions capable of enforcing and stabilising

expectations, and complex legal structures, the material aspect of money lost significance, culminating in fiat money: a purely social construct whose value rests on the credibility of the social system rather than the substance it represents.

Yet even fiat money, initially issued as banknotes and coins, has undergone a further dematerialisation. In modern economies, most money no longer exists in physical form but rather as deposits recorded on bank balance sheets (McLeay, Radia, and Thomas, 2014). These deposits are not "money" in the classical, paper sense, but rather electronic entries representing claims on institutions, themselves backed by collective trust in the financial system and the state's regulatory authority. It is crucial to note that these deposits are liabilities of the bank and are supported by the bank's assets (primarily loans and securities). Consequently, trust in this form of money is not just a social abstraction but also relies on the quality of the backing assets, which are, in turn, promises to pay. Historically, when confidence in a bank's asset side vanishes, deposits lose their convertibility, triggering bank runs. In modern systems, this fragility is mitigated by deposit insurance and Central Bank supervision, which effectively substitute the state's balance sheet for that of the commercial bank in times of crisis.

The transformation from fiat to immaterial money intensifies the abstraction: value is no longer embodied in any physical object but in a shared belief in the liquidity and solvency of interconnected institutions. Nevertheless, it would be inaccurate to view modern fiat currency as entirely unbacked or purely a "shared fiction." While it lacks convertibility into precious metals, the monetary base is supported by the assets held by the Central Bank, historically government bonds and foreign reserves, and recently, through Quantitative Easing, a broader portfolio of private credit and long-term securities. Thus, the "belief" that sustains the system is not merely circular; it is sustained by the Central Bank's credibility and solvency. Nevertheless, the coordination aspect remains vital. Even with asset backing, the system functions because agents coordinate on a specific equilibrium. What individuals perceive as money is thus a network of promises, coordinated through technological and institutional infrastructures that sustain faith in convertibility, in the case of overnight deposits to banknotes. The system functions because no single agent questions its validity, because everyone assumes that everyone else will continue to accept these claims as money.

The evolution from barter to money is not merely a technological innovation that improves transactional efficiency; it is a social coordination achievement. Money embodies the collective capacity to converge on a shared fiction and sustain it over time. Money transforms the uncertainty of double coincidence of individual wants

into the stability of social expectation; the acceptance of money thus rests on a recursive belief: I accept it today because I trust that others will accept it tomorrow. This coordination is what sustains monetary systems and enables the emergence of complex markets and intertemporal exchanges.

In the overlapping-generations model of Samuelson (1958), this sociological insight finds formal expression. The model demonstrates that when agents collectively believe in the future acceptance of money, money can serve as a store of value, improving efficiency across generations. The theoretical result mirrors the historical evolution: belief, rather than substance, sustains value. Just as early societies coordinated around commodity symbols and institutional promises, Samuelson's framework demonstrates that rational agents can coordinate around a purely intrinsically valueless medium. This theoretical result is valuable for understanding assets such as cryptocurrencies, which, unlike sovereign currencies, lack the backing of a Central Bank. Samuelson's model proves that "social contrivance" alone can sustain value and that rational agents can coordinate around fiat, and even virtual money, rendering the social contrivance of money both possible and efficient.

Chapter 2 - The Rational Basis of Fiat Money

"But cannot men during their productive years give up some of their product to bribe other men to support them in their retirement years? Thus, forty-year-old A gives some of his product to twenty-year-old B, so that when A gets to be seventy-five he can receive some of the product that B is then producing." (Samuelson, 1958)

2.1 Why Samuelson Is Central to This Thesis


Samuelson's pivotal contribution, "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money", represents the first formal demonstration of how an intrinsically valueless asset, fiat money, can improve efficiency. By utilising the overlapping generations framework, Samuelson showed the central role of money in shifting society from a suboptimal equilibrium to the optimal "biological interest rate" equilibrium. He shows that, even in the absence of complex welfare institutions, money serves as an efficient clearing mechanism for intertemporal transfers, functioning both as a store of value and as a medium of exchange.

Samuelson established that in a world where goods cannot be stored (i.e., agents cannot save in physical form) and agents wish to smooth consumption across their life cycle, the only feasible option for saving is to lend to another generation. However, he demonstrated that competitive markets do not naturally create a functional intergenerational loan market in this scenario, resulting in interest rates collapsing into negative territory. Because no physical asset exists to transfer resources through time, the economy fails to reach an optimal allocation. To solve this, Samuelson argues that a special device is required: the "social contrivance" of money. Although money possesses no intrinsic value, it acts as a durable asset for intertemporal trade, provided its acceptability is sustained by social convention.

This foundational result paved the way for later contributions such as P. A. Diamond (1965) and Tirole (1985), who showed that other intrinsically valueless assets, respectively: government debt and pure bubbles, can play an analogous role in sustaining intergenerational transfers and improving welfare in overlapping generation economies.

2.2 The Model Setup

Samuelson constructs his economy based on a discrete-time overlapping generations (OLG) framework, designed to capture the frictions of human life cycles. In this stylised world, individuals live for three distinct periods. The first two periods



Each agent produces one unit of consumption good in each period of his productive life (period 1 and period 2) and produces none during retirement (period 3). Crucially, all goods produced are perishable, meaning that they are not suitable for consumption in the subsequent period (i.e., no physical way to store value from t to $t + 1$). In Samuelson's words: "no trade with Mother Nature". In the absence of durable goods, agents must resort to consumption loans to transfer consumption across periods.

Furthermore, individual preferences for consumption over time are homogeneous

within and between generations, with each agent maximising $U = U(C_1, C_2, C_3)$ with standard convex indifference curves. In this specification, Samuelson neutralises Böhm's causes of interest to isolate the consumption-smoothing motive. He assumes no systematic subjective time preference (Böhm's second cause), implying that agents do not necessarily discount future satisfaction and may even overvalue it. He also excludes the expectation of rising prosperity (the first cause), as agents face diminishing endowments in retirement. Finally, the perishability of goods eliminates the technological productivity of capital (the third cause), leaving an economy in which interest rates are determined solely by the interplay between biological growth and the desire to transfer purchasing power into retirement.

Defining R_t as the discount factor¹ of consumption between period t and $t + 1$, where t refers to the time: $R_t = \frac{1}{1+i_t}$, each agent seeks to maximise their lifetime utility under the following budget constraint: $C_1 + C_2 R_t + C_3 R_t R_{t+1} = 1 + 1R_t + 0R_t R_{t+1}$; hence the consumption decisions of agents are affected by the rate of discount between periods.

To avoid confusion between the longitudinal (life-cycle) and cross-sectional (calendar time) dimensions, the subscript $i = \{1, 2, 3\}$ refers to the stage of life(age) of the agent, and t to the calendar time, e.g. at $t = 5$ there would be: young agent A consuming $C_{1,5}^A$ born at $t = 5$ planning for consumption in 5, 6, 7, middle-aged agent B consuming $C_{2,5}^B$ born at $t = 4$ and old agent C consuming $C_{3,5}^C$ born at $t = 3$.

Defining net savings $S_{i,t}$ at time t of a representative agent in the i -th period of their life, as the difference between the endowment in that period and the chosen consumption level in that period $C_{i,t}$. In a period t , the economy comprises three generations, born respectively in $t, t - 1, t - 2$. Their respective savings functions are determined by the interest rates relevant to their specific lifespans:

$$\begin{aligned} \text{Young : } S_{1,t}(R_t, R_{t+1}) &= 1 - C_{1,t}(R_t, R_{t+1}) \\ \text{Middle-aged : } S_{2,t}(R_{t-1}, R_t) &= 1 - C_{2,t}(R_{t-1}, R_t) \\ \text{Old : } S_{3,t}(R_{t-2}, R_{t-1}) &= 0 - C_{3,t}(R_{t-2}, R_{t-1}) \end{aligned}$$

where $S_{1,t}$ represents the saving of a young agent, $S_{2,t}$ represents the saving of a middle-aged agent (which depends on the past rate R_{t-1} and current rate R_t), and $S_{3,t}$ represents the saving of an old agent, note that in $S_{3,t}$ we have past interest rates, meaning that old agent decisions no longer influence current or future interest rates.

¹Throughout, R_t denotes the discount factor on consumption between periods; $R > 1$ corresponds to a negative net interest rate. This follows Samuelson's own price-ratio notation

Since the savings of an individual across its lifetime must sum to zero, the individual budget constraint, related to savings, is given by: $S_1 + R_t S_2 + R_t R_{t+1} S_3 = 0$, meaning that: for any cohort, the present value of lifetime savings equals zero.

Given the only way of saving in this economy is by lending and that loans are nothing more than a promise of repayment, and assuming at time t , there exist B_i agents each aged i , the clearing condition for the lending market is that the sum of aggregated savings in every period must cancel out to zero, thus: $B_{1,t} S_{1,t} + B_{2,t} S_{2,t} + B_{3,t} S_{3,t} = 0$.

Samuelson identifies a fundamental mathematical indeterminacy in the overlapping generations model, noting that no unique equilibrium can be found within a finite time horizon because the number of unknown interest rates always exceeds the number of market-clearing equations.

2.3 The Stationary Population Case

In this scenario, to avoid the mathematical difficulty of "planning-until-infinity", Samuelson assumes that the population always remains constant with $B_{t-1} = B_t = B_{t+1}$, i.e., all the deaths are covered by births. Thus, at any time t there would be B agents of age one, B agents of age two and B agents of age three, with an aggregated output of $B + B$ in every period. In this stationary case Samuelson conjectures $R_{t-1} = R_t = R_{t+1} = R$.

Each agent seeks to maximise its lifetime utility $U(C_1, C_2, C_3)$ under its budget constraint (that arises from endowments), thus the representative agent faces the following lifetime budget constraint: $C_1 + RC_2 + R^2 C_3 = 1 + 1R + 0R^2$ defining savings as: $S_1 = 1 - C_1$; $S_2 = 1 - C_2$; $S_3 = 0 - C_3$ and substituting into the budget constraint, we get the budget identity: $S_1(R) + RS_2(R) + R^2 S_3(R) = 0$. Given R , the agent chooses an optimal consumption path (C_1, C_2, C_3) to maximise U .

In a stationary equilibrium, the aggregate demand for consumption loans must equal the aggregate supply in every period, and total net saving across all coexisting generations must sum to zero. Thus, the market clearing condition is $BS_1(R) + BS_2(R) + BS_3(R) = 0$, at $R = 1$, the market-clearing condition reduces to the individual budget identity, which holds by construction; hence $R = 1$ or $i = 0$ is a solution: in a stationary economy with no growth and no physical storage, every generation is identical, so the representative life-cycle budget identity ensures that the aggregate savings across all ages nets to zero. Samuelson is setting up for two later results: this "zero-interest" solution is not the only one; the model admits multiple equilibria.

The competitive market does not converge to the zero-interest equilibrium. In fact, the market converges to a negative-interest-rate equilibrium. This is the heart of Samuelson's paradox: the laissez-faire equilibrium is dynamically inefficient and socially inferior, unless a social contrivance is introduced.

2.4 The Growing Population Case

In the case of a geometrically changing population, growing or declining, at a rate of change m , where $B_t = (1 + m)B_{t-1}$, by design: when $m > 0$ the population is growing, when $m < 0$ the population is shrinking and when $m = 0$ the population is in the abovementioned stationary case.

Samuelson again conjectures a constant R across time, and taking $R = (1 + m)^{-1}$ as a candidate solution, now the market clearing condition becomes $B(1 + m)^t S_1(R, R) + B(1 + m)^{t-1} S_2(R, R) + B(1 + m)^{t-2} S_3(R, R) = 0$, recalling the individual budget constraint: $R = (1 + m)^{-1}$ or $i = m$ is a solution that clear the market.

"Every geometrically growing consumption-loan economy has an equilibrium market rate of interest exactly equal to its biological percentage growth rate."(Samuelson, 1958)

This theorem captures the core intuition of the social contrivance: the return on money is neither transcendental nor derived from physical production. Instead, it is rooted in the biological expansion of society. However, while this optimal equilibrium exists mathematically, a competitive market without money cannot achieve it. The social device of money is therefore the institutional bridge that allows society to achieve the most desirable outcome.

2.5 The Biological Interest Rate

After proving mathematically that the equilibrium market rate in a changing population economy is equal to the rate of population change ($i = m$), Samuelson proceeds to investigate the welfare implications of this result. He poses a critical question: does the allocation of resources resulting from this "biological" interest rate represent a social optimum?

To answer this, Samuelson constructs a "technocratic welfare problem" in which a "cross-sectional family or clan" seeks to maximise the lifetime utility of each agent, $U(C_1, C_2, C_3)$, subject only to the resource constraints of the model (i.e. one unit of output per productive period), without reference to interest rates.

In a stationary population ($m = 0$), the constraint is: the sum of consumption over a lifetime cannot exceed total production ($1 + 1$). However, in a growing population ($m > 0$), the demographic structure shifts. Because the population grows geometrically, the age distribution becomes skewed in favour of the young generation, i.e. there are always more young workers than pensioners. Specifically, for every old individual, there are $(1 + m)$ middle-aged individuals and $(1 + m)^2$ young.

Samuelson demonstrates that the per capita resources available to the representative man in such a society are governed by the constraint:

$$C_1 + \frac{1}{1 + m}C_2 + \frac{1}{(1 + m)^2}C_3 = 1 + \frac{1}{1 + m}$$

When the utility function is maximised subject to this constraint, the resulting first-order conditions are identical to those of a competitive market operating at an interest rate $i = m$.

This leads to the paper's fundamental "Optimality Theorem": the biological interest rate is both the equilibrium and the social optimum. It represents the most efficient allocation of resources across generations, maximising the "total social utility enjoyed over a long period of time". This finding is crucial for the thesis argument because it establishes the benchmark against which market failures can be measured. It implies that if a competitive market cannot naturally achieve $i = m$ (as Samuelson later proves it cannot without "social contrivances"), that market is inefficient.

Samuelson acknowledges that the result, which relates interest rates to human reproduction rather than to capital productivity, is "astonishing" and seeks a "common-sense" explanation. He dismisses earlier biological theories, and offers an explanation based on intergenerational bargaining power.

Samuelson proposes that the positive interest rate in a growing population stems from the changing ratio of workers to retirees. In a growing society, the "supply" of young workers looking to trade current goods for future claims is abundant relative to the "demand" from older generations seeking to secure their retirement. Samuelson theorises that this demographic imbalance confers a competitive advantage on older men, allowing them to extract a "bribe" or premium from the young, which manifests as a positive interest rate.

However, Samuelson distances himself from this intuition, noting that while it qualitatively explains why i might be positive for a growing population, it falls short of explaining the "remarkable quantitative identity" where i exactly equals m , em-

phasising that "cold and selfish competitive markets" do not operate on altruism or respect for the elderly; they operate strictly on supply and demand.

2.6 The Infinity Paradox: A Failure of Coordination

While the biological interest rate ($i = m$) is a valid mathematical solution and a social optimum, Samuelson evaluates whether the free market can reach it. The "common-sense" explanation that the middle-aged bribe the young to support them when old relies on the assumption that the young are overconsuming relative to their output (i.e., that they are net borrowers).

However, an analysis of individual optimisation reveals a fundamental contradiction. If we assume a standard symmetric utility function among all agents, where agents wish to smooth consumption evenly across their lives ($C_1 = C_2 = C_3$), the choices of the young generations prevent the market from clearing at the optimal rate, since there is a mismatch in demand and supply of loans.

In a stationary economy ($R = 1$) where a worker produces 1 unit in period 1 and 1 unit in period 2, and total lifetime income is 2. Perfect consumption smoothing implies consuming $\frac{2}{3}$ in each period. Consequently, the young agent produces 1 but consumes only $\frac{2}{3}$, resulting in positive net saving ($S_1 = +\frac{1}{3}$).

This creates a paradox of coordination: the middle-aged wish to save for retirement, and the young also wish to save for their future, while only the old wish to dissave. Although the aggregate savings of the three generations sum to zero, so that $R = 1$ is a mathematical solution to the market-clearing equation, this allocation cannot be implemented through voluntary bilateral trade: the only would-be borrowers are the old, who have no productive future from which to repay. The intergenerational chain of mutually advantageous exchanges required to sustain the biological interest rate cannot close. This "Infinity Paradox" suggests that the "invisible hand" of the market acts perversely, driving the economy away from the social optimum ($i = m$) toward a suboptimal equilibrium.

To prove the failure of the competitive market and to remark the limitations of voluntary exchange, Samuelson reduces the model to its most elementary form: a two-period life cycle consisting of one working period and one retirement period.

In a two-period model, a worker produces 1 unit in the first period and 0 in the second. Because goods are perishable, the agent has no physical means to carry purchasing power into the future. Furthermore, unlike the three-period case, there is no overlapping generation available for trade during the working period that can

credibly promise repayment in the retirement period.

Consequently, the economy collapses into autarky. Agents are forced to consume exactly what they produce in each period: $C_1 = 1$ and $C_2 = 0$. The saving rate is forced to zero ($S_1 = 0$), and the implied market interest rate crashes to $R = \infty$. Desperate to avoid starvation in retirement, agents would theoretically offer infinite amounts of current goods for any future return, yet no trade is possible.

This extreme case serves as proof of the general market failure. It demonstrates that in a pure consumption-loan world, the competitive equilibrium does not naturally converge toward the biological optimum. Instead, without a "social contrivance" to serve as a store of value, the market clears at a socially inefficient level, either resulting in autarky or, in the general n -period case, settling at a negative interest rate, with wealth eroding over time. This failure establishes the necessity of money not merely as a medium of exchange, but as a social compact required to rescue the economy from dynamic inefficiency.

To test if the overlapping nature of the three-period model allows the "Invisible Hand" to avoid this outcome, Samuelson drops the assumption of an eternal past and introduces a specific "beginning" to history. By analysing the recursive equilibrium equations starting from the first generation, the objective is to determine if the market dynamics will naturally converge over time toward the socially optimal biological interest rate or remain trapped in the inefficient autarkic solution.

One might conjecture that the friction in the two-period case is resolved in a three-period overlapping model through multilateral trade, in which the young support the old, anticipating support from a future generation rather than direct repayment. Samuelson, however, refutes this hope through his "Impossibility Theorem," proving that the market cannot sustain this chain of transfers.

For voluntary exchange to persist, every trade must be mutually advantageous. If a young agent A lends to an older agent B, expecting repayment from a future agent C, the chain of logic fails. Agent B cannot provide any benefit to C to induce this future payment, as B produces only in the past relative to C. Consequently, the intergenerational chain of mutually advantageous trades cannot close. Samuelson concludes that the joint pattern $S_1 > 0$ together with $R_t \rightarrow 1$ is logically impossible in a free market (absent special configurations of time preference), so the competitive economy cannot approach the biological optimum and instead settles at a suboptimal equilibrium with a negative interest rate.

To see why voluntary exchange cannot sustain the optimum, suppose history has a

beginning and consider a finite economy that ends at period T . The young in the final cohort have no successor generation to repay them; facing certain default, they will not lend, so their net savings are zero rather than positive. But then the young of period $T - 1$ cannot have lent in the expectation of repayment either, since the succeeding young, on whom their repayment ultimately depends, are now known not to reciprocate, so their lending plan collapses. The same reasoning propagates backwards, period by period, until it reaches the very first generation at $t = 0$. The conclusion is that $S_1 > 0$ cannot be sustained by voluntary bilateral trade in the limit of this truncation argument, and hence the competitive economy cannot approach $R_t \rightarrow 1$. It is worth being explicit that the work here is done by backward induction from a finite terminus, not by any direct property of the infinite economy. This distinction becomes the crux of later criticism (see Section 2.9).

Samuelson successfully demonstrated, using a numerical example (refer to the original text for the calculations), that in the stationary population case, the market interest rate would be negative rather than the biological zero interest rate corresponding to the social optimum, and similarly for a changing population where $m \neq 0$. The actual competitive market rate will always be negative and less than the biological optimal rate m .

The resulting equilibrium, characterised by negative interest rates, may initially seem counterintuitive. However, Samuelson argues that this result is consistent given the "extreme and purposely unrealistic assumptions" of the model. The construction explicitly neutralises the traditional drivers of positive interest identified by Böhm. Under these specific constraints, a negative market rate is not an anomaly but a logical necessity of the competitive equilibrium.

2.7 Why Pareto Efficiency Holds but Society Is Worse Off?

In the competitive solution (without money or social security), the market interest rate fails to reach the "biological optimum", as established by the previous reasoning. The market fails to reach the optimum because the young have no selfish motivation to voluntarily transfer goods to the old, as the old have nothing to trade in return (having already retired) and cannot interact with future generations to secure repayment. Despite this suboptimal outcome, Samuelson asserts that "free pricing gets you on the Pareto-efficiency frontier". The outcome is Pareto efficient because it is impossible to move to the "better" social optimum without making at least one group (the final generation) worse off.

The superior "social optimum" relies on a continuous chain in which the young sup-

port the old, expecting to be supported by the next generation. If humanity were to end at a finite time, the final generation of young people would support their elders but receive nothing in return when they retire. If society adopted the social optimum, the very last generation would inevitably suffer a net loss, as they would have supported their elders but would have no successors left to provide for them in return. Because implementing the better system would harm this final generation (relative to the competitive isolation where they keep their own goods), the competitive solution satisfies the strict definition of Pareto efficiency. As Samuelson notes, "Perfect competition shrugs its shoulders at such a question and (not improperly) sticks to its Pareto optimality".

This Pareto efficiency, however, is a property of each finite truncation, it is the inability to compensate a last generation that blocks the improvement. In the infinite-horizon economy, there is no terminal cohort; the negative-interest equilibrium is dynamically inefficient; and the welfare gain from money established in §1.8 is precisely the removal of that inefficiency. Because the free market traps society in a Pareto-efficient but socially inferior position, Samuelson argues that society requires "social collusions" or "social compacts" to reach the optimum.

2.8 The Social Contrivance of Money

The "Impossibility Theorem" establishes that a pure consumption-loan market fails to reach the biological optimum, and society remains trapped in a suboptimal equilibrium in which every generation is worse off than it would be under the biological interest-rate configuration. To escape this deadlock, Samuelson argues that society must transcend the limitations of bilateral trade through a social contract.

This social compact fundamentally alters the rules of the game. It replaces the immediate necessity of a counterparty with a collective intertemporal guarantee. Under this agreement, the productive young agree to provide subsistence to the generation of retirees. In exchange, they do not receive an immediate material payment, but rather a binding assurance that they, in turn, will be supported in their own retirement. Crucially, this support is guaranteed by a "draft on the yet-unborn". By creating a claim against future generations, the social contract bypasses the constraints of the two-period or finite-horizon models, allowing the social optimum to be achieved within a single lifetime. Samuelson emphasised that this contract cannot rest solely on the good faith of agents and requires a mechanism of enforcement: society must enforce the contract politically, using force to transform an unstable equilibrium into a stable one, ensuring that the draft on the unborn is honoured when it comes due.

Samuelson operationalises this compact by introducing a physical instrument: money. Since, by design in this model and in the real world, consumption goods perish, individuals cannot store value to transfer purchasing power from their working years to their retirement. To circumvent this biological reality, Samuelson posits that men, either officially through the state or unofficially through custom, can resort to "printing oblongs of paper or stamping circles of shell".

These tokens introduce a crucial innovation to the economy: unlike the consumption goods, money can be spent across time. Samuelson is careful to distinguish this "social contrivance" from commodity money or durable capital. He insists that these new assets must be "themselves quite worthless for consumption". Their value does not derive from intrinsic utility, but solely from their value in exchange. Money thus acts as a pure store of wealth, bridging the temporal gap between production and consumption that perishable goods cannot span.

The introduction of this durable medium transforms the mechanism of saving. In the absence of money, saving was impossible or yielded negative returns. With the social contrivance, the productive young and middle-aged agents can now invest their savings in currency. When they retire, they disinvest in currency, transferring it to the new generation of productive workers in exchange for consumption. As long as the new generation accepts the currency, believing they, too, can use it for their future retirement, the "pieces of paper" function successfully as a claim on future labour. Thus, without the need for elaborate legislative machinery, society can utilise money to move from the "non-optimal negative-interest-rate configuration" to the socially optimal configuration.

For analytical tractability and in line with Ljungqvist and Sargent's standard treatment, we now adopt a two-period formulation of Samuelson's OLG structure. By introducing money, agents can now allocate their endowment between consumption and real money balances, let p_t denote the price level of consumption at time t in terms of money, let f_t represents the nominal money balances acquired by a young agent born at time t and carried into old age., and $C_{i,t}$ where the subscript t represents the time period, and the subscript $i = \{1, 2\}$ define the age of the agent, let F be the fixed aggregate stock of money endowed to the initial old generation. Two sequential budget constraints emerge at time t :

$$C_{1,t} + \frac{f_t}{p_t} \leq y_{1,t} \quad (\text{Young at time } t)$$

$$C_{2,t+1} \leq \frac{f_t}{p_{t+1}} + y_{2,t+1} \quad (\text{Old at time } t+1)$$

Combining these two inequalities (by substituting for f_t) yields the following life-time budget constraint:

$$C_{1,t} + \left(\frac{p_{t+1}}{p_t} \right) C_{2,t+1} \leq y_{1,t} + \left(\frac{p_{t+1}}{p_t} \right) y_{2,t+1}$$

Let $\frac{p_{t+1}}{p_t} = \frac{1}{G_t}$ where G_t represents the gross real return on money. If prices are falling (deflation), money gains value over time, providing a positive real return ($G_t > 1$). If prices are constant, the return is unity ($G_t = 1$), corresponding to a zero net interest rate ($i = 0$).

Following Ljungqvist and Sargent (2018), we must distinguish between a monetary equilibrium and a non-monetary one. A monetary equilibrium requires that money has a positive exchange value in all periods. An equilibrium with valued fiat money is composed of a feasible allocation sequence for a representative agent $\{C_{1,t}, C_{2,t}\}_{t=1}^{\infty}$ and a sequence of nominal prices $\{p_t\}_{t=1}^{\infty}$ such that:

1. Prices are finite, $0 < p_t < \infty$.
2. Given the price sequence, the allocation maximises the utility of the representative agent subject to the budget constraints.
3. The demand for real balances by the young equals the real supply of money held by the old. In a stationary setting, with F fixed,

$$\frac{F}{p_t} = y_{1,t}^{agg} - C_{1,t}^{agg}$$

where $y_{1,t}^{agg} - C_{1,t}^{agg}$ represent the aggregated real savings of the young in period t .

By contrast, the non-monetary equilibrium is defined by $p_t = +\infty$ (equivalently, $1/p_t = 0$) for all t , so that real balances are identically zero and no intertemporal trade occurs. The allocation in this case coincides with autarky: $c_{1,t} = y_{1,t}$ and $c_{2,t+1} = y_{2,t+1}$. We can now derive the equilibrium for the case of a stationary population ($m = 0$) and a fixed money supply F . In the stationary case, consumption and real balances are constant:

$$C_{1,t} = C_y \quad C_{2,t} = C_o \quad \frac{F}{p_t} = \frac{F}{p_{t+1}}$$

for real balances F/p_t to be constant while the money stock F is fixed, the price level must be constant: $p_t = p_{t+1} = p^*$. Substituting this into the rate of return expression derived from the budget constraint: $G = \frac{p_t}{p_{t+1}} = 1$; consequently, $i = 0$.

The intertemporal budget constraint simplifies to:

$$C_y + C_o = y_y + y_o$$

This equation implies that the agent can transfer purchasing power from youth to old age at a one-to-one ratio. This corresponds exactly to Samuelson's "biological optimum" for a stationary population.

The market clearing condition for money in this stationary state becomes:

$$\frac{F}{p^*} = y_y^{agg} - C_y^{agg}$$

where $y_y^{agg} - C_y^{agg}$ correspond to the savings of young when the gross rate of return is 1. As long as the young desire to save (i.e., $y_y^{agg} - C_y^{agg} > 0$) at zero interest, there exists a valid monetary equilibrium price p^* that sustains this social optimum.

The introduction of money does more than simply provide a storage mechanism; it aligns the market interest rate with the social optimum. Samuelson demonstrates that under a fixed money supply, the natural adjustment of the price level creates a real rate of return on currency that exactly mirrors the biological growth rate of the population.

In the case of a stationary population where $m = 0$, if the total stock of money F is held fixed, the price level will stabilise. In this environment, a unit of currency saved during the working years retains its purchasing power upon retirement. Consequently, the real interest rate on money savings is zero $i = 0$, which corresponds precisely to the biological optimum for a stationary society.

The mechanism becomes even more elegant when applied to a growing population ($m > 0$). As the population grows geometrically, the economy's total output grows at the same rate. With a fixed supply of money chasing an ever-expanding supply of goods, the price level must decline, falling proportionally to $\frac{1}{(1+m)^t}$. This deflationary trend implies that money gains purchasing power over time, yielding a real interest rate exactly equal to m per period.

Thus, the "social contrivance" of money automates the optimal allocation. Without any central planning or explicit legislation, the monetary mechanism ensures that the return on savings equals the rate of population growth. Even in the case of a declining population ($m < 0$), the mechanism holds: a shrinking output, interacting with a fixed money supply, results in rising prices, thereby generating the negative real interest rate needed to match the negative biological growth rate. In all cases,

money serves as the vehicle that moves the economy from the inefficient competitive equilibrium to the biological social-optimality configuration.

The stability of this monetary equilibrium relies solely on a "grand consensus," established by state decree or custom, rather than any intrinsic value. Samuelson identifies the inherent risk of this arrangement: based on "immediate self-interest," the productive young could theoretically "unilaterally repudiate" the worthless money held by the aged to maximise their own consumption.

However, the system persists because the young are trapped in the same intertemporal bind; they accept the currency today only because they expect the next generation to honour it tomorrow. Thus, money effectively becomes the "social compact" itself. This self-enforcing equilibrium transforms the "social contrivance" into a functioning store of value, securing the biological optimum through a shared belief in future acceptance.

2.9 Critiques of Samuelson

While Samuelson's contribution provided a groundbreaking rationale for fiat money, it faced profound theoretical challenges from the high theory of General Equilibrium and subsequent monetary schools. The first major critique argues that the "market failure" Samuelson identified, which justified the introduction of money, was a mathematical artefact of the "double infinity" inherent in the model: the presence of an infinite number of time periods combined with an infinite stream of distinct agents. Karl Shell (1971) shows that the "double infinity" of agents and dated commodities in an OLG setting creates an indeterminacy at infinity, often described in later literature as a "missing market at infinity". In standard finite economies, equilibrium price systems are pinned down by terminal boundary conditions; in OLG, the tail remains untethered. The fiat-money "social contrivance" exploits exactly this slack, and Tirole (1985) later shows that such monetary equilibria are rational bubbles, sustainable precisely when the interest rate does not exceed the growth rate.

A second fundamental critique concerned the asset characteristics of money, formalised by Neil Wallace (1980) as the "Rate of Return Dominance" problem. Wallace accepted the OLG framework's premise of "intrinsic uselessness". However, he demonstrated a fragility in Samuelson's construction: because money in this model functions strictly as a vehicle for intertemporal saving, it must compete with other assets. Wallace proved that in any economy with a physical storage technology (capital) that yields a return r greater than the population growth rate m , fiat money cannot exist because rational agents will always prefer the higher-yielding asset.

This implies that Samuelson's theory of money is valid only in dynamically inefficient economies ($r < m$) or in those that lack productive storage entirely, thereby limiting its empirical relevance in modern economies where non-interest-bearing currency coexists with high-yield capital. More fundamentally, Wallace's point is one of indeterminacy rather than mere dominance. Once a storage technology with the same payoff structure as money is admitted, the OLG model contains nothing to pin down why agents would hold the dominated asset at all; the framework cannot endogenously select money for circulation alongside capital.

The third major line of criticism targets the functional definition of money, which originates from Kiyotaki and Wright (1989). They argued that OLG models focus on the "store of value" function while ignoring the "medium of exchange" function. In the Consumption-Loan model, money solves a timing problem (transferring purchasing power to the future) rather than an exchange friction problem; it is used only twice in an agent's life. Kiyotaki and Wright posited that a true theory of money arises from "search frictions" (the classic double coincidence of wants problem) rather than intertemporal smoothing. Samuelson abstracted away the very coordination and trust problems that make money essential; on this view, Samuelson's model treats money as a store of value or a public debt instrument, but abstracts from the coordination and acceptance frictions that make it essential as a medium of exchange.

The model has also been critiqued for its structural instability and the prevalence of indeterminacy. Unlike standard representative-agent models that typically yield unique, stable equilibria, the OLG structure admits a continuum of equilibria.

A final critique targets the stability of the "social contrivance" itself, challenging the plausibility of the cooperative fiction that sustains the monetary equilibrium. Game theorists argue that Samuelson's model relies on a precarious suspension of disbelief regarding the "end of the world". This critique is best understood through the logic of backward induction within a finite horizon framework. If one assumes a known, finite end to the economy at time T , even if distant, the rational expectations equilibrium collapses. The final generation of young agents at time T , knowing there is no generation $T + 1$ to whom they can sell money, will assign a value of zero to the currency. Rational agents at time $T - 1$, anticipating that money will be worthless in the subsequent period, will refuse to accept it. This logic cascades instantly backwards through time, causing the value of money to collapse to zero at $t = 0$. Samuelson avoids this problem by assuming a strictly infinite horizon, so that no final period exists, and the backward induction argument cannot be applied. Critics interpret this as showing that the value of fiat money in the model is entirely

belief-dependent: its role as a "social contrivance" is sustained only so long as each generation expects future generations to continue accepting it. If, at some finite date, it became common knowledge that the currency would be repudiated, the monetary equilibrium would collapse immediately, not just at the announced terminal date. In my own reading, it is precisely the role of credible institutions that back the fiat currency to dissipate this threat to the equilibrium.

Nonetheless, Samuelson provided that in an economy lacking durable goods, fiat money serves as an essential bridge for intertemporal trade, allowing the economy to reach the biological optimum. Yet, as P. A. Diamond (1965) observes what happens in a world where physical production and capital accumulation are feasible. By modelling neoclassical production within the OLG framework, Diamond creates a world where agents can save for retirement not just by buying money but also by lending to the government or entrepreneurs. In this context, the thesis turns to the role of debt.

Chapter 3 - Claims on the Future: Debt as a Store of Value

3.1 The Importance of Diamond's Extension

In "National Debt in a Neoclassical Growth Model", Diamond not only shows the impact of national debt on the long-run equilibrium and on the growth of an economy, but he also demonstrates that public debt facilitates intertemporal transfers. By designing a system where state debt is continuously refinanced and effectively floats indefinitely, Diamond establishes debt as a viable store of value that competes with capital. However, this shifts the institutional nature of the asset: Samuelson's money relies on "social compact" or custom (which may or may not be enforced by the state); Diamond's debt is anchored by the state's fiscal authority. While this provides an institutional backing, it introduces the role of taxation.

Unlike Samuelson's money, Diamond's debt enters an environment already populated by physical capital. This setting allows investigating whether government debt acts as a substitute for physical capital, crowding it out, and how this substitution affects the economy's Golden Rule path.

In Diamond's extension, government debt serves as a store of value by acting as a saving mechanism and configuring as an alternative to physical capital. An individual faces a portfolio choice: they may allocate their savings in physical capital or government debt. Since both assets are assumed to pay the competitive interest rate, they are perfect substitutes in the individual's budget constraint. Consequently, the issuance of debt does not merely facilitate intertemporal transfer as in Samuelson's pure exchange economy; rather, it induces a substitution effect where 'paper' wealth displaces physical capital accumulation. This confirms that debt operates as a store of value competing with productive investment.

3.2 Diamond's OLG Model

Diamond's economy has an infinite time horizon and is characterised by agents who live for two periods, working in the first (age 1) and retiring in the second (age 2). In such an economy, agents must save for retirement; their consumption-smoothing preferences are described by the utility function $U(c_1, c_2)$. The output of one's work is both consumable and investable, meaning that output and capital are the same commodity. The population changes geometrically: $L_t = L_0(1 + n)^t$, where L_t is the population born in the t -period, L_0 is the starting population, and n is the exogenous rate of population change.

The technology of the economy is defined by $F(K, L)$, which exhibits constant returns to scale, is twice differentiable and has positive marginal products and dimin-

ishing marginal rate of substitution everywhere. Output in period t is thus defined as $Y_t = F(K_t, L_t)$, where K_t is the capital saved in $t - 1$ to be employed for production in t , capital does not depreciate and can be rolled over for production, and L_t is the t -period workforce. Diamond's economy is a "100% equity" economy, meaning that: $S_t = K_{t+1}$.

As a baseline for the model, Diamond uses a centrally planned economy, in which the social planner controls the capital stock K_t and the distribution of output Y_t in each period. The planner sets aggregate consumption C_t and thus consumption per worker $c_t = \frac{C_t}{L_t}$; and capital in period $t + 1$ (or savings per worker in period t) $k_{t+1} = \frac{K_{t+1}}{L_{t+1}}$ where $K_{t+1} = Y_t + K_t - C_t$. Assuming a fixed consumption for worker based on their generation, we have that: $C_t = c_t^1 L_t + c_t^2 L_{t-1}$. Formalising:

$$Y_t + K_t = K_{t+1} + C_t = K_{t+1} + c_t^1 L_t + c_t^2 L_{t-1}$$

The central planner, may decide to keep a constant output per worker: $y_{t-1} = y_t = \frac{Y_t}{L_t} = y_{t+1} = y$, to do so would require a fixed capital per worker, and the planner must account for the population change, meaning: $k_t = \frac{K_t}{L_t}$ and $K_{t+1} = (1+n)K_t$ so that $k_{t-1} = k_t = k_{t+1} = k$. Formalising the resource constraint faced by the central planner:

$$c^1 + \frac{c^2}{1+n} = y - nk$$

where nk is the break-even investment (savings per worker), necessary to maintain a constant y by providing capital to the new generation of workers the next productive period (assuming a growing population, $n > 0$). When the social planner sets a fixed capital-labour ratio and, consequently, a fixed output-labour ratio, the economy enters a Golden Age path.

The central planner may choose any k and consequently any y (n.b. k must be set so that output-capital is greater than n , so that the saving rate does not exceed 1). Assuming the benevolence of the central planner, meaning that it seeks to maximise the lifetime utility of each agent, the central planner seeks to select the optimal k so that for a representative agent: $U(c_1, c_2)$ is maximised subject to $c_1 + \frac{c_2}{1+n} = y - kn$.

The maximisation decomposes into two problems: that of setting an optimal capital intensity k^* . This yields the Golden Rule condition where the marginal product of capital equals the rate of population growth: $f'(k) = n$, where f is the per-worker output function. The optimality of this condition is independent of how the output is divided between agents. At k^* , as illustrated in Figure 2, aggregate consumption is maximised.

"Thus the solution of the problem of selecting an optimal Golden Age Path treats the allocation of consumption over the lifetime of an individual in a similar fashion to the allocation of consumption, in a single year, between individuals of different ages." (P. A. Diamond, 1965)

Secondly, aggregate consumption must be allocated between the young (c_1) and the elderly (c_2). Individual utility maximisation requires: $\frac{\partial U}{\partial c_1} = (1+n)\frac{\partial U}{\partial c_2}$, implying that the rate of interest used in consumer decisions equals the rate of population growth n , which corresponds to Samuelson's biological interest rate. Thus, in the centrally planned optimum, the implicit interest rate is determined entirely by demographic growth (n). As Diamond notes, while the planner forces the marginal product of capital to equal n , a competitive market offers no such guarantee, allowing the interest rate to diverge from the growth rate.




Figure 2: Selection of the Golden Rule Path

Figure 2 illustrates the maximisation process faced by the central planner, in which consumption is maximised and the marginal product of capital equals the rate of population growth, which coincides with the natural interest rate. c^* represents the consumption per worker at the optimum and nk^* the break even investment (savings) per worker. Any deviation from k^* results in suboptimal consumption. At a lower capital stock (k_L), the economy is under-capitalised; while the marginal product of capital is high, total output is insufficient to support maximum consumption ($c_L < c^*$). Conversely, at a higher capital stock (k_H), the economy is dynamically inefficient. Although output is higher, a disproportionately large share of resources must be reinvested (nk) simply to maintain the capital stock, crowding out consumption ($c_H < c^*$).

Having established the centrally planned benchmark, one must return to the market

context, where the existence of an "invisible hand" that steers markets towards the optimum cannot be assumed. While the planner's solution provides an upper bound for consumption, it relies on the strong assumption of a "benevolent omniscient dictator" capable of aggregating and processing all societal preferences. As Hayek notes in *Freedom and the Economic System*, centrally planned efficiency requires a consensus on a single hierarchy of ends, a consensus that may be absent in a free society (F. A. v. Hayek, 1939).

Consequently, the competitive economy does not rely on a single authority's will but on market processes to determine savings in each period and, thus, the capital available for production in the next period. Diamond focuses on assessing the alternative Golden Age paths to which the economy converges, with different levels of outstanding national debt.

3.3 Competitive Equilibrium and the Possibility of Failure

Observing an agent born at time t that works in time t and receives wage w_t equal to the marginal product of labour, the agent allocates the wage between consumption in the current period (c_t) and consumption during retirement (c_{t+1}) maximising its lifetime utility. To smooth its consumption across periods, the agent can save by lending to entrepreneurs (capital demanders). Savings between period t and $t+1$ are defined as $s_t = w_t - c_t$, and pay the competitive interest rate, which is determined by the marginal product of capital, $r_{t+1} = f'(k_{t+1})$. Thus, the individual budget constraint takes the form of:

$$c_t + \frac{c_{t+1}}{1 + r_{t+1}} = w_t$$

Hence, equivalently to the social planner case, consumption will be allocated so that:

$$\frac{\partial U}{\partial c_1} = (1 + r) \frac{\partial U}{\partial c_2}$$

By analysing a single agent, the relevant market relations can be traced out. Since the technology function exhibits constant returns to scale ($Y = Lf(k)$) there exist a relationship between the interest rate (cost of borrowing capital) and the wage earned by workers (residual output after capital is compensated), since $r = f'(k)$ and $w = f(k) - kf'(k)$, Diamond encloses this relationship (factor-price frontier) into: $w = \phi(r)$. Necessarily for a given technology higher interest rate (more expensive capital) necessarily implies a lower wage for a given technology ($\frac{dw}{dr} = \phi'(r) = -k$). Therefore, savings can be expressed as a function of wage and interest level $s_t = s(w_t, r_{t+1})$. Since consumption in every period is a normal good, $0 < \frac{\partial s}{\partial w} < 1$, and

due to income and substitution effects, the sign of $\frac{\partial s}{\partial r}$ is unclear. From the reasoning above, the utility function has:

$$\frac{\partial U}{\partial w} = \frac{\partial U}{\partial c_1} \quad \text{and} \quad \frac{\partial U}{\partial r} = \frac{s}{(1+r)} \frac{\partial U}{\partial c_1}$$

The supply side of the capital market is thus composed of the aggregate supply of savings, which can be modelled as $S_t = s_t L_t = L_t s(w_t, r_{t+1})$. The demand side is made up of entrepreneurs demanding capital for production in the next period and willing to pay the marginal product of capital, $r_{t+1} = f'(K_{t+1}/L_{t+1})$.

Equating demand and supply yields the equilibrium conditions for clearing the capital and savings market. Set $S_t = K_{t+1}$, since the capital market clears when today's savings equal the total capital stock in the next period; expanding savings $L_t s(w_t, r_{t+1}) = K_{t+1}$, dividing by L_{t+1} we get:

$$\underbrace{\frac{L_t}{L_{t+1}}}_{1/(1+n)} s(w_t, r_{t+1}) = \underbrace{\frac{K_{t+1}}{L_{t+1}}}_{k_{t+1}}$$

substituting and rearranging:

$$r_{t+1} = f' \left(\frac{s(w_t, r_{t+1})}{1+n} \right)$$

The equilibrium condition in the capital market relates the interest rate to the wage rate of the previous period; r_{t+1} inside the savings function $s(w_t, r_{t+1})$ represents the anticipated return that the young generation expects to earn when they retire. The r_{t+1} on the left side is the actual marginal product of capital that entrepreneurs will pay tomorrow. Diamond assumes agents have perfect foresight; the rate they expect must exactly equal the rate that actually clears the market. Thus, the young decide how much to save based on tomorrow's interest rate, and tomorrow's interest rate is determined by how much the young saved.

By design, the demand curve is downward-sloping, but because the elasticity of savings with respect to the interest rate that can be positive or negative, the slope of the supply curve is unclear. Consequently, when an increase in wages shifts the savings supply curve to the right, the new equilibrium savings depend on the relative slopes of the demand and supply curves. Diamond identifies two possibilities: a "normal" case where higher income yields higher equilibrium savings, and a 'perverse' case where a strong negative interest elasticity of savings causes a rise in income to paradoxically lower equilibrium savings. For the purposes of this analysis, I will restrict

the focus to the "normal" case. This restriction is justified because it aligns with standard macroeconomic intuition regarding income effects, but, as Diamond notes, the "perverse" scenario can be eliminated under the assumption of Walrasian stability in the capital market; this ensures that the mechanics of debt's impact remain clear without unnecessary complications.




Figure 3: Equilibrium and supply shift

As shown by Figure 3, an increase in wage ($w' > w$) causes the supply curve to shift to the right; thus, a higher income level translates to an increase in savings and consequently a lower interest rate. Altering the wage in period t changes the interest rate in $t + 1$; this relation is defined as: $r_{t+1} = \psi(w_t)$.

$\phi(r)$ and $\psi(w_t)$ interact to map the evolution of the economy across generations. The function $w = \phi(r_t)$ represents the factor-price frontier, which defines the inverse relationship between the wage rate and the interest rate in a single period, given the firm's production technology. Meanwhile, the function $r_{t+1} = \psi(w_t)$ represents the capital market equilibrium, linking the current generation's wage to the subsequent period's interest rate through their savings behaviour. Together, they trace out the entire time path of the economy toward its long-run equilibrium.




Figure 4: Time path of the competitive solution

Starting from I where wage equals w_1 and the interest rate equals r_1 , ψ gives the period two interest rate (r_2) and ϕ of r_2 gives w_2 . Clearly, as time progresses, the economy converges towards E where $w = w_E$ and r_E , and ϕ and ψ intersect. To derive the unique equilibrium we find r_{t+1} as a function of r_t : $r_{t+1} = \psi(\phi(r_t))$.

To derive the stability condition, since:

$$\frac{dr_{t+1}}{dw_t} = \psi' = \frac{f'' \frac{\partial s}{\partial w}}{1 + n - f'' \frac{\partial s}{\partial r}} < 0$$

and $\phi' = -k$ from the above discussion, the necessary stability condition for the equilibrium is:

$$0 < \frac{dr_{t+1}}{dr_t} = \phi' \psi' = \frac{-k f'' \frac{\partial s}{\partial w}}{1 + n - f'' \frac{\partial s}{\partial r}} \leq 1$$

As previously mentioned, in a competitive framework the economy may not converge toward what is socially desirable and may settle in a state of dynamic inefficiency.

Diamond demonstrates, using an example with Cobb-Douglas production and utility functions, that a competitive economy can converge to a long-run equilibrium in which the interest rate (r_E) is either greater than or less than the population growth rate (n). The case where $r_E > n$ corresponds to dynamic efficiency: because the interest rate equals the marginal product of capital, such equilibrium implies that the economy has settled at a capital-labour ratio that is strictly below the Golden Rule level. While this state might be viewed as an "under-accumulation" of capital, because aggregate steady-state consumption is not at its theoretical maximum, it

is Pareto efficient. In this case, any reallocation that raises the consumption of the currently living generations necessarily requires a reduction in savings. Since savings in period t determine the capital stock k_{t+1} inherited by the next cohort, such a reduction drops the capital-labour ratio below its steady-state level, strictly reducing the output and consumption available to at least one future generation. No Pareto-improving reallocation therefore exists, and the allocation, despite falling short of the Golden Rule maximum of steady-state consumption, is dynamically efficient.

When $r_E < n$ the economy has over-accumulated capital to the point that the return on that capital is lower than the biological growth rate of the economy; society is saving so much that a disproportionate amount of output is wasted equipping the growing population with machines, rather than being consumed. Because capital and output are the same commodity, a central planner could discard (or consume) capital to lower the capital-labour ratio down to the Golden Rule level; this would permit a higher level of consumption in each period forever, without hurting any future generation; thus the economy is in a state of dynamic inefficiency.

3.4 Debt as Competing Assets and Its Dual Nature

Having established the benchmark for a competitive setup, Diamond turns to analysing the impact of debt on the economy's long-run path. To evaluate the effects of national debt, Diamond distinguishes between balanced-budget incidence and differential incidence. Balanced-budget incidence evaluates the effects of changes in both government expenditures and their financing methods; differential incidence isolates the impact of financing mechanisms for a fixed level of expenditure. To cleanly isolate the long-run impact of the debt itself, Diamond restricts his analysis entirely to differential incidence, specifically examining the substitution of debt for tax finance. By assuming the government issues debt to finance a temporary windfall payment, the example employed by Diamond is a veteran's bonus; issuance of such a bonus has no permanent effect on the economy's long-run equilibrium. Consequently, by comparing the original equilibrium to one where this debt exists, any resulting shifts in the economy's steady state can be attributed purely to the introduction of the debt.

To introduce national debt in its model, Diamond assumes a 1-period maturity to avoid the problem of expected capital gains, and that debt that refloats and needs to be rolled over pays the current interest rate; for internal debt, this assumption is crucial since, under perfect certainty, for agents to hold both debt and capital in their portfolio, debt must pay the same return as capital. If there is no risk, any yield

difference would cause agents to hold only the higher-yielding asset. This assumption also holds for external debt; reasonably, why would a foreign investor buy debt if domestic capital pays better? Thus, to ease the comparison between internal and external debt, Diamond sets their returns as equal. Diamond acknowledges that a separate supply curve for foreign capital would be a "more natural assumption" (explored further in the appendices of his work), but maintains the symmetry to cleanly isolate the differential incidence.

A fixed amount of debt on an indefinitely growing population would have no long-run effect; therefore, Diamond assumes the debt-labour ratio is held constant over time, so debt grows symmetrically to the population at a rate n . The government maintains this by financing a portion of the interest costs through the issuance of additional debt, while financing the remainder through taxes levied as a lump sum on workers. Hence, the total stock of debt in the market can be expressed as:

$$G_t = gL_0(1 + n)^t$$

Where G_t is the debt at the beginning of production period t , g is the constant per-worker debt ratio, and L_0 is the tax base (labour force) "at the beginning of time". Note that when the economy is dynamically inefficient, the interest cost is less than the rollover; this setup implies negative taxes, i.e. net transfers to the younger generation.

In what follows, g_1 and g_2 denote the per-worker ratios of externally and internally held debt, respectively.

3.4.1 External Debt - The Benchmark

The effects of externally held debt on the domestic economy operate exclusively through the fiscal channel. Since foreign creditors hold the debt, the only domestic consequence is the tax the government must levy to sustain it. There is no alteration of the demand side of the capital market; foreign savers, not domestic households, absorb the debt, so the factor-price frontier $\phi(r)$ remains entirely unchanged. This separation is what makes external debt the benchmark: it isolates a single transmission mechanism before the more complex internal-debt case is introduced.

With an external debt-labour ratio g_1 , the government must service the outstanding stock each period. Since the debt-labour ratio is held constant and debt pays interest at the current rate r_t , the total interest due per worker is $r_t g_1$. Of this, a fraction corresponding to the population growth rate n is covered by rolling over new debt; the government issues ng_1 in new debt per worker to keep the debt-labour ratio

constant. The residual must be financed by taxation. The lump-sum tax levied on each young worker in period t is therefore: $\tau_t = (r_t - n)g_1$.

The sign of τ_t governs the direction of the subsequent results in this section. In the normal or efficient case, where the equilibrium interest rate exceeds the growth rate ($r > n$), the tax is strictly positive: workers remit a net payment to foreign creditors. At the Golden Rule ($r = n$), the rollover exactly covers the entire interest bill and no tax is needed. In the inefficient case ($r < n$), the rollover revenue exceeds the interest cost, so τ_t turns negative; the government distributes a net transfer to the young generation rather than extracting one. This reversal is not a modelling anomaly; it is a consequence of the debt-labour ratio convention.

Since taxes are settled concurrently with the receipt of factor payments, the young worker's disposable wage is $\hat{w}_t = w_t - (r_t - n)g_1$; now the individual budget constraint becomes $c_1 + c_2/(1 + r_{t+1}) = \hat{w}_t$, thus the individual savings function becomes:

$$s_t = s(\hat{w}_t, r_{t+1}) = s(w_t - (r_t - n)g_1, r_{t+1})$$

The normality assumption on consumption, $0 < \partial s / \partial w < 1$, implies that a reduction in disposable income reduces saving. This is the link between the tax and the capital market: in the efficient case, positive taxes compress $\hat{w}_t$, which in turn compresses s_t , reducing the supply of funds available for investment. The entire supply-side contraction flows through this single channel. However, because external debt does not appear on the demand side of the capital market, this is the only channel at work, contrary to internal debt.

Having established how external debt alters the savings function, note that with external debt, young workers save from their net wage $\hat{w}_t$ rather than their gross wage w_t . Since foreign creditors absorb the government debt, the entire stock of domestic savings continues to flow into physical capital; the demand side of the market is unaffected. The equilibrium condition therefore becomes:

$$r_{t+1} = f' \left(\frac{s(\hat{w}_t, r_{t+1})}{1 + n} \right)$$

Its structure is identical to the baseline condition, with the sole modification that w_t is replaced by $\hat{w}_t = w_t - (r_t - n)g_1$ as the argument of the savings function. This substitution encodes the transmission mechanism of external debt: taxes reduce disposable income, disposable income governs saving, and saving determines the capital stock available for production in the next period.

It is useful to separate the newly derived equilibrium equation into its two con-

stituent relationships. The factor-price frontier $\phi(r)$,

$$w_t = \phi(r_t), \quad \phi'(r) = -k < 0$$

is entirely unaffected by the introduction of external debt, since it depends only on the production function and the capital-labour ratio. The second relationship, denoted ψ , maps last period's wage into the current interest rate through the capital market:

$$r_{t+1} = \psi(w_t; g_1)$$

Only ψ is modified by external debt. The economy's dynamics are governed by the composition $r_{t+1} = \psi(\phi(r_t); g_1)$, and the long-run equilibrium r is the fixed point of this mapping. The introduction of g_1 shifts ψ and thus displaces r .

The modification of the savings argument also alters the stability condition for the capital-market equilibrium. Implicitly differentiating the equilibrium equation with respect to r_t , and applying the chain rule through both the net-wage tax term $-(r_t - n)g_1$ and the factor-price frontier $\phi(r_t)$, yields:

$$0 < \frac{dr_{t+1}}{dr_t} = \frac{-f''(k + g_1) \frac{\partial s}{\partial w}}{1 + n - f'' \frac{\partial s}{\partial r}} \leq 1$$

The structure of this condition mirrors the baseline stability condition, with k replaced by $(k + g_1)$ in the numerator. Although Diamond states it without further commentary, the substitution admits a transparent decomposition; observe that the numerator is governed by the total derivative of the net wage with respect to the current interest rate:

$$\frac{d\hat{w}_t}{dr_t} = \phi'(r_t) - g_1 = -(k + g_1),$$

which combines two distinct channels through which r_t moves $\hat{w}_t$: the first is the factor-price-frontier channel: along the equilibrium path $w_t = \phi(r_t)$, a higher current interest rate lowers the gross wage at rate $\phi'(r_t) = -k$. The second is the fiscal channel: a higher r_t raises the per-worker tax bill, compressing the net wage. Both effects propagate through the savings function via the wage-income response $\partial s / \partial w$, which is why the numerator carries the combined sensitivity $-f''(k + g_1) \partial s / \partial w$. The term $\partial s / \partial r$, by contrast, captures savings' direct response to the future interest rate; because r_{t+1} appears on both sides of the equilibrium equation, this term is obtained by the implicit-differentiation and ends up in the denominator, accounting for the deviation of $1 + n - f'' \partial s / \partial r$ from the baseline factor $1 + n$. The capital-market

response to an interest-rate change is thus amplified by the presence of the external debt stock. However, the amplification operates entirely through the wage-income channel, not through the savings-response channel.

The signs are consistent with the assumptions maintained throughout. Since $f'' < 0$ and $\partial s/\partial w > 0$, the numerator is positive. Under the assumed relative slopes of the demand and supply curves in the capital market, the denominator is also positive, and the stability bound of unity is satisfied. Diamond assumes that a unique stable equilibrium exists both with and without the debt, and this condition formalises that assumption in the presence of g_1 .

To determine how the introduction of external debt displaces the ψ curve, we implicitly differentiate the equilibrium condition with respect to g_1 , holding w_t fixed and treating r_{t+1} as a function of g_1 :

$$\frac{\partial r_{t+1}}{\partial g_1} = \frac{f''(n-r) \frac{\partial s}{\partial w}}{1+n-f'' \frac{\partial s}{\partial r}}$$

Since the denominator is positive by the stability condition, the sign of this expression is determined entirely by the numerator. Given $f'' < 0$ and $\partial s/\partial w > 0$, the sign of $\partial r_{t+1}/\partial g_1$ equals the sign of $(r-n)$. At any w_t such that the implied r_t exceeds n , the new ψ lies above the old; at w_t values implying $r_t < n$ the new ψ lies below. The two curves intersect at the wage corresponding to $r_t = n$.

In long-run equilibrium the interest rate is stationary, so the equilibrium condition collapses to a relationship between r and g_1 :

$$r = f' \left(\frac{s(\phi(r) - (r-n)g_1, r)}{1+n} \right)$$

Implicitly differentiating with respect to g_1 and collecting terms yields:

$$\frac{dr}{dg_1} = \frac{-f''(r-n) \frac{\partial s}{\partial w}}{1+n-f'' \frac{\partial s}{\partial r} + f''(k+g_1) \frac{\partial s}{\partial w}}$$

The denominator is positive by the stability condition. Since $f'' < 0$ and $\partial s/\partial w > 0$, the numerator has the same sign as $(r-n)$, and therefore so does dr/dg_1 . Geometrically, this sign reversal manifests as a rotation of ψ around the point $r = n$. External debt shifts the equilibrium interest rate away from the Golden Rule: it raises r in the efficient case, thereby compressing the capital-labour ratio, and lowers r in the

inefficient case, thereby expanding it. With the factor-price frontier unchanged, the equilibrium wage moves in the opposite direction to the interest rate in both cases, as depicted in Figure 5.




Figure 5: Shift in the ψ curve caused by external debt

Thus, the introduction of external debt displaces the economy's long-run path; if, prior to the issuance of debt, $r \neq n$, the issuance widens the gap between r and n ; otherwise, if $r = n$, no displacement occurs.

Having established how external debt shifts the long-run equilibrium interest rate, it is possible to determine its effect on the utility of an individual living in that equilibrium.

Recall from the competitive framework that the derived utility function satisfies:

$$\frac{\partial U}{\partial w} = \frac{\partial U}{\partial c_1}, \quad \frac{\partial U}{\partial r} = \frac{s}{1+r} \cdot \frac{\partial U}{\partial c_1}$$

The total change in utility as g_1 varies is obtained by the chain rule through both factor prices:

$$\frac{dU}{dg_1} = \frac{\partial U}{\partial c_1} \left[\frac{d\hat{w}}{dg_1} + \frac{s}{1+r} \frac{dr}{dg_1} \right]$$

The net wage $\hat{w} = \phi(r) - (r - n)g_1$ depends on g_1 both directly, through the tax term, and indirectly, through the equilibrium interest rate r . Differentiating with respect to g_1 :

$$\frac{d\hat{w}}{dg_1} = \phi'(r) \frac{dr}{dg_1} - g_1 \frac{dr}{dg_1} - (r - n) = -(k + g_1) \frac{dr}{dg_1} - (r - n)$$

where the substitution $\phi'(r) = -k$ follows from the factor-price frontier.

Substituting this into the utility expression and collecting terms:

$$\frac{dU}{dg_1} = -\frac{\partial U}{\partial c_1} \left[(r - n) + g_1 \frac{dr}{dg_1} + \left(k - \frac{s}{1+r} \right) \frac{dr}{dg_1} \right]$$

Since $\partial U/\partial c_1 > 0$ by non-satiation, the sign of dU/dg_1 is the opposite of the bracketed expression. Diamond decomposes the bracket into three terms, each with a distinct economic interpretation; in the efficient case each is positive, so utility falls unambiguously.

Term I: the direct tax cost, $(r - n)$. This is the per-worker lump-sum tax required to service the marginal unit of debt. In the efficient case ($r > n$) it is positive, reflecting a direct compression of lifetime consumption; in the inefficient case it is negative, since the tax becomes a net transfer to the young.

Term II: the burden on existing debt, $g_1 dr/dg_1$. As the equilibrium interest rate adjusts, the cost of servicing the already-outstanding stock g_1 changes proportionally. In the efficient case $dr/dg_1 > 0$, so this term is positive: the rising interest rate increases the tax burden on existing debt, further reducing utility.

Term III: the factor-price effect, $(k - s/(1+r)) dr/dg_1$. This term captures the welfare consequence of the interest rate moving along the factor-price frontier. Its sign is not immediately transparent, and requires the following simplification. Using the capital-market clearing condition $k = s/(1+n)$:

$$k - \frac{s}{1+r} = \frac{s}{1+n} - \frac{s}{1+r} = s \cdot \frac{r-n}{(1+n)(1+r)} = \frac{k(r-n)}{1+r}$$

Substituting into Term III:

$$\left(k - \frac{s}{1+r} \right) \frac{dr}{dg_1} = \frac{k(r-n)}{1+r} \cdot \frac{dr}{dg_1}.$$

Since dr/dg_1 has the same sign as $(r - n)$, term III is positive in both regimes: the equilibrium is moving away from the Golden Rule, and any displacement worsens welfare. Geometrically, the factor-price frontier and the derived indifference curve are tangent precisely at $r = n$, where their slopes $-k$ and $-s/(1+r)$ coincide; on either side of this point the slopes diverge, and the economy slides off the welfare-maximising tangency.

Combining the three terms, in the efficient case the bracket is the sum of three positive quantities; with the leading negative sign, $dU/dg_1 < 0$ unambiguously. In

the normal case, where the competitive equilibrium is efficient ($r > n$), external debt unambiguously reduces the long-run utility of individuals living in equilibrium. In the inefficient case ($r < n$), Terms I and II become negative while Term III remains positive, so the bracket has an ambiguous sign, and no a priori welfare ranking is available.

The ambiguity in the inefficient case is not a weakness of the model. It accurately reflects the non-standard nature of a dynamically inefficient equilibrium: the negative tax constitutes a transfer to the young that partially or wholly offsets the factor-price distortion, and the net balance depends on the economy's specific parameter values. The normal case, by contrast, admits a clean and unambiguous conclusion, one that will serve as the benchmark against which the welfare cost of internal debt is measured in the section that follows.

The analysis of external debt can be summarised compactly. The introduction of a positive external debt-labour ratio g_1 sets in motion a chain of causation. The lump-sum tax $(r_t - n)g_1$ reduces the disposable wage of the young, compressing savings, contracting the capital supply, and shifting the ψ curve. In the efficient case, this shift raises the equilibrium interest rate, widening the gap from the Golden Rule, and each of the three terms in the welfare expression is negative. The conclusion is unambiguous: external debt reduces long-run utility in the normal case.

Object	Effect of $g_1 > 0$	Sign (efficient case)
Per-worker tax τ	$(r - n)g_1$	+
Net wage $\hat{w}$	Falls	—
Savings $s(\hat{w}, r)$	Falls	—
ψ curve	Shifts upward	↑
Equilibrium r^*	Moves away from n	↑
Equilibrium k^*	Falls	—
dU/dg_1	Negative	—

Table 1: Comparative statics of external debt in the efficient case

What the external debt case deliberately abstracts from is equally important. Throughout this subsection, the demand side of the capital market has remained untouched: domestic savers allocate their entire savings to physical capital, and the government paper is absorbed abroad. The factor-price frontier is unchanged, the capital-market clearing condition retains its baseline form on the demand side, and the welfare loss arises solely from the fiscal compression of saving. This abstraction will not survive the introduction of internal debt, as shown in the next subsection.

3.4.2 Internal Debt - The Portfolio Substitution Channel

The structural difference from the external case lies in a single feature: the bondholders are domestic savers rather than foreign agents. The implication, however, is consequential, because it forces the government onto the demand side of the capital market for the first time; internal debt requires the young's savings to finance both physical capital and government paper. The supply-side apparatus therefore carries over without modification, but a second, demand-side channel must now be added.

Denote the internal debt-labour ratio by g_2 , defined in parallel with g_1 as the per-worker stock of public debt outstanding at the start of period t . Throughout the present section we hold $g_1 = 0$ and study internal debt in isolation; the two ratios will coexist further in the text. Two results, sharper than their external-debt counterparts, will be established. First, internal debt unambiguously raises the long-run equilibrium interest rate, irrespective of efficiency. Second, it lowers steady-state utility in the efficient case and raises it in the inefficient case.

The supply side of the capital market is altered under internal debt in precisely the same fashion as under external debt. From the standpoint of the individual taxpayer, the destination of the tax revenue, whether it flows abroad to foreign bondholders or to domestic bondholders of the same generation, is immaterial to the budget constraint he faces. The young pay lump-sum taxes of $(r_t - n)g_2$ per worker to finance the portion of interest costs not rolled over into new debt, and they solve the same intertemporal optimisation problem as in the external debt case with the disposable wage $\hat{w}_t = w_t - (r_t - n)g_2$ substituted for w_t in the savings decision. The savings function therefore takes the form $s(\hat{w}_t, r_{t+1}) = s(w_t - (r_t - n)g_2, r_{t+1})$ structurally identical to its external-debt counterpart the external-debt counterpart. Normality of present consumption, $0 < \partial s / \partial w < 1$, implies once again that an increase in g_2 depresses savings whenever $r > n$ and raises them whenever $r < n$. The structural change wrought by internal debt appears not in the savings function but in the asset-market identity, to which we now turn.

The structural departure from the external-debt case lies on the demand side of the capital market. When the government issues debt to domestic savers, the young's accumulated wealth must finance both the productive capital stock and the outstanding stock of government paper. The market identity therefore reads

$$S_t = K_{t+1} + G_{t+1},$$

where $G_{t+1} = g_2 L_0 (1 + n)^{t+1}$ by the law of motion established above. Dividing

through by $L_{t+1} = L_0(1+n)^{t+1}$ yields the per-worker form

$$\frac{s_t}{1+n} = k_{t+1} + g_2.$$

The latter equation is the analogue, under internal debt, of the baseline market-clearing condition $s_t/(1+n) = k_{t+1}$ derived in competitive equilibrium: a fixed margin g_2 now separates the per-worker savings of the young from the capital-labour ratio that prevails in the next period. This margin measures, period by period, the share of savings absorbed by government paper rather than productive capital. This is the portfolio-substitution channel: it has no counterpart in the external-debt case, where G_{t+1} was absorbed abroad, and the domestic identity reduced to $S_t = K_{t+1}$.

Combining the market-clearing equation with the marginal-productivity schedule for capital, $r_{t+1} = f'(k_{t+1})$, yields the equilibrium condition for the capital market under internal debt:

$$r_{t+1} = f' \left(\frac{s(w_t - (r_t - n)g_2, r_{t+1})}{1+n} - g_2 \right).$$

The equilibrium condition embeds both channels through which g_2 acts on the long-run equilibrium. The argument $w_t - (r_t - n)g_2$ in the savings function captures the tax channel, which is structurally identical to the external-debt case. The standalone $-g_2$ inside $f'(\cdot)$ captures the portfolio-substitution channel, which is new.

The stability requirement under internal debt follows from implicit differentiation of the equilibrium condition with respect to r_t , treating r_{t+1} as a function of r_t along the equilibrium path. The resulting necessary condition for local stability, together with the maintained assumption on the relative slopes of the demand and supply schedules in the capital market, reads:

$$0 < \frac{dr_{t+1}}{dr_t} = \frac{-f''(k + g_2) \frac{\partial s}{\partial w}}{1+n - f'' \frac{\partial s}{\partial r}} \leq 1.$$

The stability condition is structurally identical to that of the external-debt case, with g_1 replaced by g_2 : the multiplier $(k + g_2)$ in the numerator decomposes, just as before, into a factor-price-frontier component and a fiscal component. The portfolio-substitution channel does not appear in the equation, because g_2 is held constant when differentiating along the equilibrium path. Its effect will appear instead in the shift ψ .

Under external debt, the shift of ψ rotated around the point $r = n$, and the sign

of $\partial r_{t+1}/\partial g_1$ tracked the sign of $(r - n)$. Internal debt removes this ambiguity: the locus shifts upward at every interest rate, and the equilibrium interest rate rises whether the economy is dynamically efficient or inefficient.

Implicit differentiation of the equilibrium condition with respect to g_2 , holding w_t and r_t at their long-run values, yields

$$\frac{\partial r_{t+1}}{\partial g_2} = \frac{-f'' \left[(r - n) \frac{\partial s}{\partial w} + (1 + n) \right]}{1 + n - f'' \frac{\partial s}{\partial r}}.$$

The two terms inside the square brackets isolate the two channels through which internal debt acts. The first, $(r - n) \partial s / \partial w$, is the tax channel familiar from the external case: a unit increase in g_2 raises per-worker taxes by $(r - n)$, which shifts disposable wage income by the same amount and savings by $(r - n) \partial s / \partial w$ via normality of present consumption. Its sign tracks $(r - n)$, and it vanishes at the Golden Rule. The second, $(1 + n)$, is the portfolio-substitution channel, new to internal debt and absent under external debt. Its origin is mechanical: the market-clearing identity $k_{t+1} = s_t / (1 + n) - g_2$ implies that, holding savings fixed, each additional unit of g_2 displaces one unit of next-period per-worker capital. Clearing the savings denominator in the implicit-derivative calculation converts this one-for-one displacement into the $(1 + n)$ that appears inside the bracket; the factor is therefore a unit conversion between savings per old worker and capital per young worker, not an additional economic effect. The substantive content is the unit displacement itself, which is independent of r and unambiguously positive.

The sign of $\partial r_{t+1} / \partial g_2$ now follows. The denominator is positive by the stability requirement. In the numerator, $-f'' > 0$ multiplies a bracketed expression whose sign must be checked only in the dynamically inefficient region $r < n$, since for $r \geq n$ the tax-channel term is non-negative and the bracket exceeds $(1 + n) > 0$ trivially. For $r < n$, maintaining the standard regularity condition $r > -1$, normality of present consumption $0 < \partial s / \partial w < 1$ implies

$$(r - n) \frac{\partial s}{\partial w} + (1 + n) > -(n - r) + (1 + n) = 1 + r > 0.$$

Hence, $\partial r_{t+1} / \partial g_2 > 0$ for all admissible r , and the locus ψ shifts upward uniformly.

Under external debt, the tax channel was the only mechanism through which g_1 entered the equilibrium condition, and its sign inherited that of $(r - n)$; the dynamically inefficient region was therefore one in which external debt lowered the interest rate, and the locus of long-run equilibria pivoted around the Golden Rule.

Internal debt introduces a second mechanism that operates with a fixed sign, and the portfolio-substitution channel uniformly dominates the tax channel: even when the tax channel reverses, the $(1+n)$ term ensures that the net shift remains positive. The rotational ambiguity of the external-debt case is replaced by a uniform upward shift of ψ , and the long-run interest rate moves in a single direction regardless of the position of the economy relative to the Golden Rule.

To obtain the change in the long-run equilibrium interest rate, the wage must be allowed to adjust along the factor-price frontier, $w = \phi(r)$. Substituting this relation into the equilibrium condition yields the locus of long-run equilibria parametrised by g_2 ,

$$r = f' \left(\frac{s(\phi(r) - (r-n)g_2, r)}{1+n} - g_2 \right),$$

Differentiating implicitly with respect to g_2 delivers the comparative-static result

$$\frac{dr}{dg_2} = \frac{-f'' \left[(r-n) \frac{\partial s}{\partial w} + (1+n) \right]}{1+n - f'' \frac{\partial s}{\partial r} + f''(k+g_2) \frac{\partial s}{\partial w}}.$$

The numerator here reproduces that of the ψ -shift without modification: the two channels through which g_2 enters the equilibrium condition, tax and portfolio-substitution, act on the long-run equilibrium with the same weights with which they shifted ψ . The denominator, by contrast, picks up the additional term $+f''(k+g_2)\partial s/\partial w$, which captures the feedback through the factor-price frontier and the fiscal channel: a higher equilibrium interest rate lowers the wage by $\phi'(r) = -k$ along the factor-price frontier, and simultaneously raises the tax bill on the existing stock of debt by g_2 per unit of r . The two effects compound: the combined term $-(k+g_2)$ depresses savings through $\partial s/\partial w$, partially offsetting the initial impulse on r . The sign is unambiguous. The numerator is positive by the bound established above. The denominator is positive by the stability requirement, to which the new $f''(k+g_2)\partial s/\partial w$ term contributes positively (since $-f'' > 0$ and $\partial s/\partial w > 0$). Hence $\frac{dr}{dg_2} > 0$ for all admissible long-run equilibria, irrespective of the sign of $(r-n)$.

The economic content of this conclusion is the counterpart to the result on external debt. External debt shifted the equilibrium interest rate away from the growth rate: upward in the efficient region and downward in the inefficient region, with the Golden Rule as a fixed point. Internal debt moves it upward without exception. In the dynamically efficient region, the two instruments operate in the same direction; in the inefficient region, they operate in opposite directions, with internal debt pushing r toward n and external debt pushing it away from n . This asymmetry is traceable

in $\frac{dr}{dg_2}$ to the $(1+n)$ term that the portfolio-substitution channel contributes to the numerator.

The welfare consequences of internal debt are obtained by combining the utility decomposition derived in the external case, which holds for either form of debt, since it tracks only the relation between disposable wage, interest rate, and steady-state utility, with the modified market-clearing condition. Two equivalent expressions emerge. The first makes the underlying channels visible; the second collapses them into a form whose sign is immediate.

The three-term decomposition reads

$$\frac{dU}{dg_2} = -\frac{\partial U}{\partial c_1} \left[(r-n) + g_2 \frac{dr}{dg_2} + \left(k - \frac{s}{1+r} \right) \frac{dr}{dg_2} \right].$$

The interpretation parallels the external case term by term. The first term is the direct fiscal cost of servicing the debt: each additional unit of g_2 requires per-worker taxes of $(r-n)$ in the steady state, and these taxes raise (or lower, in the inefficient case) the young's lifetime tax bill. The second term is the indirect fiscal cost: the rise in the equilibrium interest rate increases the tax burden of servicing the pre-existing stock g_2 . The third term is the factor-price effect: a higher interest rate lowers the wage along the factor-price frontier, while simultaneously raising the return on accumulated savings, and the net welfare consequence depends on the sign and magnitude of $k - s/(1+r)$.

The key difference from external debt is concentrated in this third term. Under external debt, s measured the savings absorbed entirely by physical capital, and the term $k - s/(1+r)$ captured the imbalance between the wage loss and the savings return at the prevailing interest rate. Under internal debt, the market-clearing condition implies $s = (1+n)(k + g_2)$: savings now finance both capital and government paper. The consequence is that g_2 rejoins k inside the savings identity, and the indirect fiscal and factor-price terms of the welfare equation can be combined into a single expression proportional to $(r-n)$. Specifically, substituting $s = (1+n)(k + g_2)$ yields

$$g_2 + \left(k - \frac{s}{1+r} \right) = (k + g_2) \left[1 - \frac{1+n}{1+r} \right] = \frac{(k + g_2)(r-n)}{1+r},$$

which factors $(r-n)$ out of the bracket and collapses the three terms of the equation into a single product:

$$\frac{dU}{dg_2} = -\frac{\partial U}{\partial c_1} (r-n) \left[1 + \frac{k + g_2}{1+r} \frac{dr}{dg_2} \right].$$

The bracketed factor is unambiguously positive: the unit term is positive by inspection, and the second term is positive because $k + g_2 > 0$, $1 + r > 0$, and $dr/dg_2 > 0$. The marginal utility $\partial U/\partial c_1$ is positive by assumption, and the leading minus sign is fixed. The sign of dU/dg_2 is therefore the sign of $-(r - n)$, with no residual ambiguity in either dynamic regime:

$$\frac{dU}{dg_2} < 0 \text{ if } r > n, \quad \frac{dU}{dg_2} > 0 \text{ if } r < n.$$

Under internal debt, the welfare verdict reverses cleanly across the Golden Rule: negative when $r > n$, positive when $r < n$, whereas under external debt only the efficient region was sign-determinate, with the inefficient region yielding no unambiguous conclusion. Internal debt is therefore the sharper instrument: it delivers a sign in both regimes, and the sign reverses at the growth rate.

A final observation, due to Diamond, sharpens the interpretation. The clean final expression for welfare shows that the sign of dU/dg_2 is the sign of $-(r - n) dr/dg_2$. Because $dr/dg_2 > 0$, this sign coincides with the sign of $(n - r)$: any movement of the interest rate toward the growth rate raises utility, and any movement away from it lowers utility. Internal debt acts purely as a device that relocates the equilibrium along the long-run locus, and its welfare effect is governed entirely by whether that relocation approaches or recedes from the Golden Rule.

Diamond's own treatment of this conclusion is geometric, and a brief paraphrase clarifies the structural content of the welfare equation. In the consumption plane (c_1, c_2) , the consumption-possibility frontier of a society with capital-labour ratio k is the line $c_1 + c_2/(1 + n) = y - nk$, whose height is determined by r through the marginal-productivity relation $r = f'(k)$, and whose maximal value of $y - nk$ is attained at the Golden Rule $r = n$. The competitive equilibrium lies on this frontier at the point where an indifference curve is tangent to it with slope $-(1 + r)$. Normality of present and future consumption implies that, along any fixed budget line, indifference curves become steeper as one moves up and to the left.

A perturbation of g_2 that raises the interest rate from r to r' then operates through two superimposed effects, both unfavourable in the efficient case. First, the consumption-possibility frontier shifts downward, since $y - nk$ is maximised at $r = n$ and any movement of r away from n reduces it. Second, the equilibrium relocates along the lowered frontier toward a steeper tangency, and by normality this places the new equilibrium on a lower indifference contour than the point of equal slope on the upper frontier. The two effects compound to lower utility unambiguously when $r > n$

and, by the symmetric argument, raise it unambiguously when $r < n$. The geometry thus recovers without algebra: internal debt does not alter the social consumption possibilities of the economy, but it relocates the equilibrium along them, and the welfare verdict is governed entirely by whether that relocation approaches or recedes from the Golden Rule.

3.4.3 Internal Versus External Debt

With both instruments in play, the apparatus of the previous two subsections merges without fresh derivation. Total debt $g_1 + g_2$ is serviced by taxes on the young, so the entire stock enters the disposable wage through the tax channel; only internal debt sits on the demand side, so only g_2 enters the market-clearing identity as a standalone crowding-out term. The equilibrium condition for the capital market is accordingly

$$r_{t+1} = f' \left(\frac{s(w_t - (r_t - n)(g_1 + g_2), r_{t+1})}{1 + n} - g_2 \right),$$

and the stability requirement carries over with $k + g_1 + g_2$ in the numerator, its fiscal component now reflecting taxes on the combined stock. Letting the wage adjust along $w = \phi(r)$ and differentiating the long-run locus delivers the two comparative statics

$$\frac{dr}{dg_1} = \frac{-f''(r - n) \frac{\partial s}{\partial w}}{D}, \quad \frac{dr}{dg_2} = \frac{-f'' \left[(r - n) \frac{\partial s}{\partial w} + (1 + n) \right]}{D},$$

sharing the denominator $D = 1 + n - f'' \partial s / \partial r + f''(k + g_1 + g_2) \partial s / \partial w > 0$. These simply reproduce the single-instrument verdicts: external debt moves r away from the growth rate, internal debt raises it without exception. The welfare effect of each, read off the decomposition that holds for either instrument, is

$$\frac{dU}{dg_i} = -\frac{\partial U}{\partial c_1} \left[(r - n) + \frac{dr}{dg_i} (g_1 + g_2) + \frac{dr}{dg_i} \left(k - \frac{s}{1 + r} \right) \right], \quad i = 1, 2.$$

The differential-incidence experiment is the substitution of internal for external debt at constant total debt: raise g_2 and retire an equal quantity of g_1 , captured by the operator $d/dg_2 - d/dg_1$. Because total debt, and hence the interest bill and the taxes financing it, is unchanged, the supply side of the capital market is left untouched; the swap operates purely through the demand side. Subtracting the comparative

statics, the tax channel cancels, and only the portfolio-substitution term survives:

$$\frac{dr}{dg_2} - \frac{dr}{dg_1} = \frac{-f''(1+n)}{D} > 0.$$

The interest rate always rises, and the capital-labour ratio always falls: replacing foreign creditors with domestic ones compels savers to hold government paper in place of productive capital.

The welfare consequence follows from subtracting the two utility expressions. The direct tax term $(r-n)$ is common to both and cancels, as does the debt-stock weight, leaving

$$\frac{dU}{dg_2} - \frac{dU}{dg_1} = -\frac{\partial U}{\partial c_1} \left(\frac{dr}{dg_2} - \frac{dr}{dg_1} \right) \left[(g_1 + g_2) + \left(k - \frac{s}{1+r} \right) \right].$$

Evaluating the bracket with $s = (1+n)(k+g_2)$, exactly as in the internal-debt section, collapses it into a form whose sign is immediate:

$$\frac{dU}{dg_2} - \frac{dU}{dg_1} = -\frac{\partial U}{\partial c_1} \left(\frac{dr}{dg_2} - \frac{dr}{dg_1} \right) \left[g_1 + \frac{(k+g_2)(r-n)}{1+r} \right].$$

The same change may equivalently be split into the part travelling abroad through external interest payments, the g_1 term, which alters the economy's consumption possibilities, and a purely domestic part carrying the sign of $(n-r)$, the decomposition that ties the swap back to the external-debt analysis of the preceding section.

The sign now reads off directly. Marginal utility is positive, the interest-rate response is positive, so the welfare effect of the swap inherits the sign of $-[g_1 + (k+g_2)(r-n)/(1+r)]$. In the efficient case, $r > n$, both terms in the bracket are positive, and the conclusion is unambiguous: substituting internal for external debt lowers steady-state utility. This is the paper's central comparative claim: that, in the normal case, internal debt is the more costly instrument. In the inefficient case $r < n$ the second term turns negative, and the verdict depends on the relative magnitudes of g_1 and $(k+g_2)(n-r)/(1+r)$; no general sign is available.

The asymmetry admits a clean economic reading. External debt provides a relief valve: a portion of the interest burden is paid to foreigners, so part of the cost leaves the domestic consumption-possibility set rather than merely being reallocated within it. Internal debt has no such outlet. It imposes the full interest cost at home and, on top of that, substitutes paper for capital in savers' portfolios, depressing output. The debt swap removes the foreign relief valve while activating the crowding-out channel, so, in the efficient case, it can only reduce welfare. Diamond notes the mirror-image

framing: external debt is internal debt plus marginal foreign borrowing, which leaves net output unchanged but shifts relative factor prices and, through them, savings and utility.

3.5 Remarks

The signs above are those of the Walrasian-stable case maintained throughout. Diamond shows in his Appendix A that, under the alternative configuration of the capital market, the denominators reverse sign and, with them, every comparative-static verdict; the welfare ranking of internal against external debt is therefore a statement about the normal case, not a universal one.

The three instruments analysed above admit a single verdict, organised entirely by the position of the equilibrium interest rate relative to the growth rate. External debt moves r away from n : it lowers steady-state utility in the efficient case and leaves the sign open in the inefficient case, and this holds whether or not internal debt is also present. Internal debt raises r unconditionally: taken in isolation, it lowers utility in the efficient case and raises it in the inefficient case, though the presence of external debt can overturn the inefficient-case gain. The substitution of internal for external debt at constant total debt likewise raises r : it lowers utility unambiguously in the efficient case and yields no general sign in the inefficient case. The common thread is that, once the direct servicing cost is netted out, welfare rises or falls precisely as the interest rate is moved toward or away from the Golden Rule.

Diamond offers two ways of classifying these effects, each illuminating a strand of the prior literature. The first divides the welfare change into a tax component and a factor-payment component. Its content is that the taxes required to service debt fall on the young identically, whether the debt is held at home or abroad: the destination of the interest payment is immaterial to the taxpayer's budget constraint. The two forms of debt therefore differ not in their fiscal incidence but only in their consequences for factor prices. The second classification rests on the identity that internal debt is external debt followed by a debt swap. The factor-payment difference then separates cleanly along the two sides of the capital market: external debt operates on the supply side, through the savings that taxation withdraws, while the swap operates on the demand side, through the government paper that displaces capital in savers' portfolios.

This second decomposition locates Diamond's contribution against its predecessors. Bowen, Davis, and Kopf had isolated the tax effects of internal debt, the supply-side fiscal channel, and stopped there. Modigliani and Vickrey had isolated the

opposite channel, the fall in the capital stock as debt substitutes for capital in portfolios, but treated it in partial terms; Modigliani in particular assumed a one-for-one replacement at constant total wealth, setting aside the feedback by which the resulting output loss itself alters savings and hence wealth. Diamond's framework subsumes both. By embedding debt in a utility-maximising overlapping-generations economy with endogenous output and factor prices, the tax and portfolio channels emerge jointly from a single equilibrium, and the feedback Modigliani set aside is carried through to the comparative statics rather than assumed away.

The chapter has read Diamond as an extension of Samuelson, and the comparison sharpens at the close. In Samuelson's consumption-loan economy, the social contrivance, money, is a pure device for intertemporal transfer: it reallocates consumption across generations without any counterpart in the economy's productive structure. Diamond's national debt does this and one thing more. It is at once a transfer device and a substitute for capital, and it is this second role, absent from Samuelson, that generates the portfolio-substitution channel and with it the asymmetry between internal and external debt. The Golden Rule $r = n$ remains the organising benchmark in both cases: in Samuelson, it is the rate at which the transfer device sustains the optimal allocation; in Diamond, it is the pivot at which every debt instrument reverses its welfare sign. It is this dual character of debt, transfer and capital substitute, that the remainder of the thesis takes up.

3.6 Safe Assets and the Demand for Debt

The preceding sections presented public debt, in Diamond's setting, as a claim on the future that households hold to transfer purchasing power across periods, and as one that can raise welfare when the economy overaccumulates capital. Gary Gorton (2017) takes this store-of-value role as its starting point but poses a prior question: what property allows a financial claim to serve as such a store in the first place, and why does its supply impact the stability of the system as a whole? He answers that the relevant claims are safe assets, and that almost all of monetary and financial history can be read as the search for them. Coins of precious metal and, from the eighteenth century, bills of exchange and bank notes were early attempts, but Gorton stresses that they were never reliably safe: their value had to be established through assaying, monitoring, or successive endorsement, and they remained exposed to debasement, clipping, and default. The history is thus a record of how difficult it is to produce safety, and it motivates the analytical account that follows.

3.6.1 What Makes Debt Safe

A safe asset, in Gorton's definition, is one that can be exchanged at face value without fear of adverse selection: no counterparty has an incentive to acquire private information about its value, and the absence of that incentive is common knowledge. He summarises the property with the phrase "no questions asked" (NQA) and the equivalent term information-insensitivity. Because no holder expects to be fooled by a better-informed trader, a safe asset is exchanged without friction and can also be held to carry value forward in time; in this sense, it is the lifeblood of exchange.

Gorton identifies multiple facets that make debt "safe", and part of his contribution is to show the relationship between them. In one tradition, bank debt is a form of insurance that lets agents smooth consumption: real investment pays off only after a delay, agents may need to consume before then, and on-demand deposits allow those who need funds early to withdraw while the rest wait for the investment to mature. In a second, safety means pledgeability (assets whose cash flows are sufficiently verifiable to be pledged to outsiders), which insulates firms against being forced to abandon a project when a liquidity shock arrives. In a third, which Gorton makes central, safe debt is simply debt that can be used in transactions without fear of privately informed counterparties. What unites the three is that a safe instrument must hold its value across states of the world; the differences concern only whether the friction being insured against is early consumption, forced liquidation, or adverse selection in trade.

Gorton also separates two attributes that the literature sometimes conflates. "Safety" describes a long-term claim that will repay at par with very high probability (Treasury bonds and AAA-rated asset-backed securities are his examples) so that it stores value reliably over time. "Moneyiness", used interchangeably with the liquidity premium, is the NQA property as it bears on transactions, and characterises short-term instruments such as Treasury bills, commercial paper, and repurchase agreements. Short maturity and the option to convert reduce exposure to information arrival and thus confer moneyiness; longer-dated claims must instead earn their safety from the strength of their backing.

Why is debt the optimal instrument for creating a safe contract? Gorton's answer turns on the shape of debt's payoff. Debt is senior to equity and promises a fixed amount, so its value is comparatively insensitive to news as long as the backing collateral lies comfortably above the default boundary. When collateral is believed to sit far enough above that boundary, producing information about its exact value is pointless and costly, and the claim stays information-insensitive; only as the col-

lateral approaches the boundary does it begin to pay to discover what it is worth. Debt is therefore engineered to keep the underlying "state of the world" opaque. The same logic explains how safety fails. Privately produced safe debt is risky; Gorton uses "credit risk" for the possibility that it is not in fact safe, and the passage from information-insensitive to information-sensitive corresponds to a downgrade below AAA. In this reading, a credit rating is best understood as a measure of the distance to the boundary at which questions are asked.

3.6.2 Collateral, Opacity, and the Production of Safe Debt

If safe short-term debt must be credibly backed, what backs it? Gorton's answer is that safe collateral is itself debt. The short-term claim used to transact is collateralised by a longer-term claim, and because pledging an already information-insensitive bond makes the new contract even more information-insensitive, the optimal arrangement is thus "debt-on-debt". Short-term safety, therefore, presupposes a stock of long-term safe debt standing behind it: repo is secured by a specific high-grade bond, asset-backed commercial paper by ABS/MBS, and, historically, free-banking notes by state bonds, while insured demand deposits are ultimately backed by the taxing power of the state. The demand for short-term safety is a demand for safe long-term assets, and this dependence is what binds the two ends of the maturity spectrum together.

A second mechanism concerns the institution that issues the debt. For bank debt to retain the NQA property, Gorton argues, the bank must be opaque: information about its assets must be kept from outsiders, and the assets themselves must be chosen so that no expert has an incentive to investigate them. Loans to households and small firms are well suited to this purpose precisely because they are costly for outsiders to learn about, whereas large firms, whose value markets can assess, raise funds directly in capital markets. Opacity, ordinarily a vice, is here a condition for producing money-like claims. Securitisation extends the same idea to the capital market: it takes immobile bank loans and converts them into mobile AAA-rated bonds that can be traded, pledged, rehypothecated, and held as a store of value. Its complexity is functional: when a loan-by-loan valuation of a structured deal is prohibitively costly, information is simply not produced, and the senior tranche remains NQA, and the residual equity tranche is retained by the sponsor so that no information leaks into the market. Securitisation, in Gorton's account, is best understood as the private production of safe assets.

3.6.3 The Convenience Yield

Because safe debt yields liquidity and safety on top of its cash flows, investors accept a lower return for holding it; the value of these nonpecuniary properties of safe debt is the convenience yield, and much of the recent literature Gorton surveys is devoted to measuring it. He reports that, over much of the twentieth century, the yield on U.S. Treasuries ran on the order of 73 basis points below what it would otherwise have been: a premium attributable to their moneyness and safety. The premium is identified from yield spreads: the gap between a high-grade corporate bond and a Treasury of comparable maturity, for instance, widens when Treasuries are scarce, as heightened relative demand bids up their prices and lowers their yields. Different spreads isolate different components, and the evidence suggests that short-dated government debt is valued for short-term safety, that is, for moneyness, while long-term debt is valued for the near-certainty of eventual repayment, with the convenience yield tending to fall as maturity increases. One implication Gorton draws is conceptual: standard asset pricing values a claim at the expected product of its payoff and the consumer's intertemporal marginal rate of substitution, and a convenience yield is inconsistent with such models. Another is that Treasuries are non-Ricardian, and that close private substitutes such as AAA ABS/MBS carry a convenience yield of their own.

The decisive macro-financial implication follows from how the private sector responds to this premium. When safe public debt is scarce, the convenience yield rises, and financial intermediaries answer by issuing short-term or demandable claims and channelling the proceeds into illiquid loans, thus manufacturing safe assets where the government has not. Empirically, Gorton stresses two regularities: the supplies of public and of privately produced safe debt move inversely, and private issuance of securitised debt tracks the convenience-yield spread, expanding when the premium is high. The quantity of safe public debt is therefore far from neutral; it governs how much, and what kind of, private safe debt the system creates.

3.6.4 Regulation and the Rise of Shadow Banking

Because banks must be opaque and issue runnable short-term debt, Gorton observes, they are universally regulated, and the negative externalities of runs are the rationale for that regulation. Regulatory instruments include deposit insurance, which converts demand deposits into genuinely safe claims, and restrictions on entry: the requirement that banks be chartered, together with capital and reserve requirements and limits on permissible activities, restrains others from issuing short-term debt. Entry restriction has a subtler effect that Gorton emphasises. The rents it confers

give a bank a "charter value", an intangible asset that is lost if the bank fails, and the prospect of losing it disciplines risk-taking, aligning a bank's private incentives with the social goal of producing safe debt.

This mechanism helps explain the transformation that, in Gorton's reading, is central to the modern financial system. From the late 1970s, banks faced competition from non-banks such as money-market mutual funds and the junk-bond market, alongside deregulation of interest-rate ceilings and branching restrictions. As they grew less profitable, their charter values declined, weakening the incentive to avoid risk, and the shadow banking system grew up alongside this erosion, not by accident but as its counterpart. The aggregate evidence sharpens the point. The share of safe liabilities in total assets has been strikingly stable, at around a third, over more than half a century, which Gorton reads as a stable demand for safety; yet the composition of that safe debt changed profoundly. Demand deposits, which had supplied roughly four-fifths of privately produced safe debt, gave way to money-market instruments, repo, and securitised claims as the demands of wholesale investors came to dominate those of retail depositors. The change, he argues, is permanent.

There are, in the end, only two ways to make debt safe: back it with the government's taxing power or with credible private collateral. The first can be close to riskless; the second never fully is. Privately produced short-term debt is consequently always vulnerable to runs unless a public guarantee stands behind it. The mechanics are familiar but consequential: when holders come to doubt the collateral backing their claims, they demand cash, and if convertibility is not suspended, the issuers must sell assets in volume, driving prices down and inflicting mark-to-market losses that propagate into the real economy. A financial crisis, in Gorton's account, is precisely the event in which safe debt ceases to be safe: when the NQA property is lost, short-term debt switches from information-insensitive to information-sensitive, and questions that had gone unasked are suddenly asked of everyone at once. Such crises are not aberrations but a recurrent feature of market economies, and they are typically preceded by credit booms that tend to develop when safe government debt is insufficient, and the private sector has stepped in to fill the gap.

The logic of how crises are resolved confirms the diagnosis. The response is not to make debt more information-sensitive by exposing the collateral of individual institutions, but the reverse. Historically, bank clearinghouses suppressed bank-level information and issued joint liabilities that pooled members' collateral into a single portfolio; modern central banks lend against weaker collateral through anonymous facilities, shifting the quality of collateral in the economy toward government claims. The aim in both cases is to restore the NQA property across the system rather than

on a bank-by-bank basis.

The shortage of safe assets also has an international dimension. Countries differ in their capacity to produce safe claims, because pledgeable income is more limited where institutions are weaker, so the rest of the world holds U.S. Treasury and agency debt as a store of value: the United States, in effect, selling insurance to the world and running the corresponding external deficit, the phenomenon often called the global savings glut. That foreign demand has in turn deepened the domestic scarcity of safe assets.

These threads converge on monetary policy, where Gorton's conclusions depart from the conventional view that policy should focus solely on inflation and leave financial stability to regulators. In a system dominated by wholesale safe debt, an open-market operation exchanges one form of money for another, and so alters the quantity of safe public collateral, which in turn determines how much fragile private collateral is produced; too few Treasuries induce private substitution, a credit boom, and rising fragility. The two clienteles the central bank must serve, retail and wholesale, want different monies with different implications, one bearing on inflation and the other on stability. And when the supply of safe assets collapses while nominal rates sit at the zero lower bound, the real rate cannot fall far enough to clear the market for safety, so that output itself must adjust, the situation literature calls a "safety trap". A shortage of safe assets is, in this sense, a structural source of fragility rather than a passing disturbance.

3.6.5 Implications for Diamond's Framework

Gorton's synthesis sharpens, rather than displaces, the role that the preceding section assigned to public debt. There, government debt is a store of value that competes with productive capital and can be welfare-improving in a dynamically inefficient economy. Gorton's qualification is that what households and intermediaries demand is not a store of value in the abstract but a safe one, an information-insensitive claim that the private sector cannot perfectly replicate, since private collateral, unlike a sovereign's promise, is never quite riskless. Two consequences follow for the picture developed earlier. First, the relevant scarcity is not of any asset that bridges the present and the future, but of safe such assets, and the convenience yield is the price of that scarcity. Second, when the sovereign under-supplies safe debt, the gap is not closed by a smooth movement of the interest rate alone, as a frictionless model would have it, but by the manufacture of fragile private substitutes whose failure is a systemic event. The demand for a claim on the future, in other words, is at bottom a demand for a safe claim on the future and the difficulty of producing one

is what makes its supply a matter of macroeconomic, and not merely individual, consequence.

Throughout, Diamond's national debt has been a store of value with an issuer behind it: a claim the old sell to the young and the state rolls over without end, set against capital in the contest for each generation's savings. In the over-accumulating economy, the dynamically inefficient case $r_E < n$, the crowding out worked by internal debt, a cost where the economy is efficient, becomes instead the instrument that lifts the interest rate back toward the Golden Rule. What the argument has nowhere required is that the issuer be a government. Whether the same office can be discharged with no issuer at all, by an intrinsically valueless asset that rational agents agree to hold on the strength of nothing but the belief that a later generation will buy it in turn, is the question of the chapter that follows, where the inequality $r_E < n$, read here as the welfare diagnostic of the over-accumulating regime, returns in a sharper role still: as the exact condition under which such an asset can command a positive price at all.

Chapter 4 - When Bubbles Enhance Welfare

4.1 The Role of Tirole

The argument of this thesis has so far advanced through a single idea examined in two settings. Samuelson (1958) showed that money can command a positive price in a pure-exchange economy of overlapping generations, because each young cohort is willing to hold it in the rational expectation that the next cohort will do the same. Diamond (1965) carried the same logic into an economy with production and capital, where public debt rolled over from one generation to the next plays an analogous role and, in doing so, crowds out productive capital. Jean Tirole's "Asset Bubbles and Overlapping Generations" (1985) is the point at which this idea reaches its most general form. Tirole asks whether the phenomenon Samuelson isolated for money survives in the richer, Diamond-style world of capital accumulation, and crucially, whether it survives for any intrinsically valueless asset, not merely for money.

This is the question that closes the circle. Samuelson's money and Diamond's debt are, in Tirole's treatment, two instances of a single object he calls a pure bubble: a claim that pays no dividends and whose entire market value is therefore unsupported by any underlying stream of returns. By a simple no-arbitrage argument, such a claim can be held only if its price is expected to grow at the real rate of interest, since it must offer savers a return competitive with capital. Tirole's model is built directly on Diamond's, and it contains both predecessors as special cases: strip away capital and one recovers Samuelson's pure-exchange economy; remove the bubble and one returns to the Diamond economy of capital accumulation studied in the previous chapter. The contribution is therefore not a new mechanism but a unification of the two with a precise statement of the conditions under which the mechanism operates. One element, however, is genuinely new: in a production economy the bubble feeds back on the very interest rate it is obliged to pay, since the savings it absorbs are savings withdrawn from capital.

Tirole frames the problem in the language of asset pricing. Standard finance theory holds that the price of an asset should equal its market fundamental, the present discounted value of the dividends it pays; on this "fundamentalist" view, an intrinsically useless asset must be worthless. Tirole had himself shown, in an earlier contribution (Tirole, 1982), that in an economy with a finite number of infinitely lived traders this view holds exactly: no asset can trade above its fundamental. His 1985 paper asks whether the same holds under overlapping generations and finds that it does not. Because new cohorts of savers arrive on every date, an asset can be passed forward indefinitely, and its price can rest entirely on the expectation of resale to those not yet born. A bubble, the gap between an asset's price and its fundamental, is thus not a market failure or a symptom of irrationality, but a

possible feature of competitive equilibrium among rational agents.

What makes Tirole the natural endpoint of this thesis is the sharpness of his central result. He shows that a pure bubble can persist in equilibrium precisely when the bubbleless economy is dynamically inefficient, i.e., when its long-run rate of interest falls short of the rate of population growth (in the notation of the previous chapters, $r_E < m$). This is exactly the regime in which Samuelson's competitive equilibrium fails to reach the social optimum and his "social contrivance" becomes welfare-improving; it is the regime that the simulation in this thesis reproduces; and it is the same over-accumulation case that Chapter 3 identified in the Diamond economy. Where the economy saves too much, and capital is over-accumulated, a bubble absorbs part of that excess saving, raises the return to capital, and moves the economy toward the Golden Rule, turning a dynamically inefficient allocation into an efficient one. Tirole thus supplies the general principle toward which the earlier chapters were building: in an overlapping-generations economy of this kind, whenever the long-run interest rate lies below the growth rate, an intrinsically valueless asset can be priced, and its existence carries the economy toward efficiency.

4.2 The Bubbly Economy

Tirole conducts his analysis in the same two-period overlapping-generations economy as in the previous chapter, and introduces a single new element. Alongside productive capital, the young are now permitted to hold a third kind of asset, which Tirole calls a pure bubble. A pure bubble is a durable, intrinsically worthless claim available in fixed supply: a piece of paper that pays no dividend in any period and confers no direct utility. It is, in effect, Samuelson's money under another name: an asset whose only service was ever to be passed from one generation to the next, transplanted into a production economy. Tirole adopts the standard financial definition under which the bubble on an asset is the difference between its market price and its market fundamental, the present discounted value of the dividends it pays. Because a pure bubble pays nothing, its fundamental value is zero, and its price consists entirely of the bubble. As in the earlier models, such an asset can sustain a positive price only because each younger generation is willing to buy it from the older generation in the expectation of reselling it to the next.

What distinguishes Tirole's environment from its predecessors is that the bubble must now compete with capital for young people's savings. Since a saver may hold either asset, equilibrium requires that the two offer the same return; as the bubble pays no dividend, its entire return takes the form of a capital gain, and its price must therefore grow at the real rate of interest. This single requirement carries a

structural consequence that shapes the whole analysis. Each generation's savings are finite and must be divided between productive capital and the bubble. Hence, any value the market places on the bubble is value drawn away from capital: the two are rivals for one pool of saving, and the bubble's worth can never exceed what the young can set aside.¹ The economy therefore nests both of its forebears as limiting cases. Remove capital and one recovers Samuelson's pure-exchange world, in which the bubble is the only store of value; remove the bubble and one returns to the Diamond economy of the previous chapter. Tirole's framework sits above both, and the question it poses is whether a pure bubble can command a positive price once capital is available as an alternative store of value.

4.3 The Existence Condition

The previous section closed on a question: can a pure bubble command a positive price once capital stands ready to absorb the savings of the young? Tirole's answer is Proposition 1, the central result of the paper and the pivot of this chapter. Its organising object is the long-run interest rate of the bubbleless economy, the rate to which Diamond's equilibrium converges when the bubble is absent, which Tirole denotes $\bar{r}$ and which is none other than the rate written r_E in the previous chapter. The headline is a single comparison: bubbly equilibria exist if and only if $\bar{r}$ falls short of the population growth rate m ; thus if and only if the bubbleless economy is dynamically inefficient.² The inequality that has so far served this thesis as a welfare diagnostic: marking the regime in which Samuelson's social contrivance improves on the market, and the over-accumulation case identified in the Diamond economy of Chapter 3, returns in a second role. It is now an existence criterion: it separates the economies in which intrinsically valueless assets can be priced from those in which the fundamentalist view of asset pricing survives intact.

It is worth beginning where Tirole begins, with the argument that seems to rule bubbles out altogether. By the no-arbitrage logic of the previous section, the bubble's value must compound at the rate of interest, while the means of those who must buy it are finite at every date. If the asset's value grows without end, there must eventually come a generation for which it costs more than everything that generation can save; and an asset that must one day find no buyer should, working

¹Tirole's full model is richer than the version employed here. It also admits assets that pay real dividends ("rents"), and it goes on to examine institutional backing and reserve requirements, the substitution of one bubble for another, stochastic bubbles, and the complete transition dynamics toward the steady state. The present treatment restricts attention to the steady-state behaviour of a single pure bubble, which is sufficient to establish the existence and welfare results on which the argument of this thesis depends.

²Tirole sets aside the nongeneric boundary cases $\bar{r} = m$ and $\bar{r} = 0$, and this chapter follows him.

backwards, never find a first one. Tirole grants the argument its grain of truth: in an economy that does not grow, it is essentially correct, and bubbles are impossible.³ What it overlooks is growth. The resources available to carry the bubble are the young's savings, which expand at rate m with the population. The race that decides the bubble's fate is therefore not between a compounding price and a fixed purse, but between two growth rates: the rate of interest, at which the bubble is condemned to grow, and the rate m , at which the economy's capacity to absorb its growth.

The basic argument fails in a second, more nuanced way by treating the interest rate as an exogenous variable rather than something determined within the system. This is where factoring in the real production economy becomes essential. In Tirole's economy, the interest rate is a price, and the bubble itself moves it. Every unit of saving the young place in worthless paper is a unit withheld from capital; a smaller capital stock carries a higher marginal product; and a higher marginal product is a higher rate of interest. The bubble thus raises the very rate at which it is obliged to grow, the feedback loop announced at the opening of this chapter, and this feedback is what gives the existence condition its exact shape. Suppose first that the bubbleless economy is efficient, $\bar{r} > m$. A bubble can only push the interest rate further above m ; condemned to grow faster than the economy as a whole, it must one day exceed the entire saving capacity of the young. Wealth constraints rule out such paths, and under perfect foresight, the exclusion operates from the very first date: no equilibrium can begin with a positive bubble. This is Proposition 1(a): in the efficient region the bubbleless equilibrium is the only equilibrium, and the fundamentalist view is valid, exactly as in the world of infinitely lived traders. Suppose instead that $\bar{r} < m$. Now there is room. A bubble of the right size absorbs just enough saving to lift the interest rate to m and no further; an asset growing at rate m in an economy growing at rate m claims a constant share of each generation's savings forever, and nothing in the arithmetic of compound interest ever forces a reckoning. The bubble is sustainable, and the question becomes not whether such an equilibrium exists, but how many do.

Tirole's answer for the intermediate case, $0 < \bar{r} < m$, is exhaustive. The equilibrium set is indexed by the size of the initial bubble and organised by a ceiling: there exists a maximum feasible initial bubble, larger naturally the more capital the economy starts with; beyond which no equilibrium can be constructed. Initial bubbles above the ceiling are impossible: compounding at an interest rate that their own crowding-out pushes above m , they would come to absorb more than the young can save and

³This loose argument was made rigorous by Scheinkman, as acknowledged by Tirole, and its validity is confined to economies that do not grow.

would drive the capital stock below zero in finite time, so that perfect foresight excludes such paths from the outset. At the ceiling exactly, and only there, runs the knife-edge path on which the bubble endures. Along it capital is held permanently below its bubbleless path, the interest rate is carried up to m , and the bubble per capita settles at its Golden Rule value, i.e., the level that, at Golden Rule capital, exactly fills the gap between what the young wish to save and what the economy needs to invest. Everywhere below the ceiling lies the third configuration: a continuum of equilibria. On these paths a bubble exists at every date, and its aggregate value even grows without bound, since it compounds at a positive rate of interest; yet its real-world impact actually fades away. Per capita, a population growing faster than the bubble itself shrinks toward zero, its weight in the economy vanishes, and the system drifts back to the Diamond steady state of the previous chapter; the path with no initial bubble at all is, of course, the Diamond equilibrium itself.

This structure is what Tirole's vocabulary is built to capture, and the distinction it draws is the one precision the headline result requires. An equilibrium is bubbly: if the bubble is positive at some date, equivalently, at every date, since under perfect foresight a bubble can neither spring from nothing nor collapse exactly to zero; it is asymptotically bubbly if the bubble per capita does not vanish in the limit. In this vocabulary, Proposition 1 says the following: bubbly equilibria exist if and only if $\bar{r} < m$, the clean sense in which dynamic inefficiency is necessary and sufficient for bubbles, but persistence is rarer than existence. In the intermediate case, a continuum of bubbly equilibria coexist, and on all but one of them the bubble is transient in per capita terms; only the knife-edge path is asymptotically bubbly. Two consequences deserve to be recorded. First, in the intermediate case the bubbly equilibrium is never the equilibrium: the bubbleless Diamond path coexists with the whole continuum of bubbly ones, and nothing in fundamentals selects among them; the choice is a matter of self-fulfilling expectations, the same indeterminacy that has accompanied this class of models since Samuelson. Second, when the opening section of this chapter spoke of a bubble that persists, the claim intended was the asymptotic one, and it is exact: an equilibrium in which the bubble never fades exists if and only if $\bar{r} < m$, and in the intermediate case there is exactly one.

There remains the third case, and it is the most striking. If capital depreciates, sufficiently severe over-accumulation can drive the bubbleless long-run rate below zero and here the bubble changes status, from possible to necessary. The result is the one point of this chapter at which the rent-bearing assets set aside in the previous section must briefly reappear, for it concerns the interaction between bubbles and

rents precisely.⁴ Let the economy contain any asset paying a real dividend, however modest. Were the bubble absent, savings would in the long run flow almost entirely into capital, and the interest rate would settle at $\bar{r} < 0$; but at a nonpositive rate of interest the present value of even the smallest perpetual dividend is unbounded, and no finite saving could purchase the asset that pays it. A bubbleless equilibrium therefore cannot exist, and Tirole shows that the same contradiction eliminates the asymptotically bubbleless paths as well. Exactly one equilibrium survives: the asymptotically bubbly one, on which the bubble, by holding the interest rate up at m , keeps the value of rent-bearing assets finite. In this configuration, the intrinsically valueless asset is no longer an ornament the economy may or may not wear; it is a precondition for the price system to clear at all.

The lineage of the result measures the distance travelled since the first chapter. Samuelson established that an intrinsically useless asset can be valued under overlapping generations, but in his economy money was the only store of value, and the suspicion lingered that the phenomenon would not survive competition from a real asset. Wallace (1980) answered the suspicion halfway: in a growing economy in which consumers may store the real good, he showed that equilibria with valued fiat money exist precisely when the population growth rate exceeds the return on storage; money survives the competition wherever storage cannot keep pace with the economy. Tirole explicitly credits the paper, and his advance over it is the one this chapter has been tracing. Replace storage, whose return is a constant of technology, with neoclassical capital, whose return falls as the stock accumulates, and Wallace's comparison between two fixed numbers becomes the equilibrium feedback described above: the bubble now acts on the rate it must match, and the whole adjustment of capital, wages, and interest along the path can be traced. This is also what brings the question of welfare within reach. In Tirole's economy, the bubble does not merely coexist with capital; it determines how much capital the economy accumulates, and therefore whether the economy accumulates too much. Thus, what the crowding out of capital does to efficiency, i.e., Proposition 2, and the return to the Golden Rule is the subject of the next section.

4.4 The Return to the Golden Rule

The existence result settles the question of whether a bubble can live; it does not say whether the economy that carries one is better or worse with it. An asset that absorbs saving which would otherwise have become capital looks like a parasite: the

⁴Tirole proves the necessity result for an economy containing an arbitrarily small rent; in the strictly rentless abstraction the dividing line for existence remains $\bar{r} < m$. The case is included here because no actual economy is rentless: land alone is enough.

young work no less and produce no more for holding it, yet part of their savings is buried in paper from which nothing will ever grow. Proposition 2, the paper's welfare result, shows that in the only region where bubbles can exist, this instinct has matters backwards, and the reason has been in place since Chapter 3. When $\bar{r} < m$ the bubbleless economy over-accumulates: it maintains a capital stock so large that the marginal unit returns less than what it costs merely to keep capital per head intact as the population grows, so that a portion of every generation's saving is swallowed by the upkeep of capital that does not repay its keep. Such an economy can consume more on some dates without consuming less on any other, simply by holding less capital. The free lunch is there for the taking; what the market lacks is an instrument for disposing of its excess accumulation. A pure bubble is that instrument.

The mechanism is the one that shaped the existence condition in the previous section; what changes is the reading. Along the asymptotically bubbly path, the bubble absorbs part of the savings of the young, capital is held below its bubbleless path, and the interest rate is carried up from $\bar{r}$ to m : the economy converges to the Golden Rule, the capital stock at which consumption per head is the largest the technology can permanently sustain. Beneath this asset-market description, the bubble is operating an intergenerational transfer. Each young cohort gives up part of its wage to the current old in buying the paper, and is repaid in old age by the cohort that follows; resources flow from young to old at every date, without end. This is Samuelson's contrivance transplanted into a production economy, and at the Golden Rule it pays Samuelson's rate: each generation's contribution is returned by a cohort larger by the factor of population growth. Hence, the chain of transfers yields exactly m : the biological rate of interest whose optimality Chapter 2 examined, and now, exactly, the return on capital. The source of the welfare gain can then be put in a sentence. When $\bar{r} < m$, capital is an inferior vehicle for carrying resources into old age: at the margin it returns $\bar{r}$, while the standing willingness of the next generation to buy offers m . The bubble lets each cohort retire part of the inferior vehicle in favour of the better one, and the savings it "wastes" were savings the economy was already wasting.

Proposition 2: if $\bar{r} < m$, the asymptotically bubbleless equilibria are inefficient, and the asymptotically bubbly equilibrium is efficient.⁵ Every path of the previous section on which the bubble fades: the Diamond equilibrium itself, and with it the

⁵Efficiency is throughout Pareto efficiency across generations: an allocation is efficient if and only if no feasible alternative improves the welfare of all generations, strictly for at least one of them. Note also what the proposition does and does not say: it classifies equilibrium paths as efficient or inefficient. It is not stated as a welfare ranking of one equilibrium over another, and this chapter claims no more.

whole continuum of transient bubbles; converges to a steady state in which interest stays below growth; the over-accumulation is never undone, and each such path fails the Pareto test. The knife-edge path, the only one on which the bubble endures, is the only one that reaches efficiency. The proof is omitted here, but whether a path is efficient is decided, by a classical criterion, by whether cumulated interest keeps pace with cumulated growth; and that race is precisely the one that decides⁶ the fate of the bubble per capita, which compounds at the rate of interest while the population that must carry it grows at m . The bubble per capita is thus the economy's own efficiency statistic: a path on which it dies out is a path on which interest has lagged behind growth, over-accumulation, while a path on which it survives is a path that passes the test. Existence and welfare are two readings of the same number. Set beside Proposition 1, this yields the alignment that gives the chapter its point: bubbles can exist exactly where they are wanted. Where the bubbleless economy is efficient and needs no correction, no equilibrium can carry a bubble at all; where it over-accumulates, bubbles can arise, and the only one that endures is the only one that heals. In the extreme case $\bar{r} < 0$ the alignment becomes total, for there the asymptotically bubbly path is the sole equilibrium of any kind: the possible and the desirable coincide. In the intermediate case, however, the welfare theorem comes with the same condition as the existence theorem: the efficient path is available, not assured. Nothing in fundamentals selects it, and whether the economy finds its way to the Golden Rule rests, as it has since Samuelson, on the expectations its agents coordinate upon.

It is Tirole himself who names the connection this thesis has been assembling. Bubbles, he observes, move resources from each young cohort to its elders, and are in that sense "the private counterpart to national debt à la Diamond". The identification holds to the letter. The rolled-over public debt of Chapter 3 was a claim sold by the old to the young and carried forward, period after period, without ever being extinguished; it crowded out capital; and in the over-accumulating economy that crowding out was its service, moving the economy toward the Golden Rule. The pure bubble does each of these things line for line, and Tirole, who acknowledges that his welfare analysis owes much to Diamond's treatment of national debt, is careful about the small print of the analogy: strictly speaking, government debt carries no bubble, since each bond is priced at its fundamental, namely the repayment it promises; debt acts as a bubble because the promise is kept bond by bond

⁶The criterion is Cass's (Cass, 1972); Tirole's appendix invokes it in the form of Theorem 5.6 of Balasko and Shell (1980), extended from pure exchange to production. Strictly, the general test reads the limiting behaviour of cumulated interest over cumulated growth, not the bare fact that the ratio reaches zero; on the equilibrium paths at hand the two coincide, because the ratio — which is, up to its initial value, the bubble per capita itself — either decays geometrically or converges to a strictly positive limit. Hence the brevity of Tirole's proof.

and broken in the aggregate, every redemption financed by a new issue, so that the debt as a whole is never repaid. What separates the two instruments is therefore not their economics but their authorship. Diamond's debt must be issued and rolled over by a government; Tirole's bubble needs no issuer at all, arising spontaneously from nothing more than the shared willingness of private agents to treat a worthless object as a store of value. The social contrivance of Chapter 1, it turns out, does not require anyone to contrive it.⁷

With Propositions 1 and 2 in place, existence has been pinned to dynamic inefficiency, persistence to a single knife-edge path, and that path to the Golden Rule; the asset Samuelson called money, the debt Diamond rolled over, and the pure bubble of this chapter have been given one existence condition and one welfare role.

4.5 Three Faces of One Phenomenon, and What Sustains Them

The argument of this thesis can now be assembled in one place. Money in Samuelson, rolled-over public debt in Diamond, and the pure bubble in Tirole are three names for a single object: an intrinsically valueless claim that commands a positive price because each generation rationally expects the next to buy it, that can be priced precisely when the economy is dynamically inefficient, and that, where it exists, raises welfare by carrying out the transfer from young to old which the market alone cannot arrange. The three differ only in dress: Samuelson's asset lives in a pure-exchange economy and is called money; Diamond's lives in a production economy and is issued by a state; Tirole's lives in the same production economy and is issued by no one. What Tirole contributes to the synthesis is generality. The phenomenon is not a peculiarity of money, nor of economies without capital, nor of fiscal policy: it is a property of overlapping-generations economies as such, switched on and off by a single inequality. The condition $\bar{r} < m$ has now served this thesis three times: as the regime in which Samuelson's contrivance improves on the market, as the over-accumulation case of the Diamond economy, and as Tirole's existence criterion; and the lesson of this chapter is that these are not three results that happen to share a formula. They are one result: wherever the social contrivance is worth having, a bubble can exist; wherever it is not, the fundamentalist view stands, and no such asset can be priced at all.

It remains to certify the reading of Samuelson on which the identification rests, and

⁷The sentence quoted opens a section in which Tirole qualifies this favourable picture with two potential inefficiencies of bubbles at the microeconomic level: resources may be wasted creating new units of a bubbly asset, and a bubble may keep an asset's real services from ever being consumed. Both concern the choice of the bubbly asset rather than the steady-state welfare question treated here, and are accordingly set aside.

Tirole himself supplies the certificate. There are, he shows, two ways of giving value to money, and for money (though not for other assets, which may mix fundamental and bubble) the two are inconsistent in the long run. On the fundamentalist view, money is held for the services it renders: it finances transactions, pays taxes, meets a reserve requirement; it must store value to do so, but it is not held for speculation, and Tirole proves that money expected to retain such transaction value cannot carry a bubble. On the bubbly view, money renders no service at all: it is a pure store of value à la Samuelson, held entirely for resale, its price growing at the real rate of interest. Samuelson's money is therefore not loosely analogous to a bubble; it is the bubbly view of money, as Tirole calls it. The dichotomy will earn its keep once more, for it is exactly the question that divides interpretations of Bitcoin: means of payment or store of value, and it frames the chapter this thesis devotes to it.

If almost anything could in principle carry a bubble, Tirole remarks that nothing stops him declaring a childhood drawing of his to be worth a thousand dollars, except that no one would invest in it; the question becomes how actual bubbles are selected, and his closing answer lists three necessary conditions: durability, scarcity, and common beliefs. The asset must be durable, or the chain of resale on which its entire value rests breaks at the first link. It must be scarce in a precise sense: new units must sell at the price of old ones, so that an asset which can be created at will cannot sustain a bubble, which is why bubbles have historically attached to objects that cannot be reproduced: gold, land, famous paintings, rare stamps.⁸ And it must be the object of common beliefs: since fundamentals single out nothing, everyone must expect everyone else to accept precisely this asset among the infinity of candidates; a focal point, which is all that Chapter 1's "social contrivance" ever was. What Samuelson's money possessed by assumption, gold and land possess by history; and the three conditions read, today, like a designer's checklist. Assets engineered to satisfy them: cryptocurrencies, durable and scarce by protocol, their common belief coordinated by the code itself, or non-fungible tokens, which attempt to manufacture uniqueness outright; present themselves as natural candidates. The suggestion is this thesis's extension, not Tirole's, who wrote decades before any of them existed; Bitcoin is examined on these terms in its own chapter, and the conclusions return to the broader question.

Tirole's own conclusion is the right note on which to end, not least because it is modest. He suggests that overlapping-generations theory should shift some of its

⁸Tirole makes the requirement quantitative: the supply of a bubbly asset may grow, but at most at the rate by which the economy's growth outstrips the bubbleless interest rate, $m - \bar{r}$. The bound will matter later in this thesis, for assets whose scarcity is a design parameter rather than a historical accident.

attention from money to assets held for speculation (the movement this thesis performs in miniature) and he concedes what the model cannot yet say: the models and the empirical evidence are too preliminary to tell whether bubbles should actually be observed, and economics has not settled whether the world of overlapping generations or the world of infinitely lived traders (or of generations bound together by bequests) is the right description of ours; a choice on which everything turns since, as the opening of this chapter recalled, bubbles can exist in the first and cannot in the second. What the paper does defend is the proposition this chapter set out to reach: that "bubbles are not inconsistent with optimising behaviour and general equilibrium". The fundamentalist view of asset pricing is not a theorem of rationality; it is a property of worlds in which buyers can run out, have a last date, or face a fixed and finite roster of traders, however long those traders live. In a world where new cohorts arrive without end, perfectly rational agents may pay a positive price for nothing at all, and the economy, when it saves too much, may be the better for it.

Chapter 5 - Computational Evidence from an OLG Simulation

To investigate and verify Samuelson's findings reported in Chapter 2, I replicate the demographic and market dynamics using an Agent-Based Model. The simulation is set up as follows. Following the three-period case from Samuelson (1958), agents are young during the first 25 years of their lives, middle-aged during the second 25 years, and old during the last 25 years; a model period corresponds to 25 years. Each agent is endowed with 1 unit of consumption per year of productive life, so that the lifetime endowment equals 50. The population grows at a constant annual rate, compounded to a per-period (25-year) rate within the simulation. A model period is set to 25 years so that the three generations span a stylised 75-year life cycle, consistent with the standard interpretation of Samuelson; trading occurs once per period because the OLG money mechanisms operate at the generational, not the intra-period, frequency.

For the first generation, preferences for intertemporal consumption are heterogeneous and follow a Beta distribution with $a = 7$ and $b = 3$, yielding $\mu = 0.7$ and $\sigma = 0.1381$. Since each model period corresponds to 25 years, these per-period discount factors imply an annualised mean discount factor of $\mu^{1/25} \approx 0.9859$, equivalent to an annual rate of time preference of approximately 1.43%, with a delta-method standard deviation of roughly 0.0078 across agents, which falls within the conventional range adopted in macroeconomic calibrations, while the Beta family provides a natural parametrisation on $[0, 1]$; the shape parameters $a = 7$ and $b = 3$ concentrate mass away from the boundaries and preserve cross-agent heterogeneity without producing degenerate discount factors.

Middle-aged agents mate with each other in order to form the new cohort of young agents; newborns intertemporal preferences are constructed as follows $\beta_{new} = \frac{1}{4}\beta_{p1} + \frac{1}{4}\beta_{p2} + \frac{1}{2}\beta_{random}$ where $\beta_{random} \sim Beta(7, 3)$, modelling the interplay between cultural transmission and randomness. The utility of each agent is a standard logarithmic utility:

$$U(C_1, C_2, C_3) = \log(C_1) + \beta \log(C_2) + \beta^2 \log(C_3)$$

Given the heterogeneity in preferences, the model allows young agents to borrow against their future income, as explicitly admitted by Samuelson (1958) in footnote 12. However, given the lack of intergenerational debt enforcement in the model and the fact that old agents' balances are zeroed at death, leaving an opportunity for a free lunch, middle-aged agents are not allowed to enter retirement with negative balances, which would imply non-positive old-age consumption. Furthermore, because the log utility requires strictly positive consumption in every period of life, the

model would crash. Given the distribution of preferences, the case where middle-aged agents try to carry debt into retirement occurs in only 0.0008% of simulations, and only in the stationary case, causing the simulation to fail.

The ABM allows for two specifications of expectations about the future real return on money. Under unconditional rational expectations, agents anticipate that the steady-state return equals the per-period population growth rate, consistent with Samuelson’s biological-rate theorem. Under damped adaptive expectations, agents update beliefs from observed prices according to $R_t^e = (1 - \lambda) R_{t-1}^e + \lambda R_{t-1}$, with smoothing parameter $\lambda \in (0, 1)$. The naive rule $R_t^e = R_{t-1}$ ($\lambda = 1$) is excluded because, as Samuelson notes in footnote 11, it generates a cobweb feedback through the young’s saving rule that amplifies small price perturbations period by period, driving the simulation away from the biological steady state. The damped variant offers a middle ground: it preserves genuine out-of-equilibrium learning dynamics, as in any boundedly rational learning rule, while attenuating the cobweb feedback below its instability threshold. This allows the model to be perturbed, for example, by monetary shocks as later in the text and to recover endogenously rather than by assumption.

For the R implementation of the model, refer to Appendix A. All Monte Carlo simulations employed 10,000 repetitions and were run in parallel utilising 68 cores on an Amazon Web Services Compute Optimised Lightsail instance (72 vCPUs, 144 GB RAM). Total wall-clock time across all four scenarios was approximately four hours.

The chapter presents the results from the four scenarios described above, organised so that each successive case adds one additional margin to the previous one, thereby isolating a distinct claim in Samuelson’s 1958 argument. The stationary no-shock case serves as the benchmark: with $m = 0$ and a constant money stock, it verifies that the ABM reproduces the zero-interest equilibrium that Samuelson derives as the solution of his consumption-loan problem. The stationary shock case then probes stability and recovery. Having introduced damped adaptive expectations precisely to allow the model to be perturbed and to recover endogenously, a monetary injection displaces the price level from its path, and the question is whether, and how quickly, the economy finds its way back. The growing no-shock case generalises the benchmark by accounting for a changing population and testing the biological-rate theorem $r = m$, which constitutes Samuelson’s central result. The growing shock case finally combines the two margins and poses the hardest question of the four: whether the biological steady state is itself an attractor under perturbation, or merely a fragile balance that survives only in the absence of disturbances.

Throughout, the sections are kept parallel. Within each section, both rational expectations and damped adaptive expectations are considered, so that the learning dynamics can be read as a departure from the perfect-foresight benchmark rather than as an independent object. All four scenarios are fully reproducible: each Monte Carlo replication is seeded deterministically as $1958 + k$ where k is the simulation index; k in one scenario shares its demographic history (preference distributions, parental mating, newborn draws) with simulation k in every other scenario. This design isolates the effect of the regime change (stationary vs growing, shock vs no-shock) from variation given by randomness.

Before analysing the results of the various scenarios, it is useful to document the generational evolution of the β discount factor; the terminal values are shown in Table 2.

Scenario	N		Mean		St. Dev.	
	gen=0	gen=70	gen=0	gen=70	gen=0	gen=70
Theoretical	–	–	0.7000		0.1381	0.0739
Growing	500	3,087,295	0.7000	0.7000	0.1381*	0.0739*
Stationary	100,000	100,000	0.7000	0.7000	0.1382*	0.0739*

* Pooled across Monte Carlo replications.

Table 2: Within-population summary statistics for intertemporal preference (β) evolution

Given the magnitude of the Monte Carlo simulation (10,000 repetitions), the standard errors on the reported statistics are negligible at the precision displayed. Figure 6 reports the cross-replication dispersion of the aggregate mean ($\bar{\beta}_t$), distinct from within-simulation agent variance, with ribbons marking the 5th and 95th percentiles.




Figure 6: Generational evolution of β

Initialised with 100,000 agents, the stationary economy's mean is tightly bounded near 0.7000 immediately via the law of large numbers. Conversely, the growing economy exhibits significant initial sampling noise due to finite-size effects. As the population grows exponentially over 70 generations, this sampling noise shrinks, and the growing economy's envelope becomes indistinguishable from that of the stationary economy.

Across all Monte Carlo replications, regardless of the demographic regime (growing or stationary), the terminal cross-sectional mean of β converges to 0.7000. This outcome confirms that the designed cultural transmission mechanism is mean-preserving.

The theoretical initial dispersion of 0.1381 compresses to a steady-state value of 0.0739, in close agreement with the analytical $\sigma_\infty \approx 0.07386$ implied by the mating rule. The mating algorithm halves parental variance each generation, while the mutational draw injects fresh dispersion; the two forces balance at σ_∞ . The resulting dispersion is small enough to avoid polarisation into extreme types but large enough to guarantee the persistent survival of a tail of impatient agents, on which the borrowing behaviour documented further in the text depends.

5.1 Stationary Economies

The empirical analysis proceeds by first establishing the stationary population scenario, which serves as the benchmark for isolating the mechanics of intertemporal exchange in the absence of demographic distortions. Under these conditions, the biological-rate theorem states that the socially optimal market rate of interest, in a stationary population, must equal zero (as the growth rate of such a population).

Statistic	Rational expectations	Adaptive expectations
Mean r_{annual}	0.00006	0.00006
St. Dev. r_{annual}	0.00047 ¹	0.00048 ¹
Mean Price Level	2.83171	2.82678
Coef. Var. Price Level	0.00142	0.00335
Mean agg. real savings	441430.91±1.43	442204.90±1.46
Borrower share	0.14258 ²	0.14258 ²

¹ Mean across replications of the within-run time-series standard deviation

² The share of young agents in the final generation whose money holdings are negative

The standard errors when not displayed are negligible at the precision displayed

Table 3: Between and within simulations summary of statistics

As reported in Table 3, a stationary economy of 100,000 agents with a starting money supply of 1,250,000 units of currency, that has been tracked across 70 generations, reaches the biological optimum both under rational and adaptive expectations; thus, the ABM confirms the biological interest rate theorem for a stationary population (Mean $r_{annual} \approx 0$). The share of young agents that optimally choose to borrow is: 14.26%. Because the ABM imposes a no-negative-balance constraint upon entering retirement, this intra-generational borrowing implies that the aggregate demand for money (441, 430) is a net figure. The remaining young agents, whose discount factors are high enough to prefer net saving, must therefore save enough to simultaneously acquire the fixed supply of money from the retiring generation and finance their impatient peers' consumption loans.




Figure 7: Panels for a stationary economy

The transition dynamics depicted in Figure 7 highlight the differing convergence paths dictated by the expectation regimes. Under rational expectations, agents perfectly anticipate the stationary $r_{\text{annual}} = 0$; thus the economy reaches the biological rate within the first active generation. The negligible initial volatility is an artefact of the population initialisation. Conversely, under damped adaptive expectations, the economy undergoes a transition phase. Because agents update their beliefs based on past observed prices, an initial mispricing of the return on money causes aggregate savings to overshoot, temporarily depressing the price level and driving the real return into negative territory. However, the damping parameter ($\lambda = 0.1$) attenuates the explosive feedback loop that Samuelson warned would drive the market toward a degenerate equilibrium. By dampening the recursive volatility of the savings function, the adaptive learning mechanism allows the system to endogenously correct and converge to the rational steady state by approximately generation 30. Furthermore, the simulation obeys the underlying macroeconomic accounting identities in the steady state. The emergent mean price level ($P \approx 2.83$) is the ratio of the exogenous nominal money supply ($M = 1,250,000$) and the endogenous aggregate real savings demand ($M/P \approx 441,430$).

The ABM thus reproduces Samuelson’s core claim that a fixed fiat money supply can serve as the ‘social contrivance’ anchoring a decentralised market at the biological rate. To evaluate the robustness of the biological equilibrium, it is possible to introduce exogenous monetary shocks to the stationary economy. The economy now

experiences four "helicopter drops" of fiat money, each occurring every 14 generations, doubling the total aggregate money supply. This fiat injection is unexpected and distributed exclusively and in proportion to the working cohorts, bypassing the retired generation.

Statistic	Rational expectations	Adaptive expectations
Mean r_{annual}	-0.00151	-0.00152
St. Dev. r_{annual}	0.00574 ¹	0.00579 ¹
Mean Price Level	17.47355	18.31616
Coef. Var. Price Level	0.88632	0.89225
Mean agg. real savings	443309.75±1.43	428160.62±1.46
Borrower share	0.14258 ²	0.16437 ²

¹ Mean across replications of the within-run time-series standard deviation

² The share of young agents in the final generation whose money holdings are negative

The standard errors when not displayed are negligible at the precision displayed

Table 4: Between and within simulations summary of statistics with monetary shocks




Figure 8: Panels for a stationary economy with monetary shocks

The introduction of periodic monetary shocks alters the economy's transitional dynamics, even when the growth rate remains at zero. Based on the summary statistics in Table 4 and the transition paths in Figure 8, the mean price level experiences an aggregate increase relative to the stationary baseline, as illustrated by panel (a).

The periodic doubling of the nominal money supply translates into discrete, step-wise jumps in the log of the price level. Because the newly printed fiat is distributed exclusively to the active cohorts, this sudden adjustment in the price level acts as an "inflation tax." It dilutes the purchasing power of the retired cohort, who hold fixed nominal balances, redistributing real wealth away from the old to the younger generations. The high coefficient of variation in the price level (≈ 0.89 in both regimes) reflects these nominal dislocations.

Although the underlying biological steady-state return is zero, the mean r_{annual} is strictly negative across both expectational regimes. As shown in Figure 8(b), the real return on money plunges into deep negative territory in the generation immediately following each monetary injection (S_1, S_2, S_3, S_4) due to the sudden adjustment in the price level, dragging the long-run statistical mean below the zero steady state.

The divergence between the two expectation regimes is most pronounced in the post-shock recovery paths. Under rational expectations, the expected return is pinned at the steady-state value and does not respond to the injection, so the young's saving rule is unperturbed. The market-clearing mechanism then adjusts the price level to the doubled money supply in a single step; the old generation absorbs the revaluation as one deeply negative reading of r_{annual} , and the real return is back at the biological rate immediately after; there is no protracted recovery because nothing in agents' decisions was displaced. Conversely, under damped adaptive expectations, the sudden collapse in r_{annual} depresses the agents' expected future returns through the backwards-looking update rule.

Under log utility, a decline in the expected gross return reduces the young cohort's saving through a wealth channel rather than a substitution channel. The consumed fraction of lifetime wealth, $1/(1 + \beta + \beta^2)$, is invariant to R^e , since the substitution and income effects on the saving rate cancel exactly. Because the young hold no financial assets when they choose, the return enters their problem only through lifetime wealth $w + w/R^e$. A lower R^e raises the present value of the middle-age endowment, expanding lifetime wealth and crowding out saving. Adaptive agents, whose expected return is depressed after each injection by the backwards-looking rule, accordingly hold persistently lower aggregate real balances, with the gap widening at each shock impact (Figure 8(c)). The borrower share under adaptive rises to 16.44% (an increase of roughly 15.28% over the no-shocks baseline), as impatient agents find it optimal to finance current consumption rather than hold a temporarily depreciating currency. Also, under adaptive expectations, the "memory" of the shock, sustained by the damping parameter ($\lambda = 0.1$), prevents the price level from instantly snapping to its new equilibrium. Because the persistent expectation

of lower returns keeps the current price level momentarily elevated, the economy undergoes a sustained deflationary recovery. This slow climb back to equilibrium effectively spreads the post-shock deflationary pressure over several generations. As observed in Figure 8(b), this manifests as a sustained period of positive annualised real returns during the intervals following the shocks, as the market endogenously dampens the cobweb oscillations and guides the economy back toward the biological steady state.

Figure 8(d) makes this asymmetry explicit: the rational forecast error is negligible except at the impact of each shock, while the adaptive error decays slowly thereafter at a rate governed by $\lambda = 0.1$. This persistence is the channel through which monetary disturbances are transmitted into real allocations, as forecast deviations enter directly into the young cohort's savings decision.

Taken together, the two experiments provide a coherent validation of the ABM as a microfoundation for Samuelson's consumption-loan economy. In the frictionless baseline, decentralised log-utility agents coordinate on the socially optimal allocation identified in the original model, with a fixed stock of fiat operating as the intergenerational social compact that sustains trade at the biological rate. The shock experiment then demonstrates that this equilibrium is locally stable rather than knife-edge: even after four successive doublings of the nominal stock, the real return reverts toward zero, and the speed and shape of that reversion is governed entirely by the informational regime of the agents. The damping parameter ($\lambda = 0.1$) performs the role Samuelson anticipated for any well-behaved updating rule, muting the cobweb oscillations that purely myopic forecasting would otherwise generate and allowing the adaptive economy to reconcile endogenously with the rational benchmark.

What the stationary case leaves unexamined, however, is the sharper and more empirically relevant claim at the heart of Samuelson's argument: in a growing population, the market rate of interest should coincide not with zero but with the biological growth rate m . The following section accordingly relaxes the stationarity assumption, reporting the same diagnostic panels and summary statistics under a geometrically expanding cohort structure, first in the absence of monetary disturbances and subsequently under the same shock sequence examined above.

5.2 Growing Economies

To empirically test Samuelson's biological-rate theorem ($r = m$), the agent-based model incorporates demographic growth. The simulation operationalises this by

initialising an economy with 500 agents and imposing a constant per-model-period population growth rate of $\approx 13.28\%$, resulting in a terminal population of 3,087,295. This corresponds to an annualised growth rate of $m = 0.5\%$, providing a credible baseline for advanced economies (International Monetary Fund, 2026). We evaluate the economy’s convergence to this theoretical biological rate under both rational and adaptive expectation regimes. This convergence is analysed first in an environment free of monetary disturbances and, secondly, under the previously established sequence of unexpected fiat money shocks.

Statistic	Rational expectations	Adaptive expectations
Between and within		
simulations mean r_{annual}	0.00508	0.00508
St. Dev. r_{annual}	0.00072 ¹	0.00072 ¹
Mean absolute deviation		
from theoretical path $\log P_t$	0.0112	0.0094

¹ Mean across replications of the within-run time-series standard deviation

The standard errors when not displayed are negligible at the precision displayed

Table 5: Returns and prices for a growing economy

As shown by Table 5, the ABM confirms Samuelson’s biological interest rate theorem for a growing population, with the rational-expectations and damped-adaptive paths converging, after the transition period, on identical aggregate statistics to the precision displayed.




Figure 9: Panels for a growing economy

The y-axis in panel (b) is truncated to ease visualisation of the transition dynamics, excluding the initial peak at 0.01, given by population initialisation

Theory predicts a per-period deflation rate of $\log(P_{t+1}/P_t) = -\log(1 + m_{\text{period}}) \approx -0.1247$. The cross-replication mean log-price path in Figure 9(a) tracks this benchmark to within roughly 0.01 nats under both regimes. The ABM thus reproduces the deflationary trend implied by Samuelson for the growing case.

As shown by Figure 9(b and d) the economy undergoes an initial settling period, this period is an initialisation effect rather than a property of the equilibrium: the initial age distribution and the initial real money balance do not sit exactly on the balanced growth path, and the economy requires several generations of births and price adjustment before the cross-sectional age profile and reach their steady-state configuration, but at around generation 30 the economy stabilises, with $r_{\text{annual}} \approx 0.005 = m$ and the expectation error approaches zero.

Panel (c) documents aggregate real savings that grow approximately exponentially, as the savings market scales with population. This is the mechanical counterpart of the deflationary trend in panel (a): fixed nominal money combined with exponentially growing real demand for balances yields steady deflation.

We now combine the two departures from the stationary benchmark studied so far by subjecting the growing economy to the same sequence of unexpected fiat-money shocks. This scenario provides the sharpest test of the model: it asks whether the biological-rate attractor survives monetary perturbation, and whether the equiva-

lence between expectation regimes documented above breaks down once agents must forecast in a non-stationary environment.

Statistic	Rational expectations	Adaptive expectations
Between and within simulations mean r_{annual}	0.00350	0.00349
St. Dev. r_{annual}	0.00589 ¹	0.00593 ¹

¹ Mean across replications of the within-run time-series standard deviation

The standard errors when not displayed are negligible at the precision displayed

Table 6: r_{annual} for a growing economy with shocks




Figure 10: Panels for a growing economy with monetary shocks

The point estimates in Table 6 are notable for their proximity: the grand mean of r_{annual} differs between the two regimes by only the fifth decimal place (0.00350 versus 0.00349). Two implications follow. First, even in a non-stationary environment subject to recurrent structural shocks, agents who update their forecasts slowly produce aggregate statistics that are indistinguishable from those of agents with rational forecasts. The regime equivalence documented in the preceding sections therefore extends to the hardest test the model admits. Second, the grand mean sits roughly thirty per cent below the biological rate of 0.005 recovered in the undisturbed growing economy: the gap is the signature of the inflation tax discussed below, and it applies symmetrically across the two regimes.

The log price path in panel (a) inherits the deflationary trend of the undisturbed growing economy: because the cohort size expands geometrically while the money stock is augmented only intermittently, P_t falls along each interval $[S_k, S_{k+1}]$, with a discrete upward displacement at each shock corresponding to the real dilution of outstanding balances. Aggregate real savings in panel (c) track this process almost invisibly: the growing denominator P_t absorbs each nominal injection smoothly, so that the shocks register at most as slight inflections rather than as the abrupt discontinuities seen in their stationary analogue.

Panel (b) shows the realised annualised return oscillating around the biological rate within the narrowband isolated by the piecewise-linear transformation. The excursions at each S_k are asymmetric in sign: a pronounced downward spike, reflecting the capital loss borne by the old generation when the price level jumps, followed by a recovery distributed over the two or three subsequent periods. The contrast with case 2 is instructive. In the stationary economy recovery converged back to $r \approx 0$, so that the shock-induced displacement was, over the horizon, a mean-zero deviation around the equilibrium; here recovery proceeds toward a strictly positive rate.

Under rational expectations, the forecast error remains inside the innermost band in every non-shock generation: the equilibrium price motion is fully anticipated, and agents' only mistake is the single-period innovation at each S_k , which is immediately absorbed. Under damped adaptive expectations with $\lambda = 0.1$, the error accumulates and decays slowly over time, producing the long asymmetric plumes visible after each shock. These pronounced individual-level errors leave no detectable trace in the grand mean of Table 5, which reflects the role of market clearing: misforecast returns produce suboptimal savings choices, but the equilibrium price path that clears the market still delivers a real return whose central tendency is pinned down by the demographic structure. However, under both expectational regimes, the mean r_{annual} is influenced by the 4 periods of deep negative real returns due to the "inflation tax" from monetary expansion. With four shock generations in a run of seventy, and with non-shock generations reverting to a real return close to the biological benchmark of 0.005, the implied annualised return inside the shock periods is of the order of ≈ -0.021 , consistent with the capital loss on nominal balances when M doubles.

Considered alongside the three preceding scenarios, this final case supports a strong reading of Samuelson's result. The biological-rate attractor is robust along two independent dimensions: it survives the introduction of monetary shocks and the replacement of rational with damped adaptive expectations. The double robustness is a stronger statement than either dimension alone: it suggests that the equilibrium

described in Samuelson (1958) is not merely the rest point of an idealised planner’s problem, but the long-run behaviour of a decentralised economy populated by bounded forecasters and hit repeatedly by unanticipated monetary interventions.

5.3 Remarks

The analysis above is deliberately scoped to the four scenarios that map one-to-one onto the claims of Samuelson (1958), and several of its modelling choices admit extensions that would sharpen the diagnostics without altering the underlying message. Three in particular seem worth recording.

First, the shock protocol adopted throughout, four doublings of the nominal stock spaced every fourteen generations, was chosen so that the monetary regime would recur often enough to accumulate meaningful statistics within a run of seventy periods. The cost of this choice is that the recovery from shock S_k is not necessarily complete before shock S_{k+1} arrives, particularly under damped adaptive expectations where the forecast error decays at the geometric rate $(1 - \lambda) = 0.9$ per period. A natural follow-up experiment would rerun the shock case with a single injection, say, at generation 30, and let the economy evolve undisturbed for the remaining 40 periods. Such a design would isolate the impulse-response function of the model, permit a clean estimate of the half-life of the forecast error and of the real return, and render the transition dynamics legible in a way that the superposition of four recovery episodes partly obscures. Varying the magnitude of the single shock, a ten-fold injection against a doubling, for instance, would further trace how the recovery horizon scales with the size of the perturbation, a question the present design cannot address.

Second, the stationary and growing scenarios are initialised asymmetrically: 100,000 agents in the stationary case, compared with 500 in the growing case, the latter expanding to roughly 3 million by generation 70. The asymmetry is a concession to the arithmetic of exponential growth; seeding the growing economy at 100,000 would produce a terminal population of hundreds of millions and a Monte-Carlo compute budget far from the four hours reported above, but it does mean that the early-period dispersion in the two regimes is not directly comparable, as Figure 6 makes plain. A more symmetric design would initialise both regimes at a common population size large enough to swamp finite-sample noise from the first generation, which would, in turn, sharpen the comparison of transition dynamics and allow the early-generation behaviour of the two demographic regimes to be read off the same scale. The obstacle is purely computational: a growing economy seeded at 100,000 and simulated under the present mating and pricing rules requires storage and solver

costs that scale with the terminal cohort size, and the Monte Carlo multiplier of 10,000 replications becomes binding very quickly. The experiment is feasible in principle, but would require either a coarser replication count, a shorter horizon, or a reimplementaion in which the price fixed-point is solved on a representative-agent aggregate and the cross-sectional distribution is updated only where it enters the clearing condition.

Third, several modelling choices have been held fixed throughout and could usefully be varied. The smoothing parameter λ is set to 0.1 because it lies safely inside the stability region of the cobweb map; a sensitivity analysis across $\lambda \in (0, 1)$ would document how the recovery horizon and the amplitude of the adaptive plumes scale between the rational ($\lambda \rightarrow 0$) and the naive ($\lambda = 1$) limits. The preference distribution is fixed at $\text{Beta}(7, 3)$, and while the borrower share is empirically well-identified under this choice, the qualitative results, in particular the equivalence of rational and adaptive regimes in the grand mean, deserve to be checked against alternative preference densities and against CRRA utilities with curvature different from unity. Finally, the "rational" regime implemented here is best described as steady-state perfect foresight: agents correctly anticipate the biological attractor but treat the monetary injections as outside their information set. A fully rational treatment that priced in the shock process as a known stochastic law would be a meaningful refinement, although it would move the analysis away from the minimal extension of Samuelson (1958) that the present design was chosen to preserve. Taken together, the three items above outline a research agenda rather than a catalogue of objections: the simulation reported in this chapter is a testbed for the specific claims of the 1958 paper, and the proposed extensions would build on that testbed rather than displace it.

Ultimately, the present agent-based implementation successfully corroborates the core analytical propositions advanced by Samuelson (1958). In undisturbed baseline environments, the simulated market-clearing real interest rate converges strictly to the underlying population growth rate, thereby explicitly validating the biological interest rate theorem in both stationary and exponentially growing demographic regimes. Crucially, the introduction of periodic nominal injections demonstrates that this biological rate serves as a robust long-run anchor for the economy, especially under damped adaptive expectations. Even when the system is displaced by significant monetary shocks and governed by boundedly rational learning dynamics, the intertemporal terms of trade invariably regress toward this fundamental demographic attractor once the transient forecast errors dissipate.

Chapter 6 - Is Bitcoin Samuelson Money?

The simulation reported in the previous chapter delivered a sharp result. In a frictionless overlapping-generations economy populated by decentralised log-utility agents, a fixed stock of intrinsically worthless tokens sustained the biologically optimal allocation, with the market rate of interest converging to the population growth rate under both rational and damped-adaptive expectations. The mechanism, faithful to Samuelson, required no enforcement technology, no fundamental backing, and no central coordination. Each cohort accepted the token solely on the expectation that the succeeding cohort would do the same. The Pareto-improving equilibrium was therefore self-enforcing, secured by a forward-looking belief structure rather than by any property of the asset itself.

It is precisely this configuration: fixed supply, absent intrinsic value, decentralised issuance, and value resting entirely on intergenerational acceptance; that Bitcoin presents itself as embodying. The protocol caps the terminal stock at 21 million units, denies any claim on real resources or sovereign liabilities, and dispenses with the state altogether as an enforcement mechanism. Its valuation, in the proponents' own framing, is sustained by nothing more than collective forward-looking belief. Bitcoin appears, at least in terms of formal preconditions, as a real-world instance of the social contrivance Samuelson had described in the abstract. This convergence between the analytical structure and the empirical artefact poses the organising question of the present chapter. Does the OLG framework justify Bitcoin as money? And, conversely, if Bitcoin satisfies the formal preconditions of Samuelson money, why does it nevertheless fail in practice on virtually every empirical metric of monetary performance?

To address the question without sliding into either advocacy or dismissal, the chapter proceeds methodologically. After steelmanning the OLG case for Bitcoin, the analysis evaluates the cryptocurrency against the three classical functions of money: medium of exchange, store of value, and unit of account, following Yermack's framing. Throughout, the OLG lens remains active: each empirical shortfall is read not only as a failure relative to the functional benchmark, but as evidence of what the model must abstract away in order to deliver its result.

6.1 Bitcoin as Samuelson Money

To pose the question honestly, one must first state the strongest case for Bitcoin as money. The argument, when reconstructed in formal terms, is not marginal. Samuelson's 1958 result is structurally permissive: it establishes that an asset without intrinsic value, without state backing, and without any enforcement technology can nonetheless sustain a Pareto-improving equilibrium, provided successive gener-

ations coordinate on a shared belief in its future acceptance. The claim of Bitcoin's defenders, stripped of its rhetorical framing, is that the cryptocurrency satisfies precisely those preconditions. This section reconstructs the parallels with deliberate generosity.

Samuelson's fiat token is, by design, a social contrivance: a unit possessing no intrinsic value, no productive use, and no claim on any real resource. The model deliberately abstracts from any storage technology that might compete with the token as a vehicle for intertemporal transfers, and from any sovereign authority that might mandate its acceptance. Bitcoin meets these specifications cleanly. A bitcoin is an entry on a distributed ledger; it confers no right to a future cash flow, embodies no liability of any issuer, and grants no claim on any real or financial asset (Yermack, 2015). Its sole economic content is the protocol's guarantee that the entry exists and is transferable. Like Samuelson money, its valuation, if positive, must be sustained by something other than fundamentals.

The supply rule reinforces the analogy. In Samuelson's analysis, the stock M is fixed by assumption, and the elegance of the biological-rate result depends on this immunity to discretionary inflation: a central bank with latitude to issue at will would, in principle, dislodge the equilibrium the model identifies. Bitcoin's protocol performs the work that Samuelson's assumption performs. The terminal supply is hard-coded at 21 million units, the halving mechanism sets the issuance schedule, and any modification requires near-universal consensus among independent operators. Whatever else Bitcoin is, it constitutes a closer institutional approximation to a fixed-stock token than any historical fiat currency has been.

Issuance itself is decentralised. New units enter circulation through proof-of-work mining, an open process whose participation is not contingent on sovereign permission. There is no central authority responsible for monetary policy, no balance sheet against which bitcoin claims are issued, and no jurisdictional perimeter within which acceptance is mandated. Samuelson's model is explicit that the social contrivance functions "without any central planning or explicit legislation"; Bitcoin's institutional design realises this absence of centralised authority more fully than the fiat moneys Samuelson was theorising about.

These correspondences are not incidental. Together, they constitute the conditions under which Samuelson's result obtains. The proof requires only that the asset be intrinsically valueless, that its stock be fixed (or at least follow a rule whose implied rate of return does not dominate the population growth rate), and that successive generations coordinate on accepting it. No further property is asked

of the asset; in particular, none about the institution that issues it. The OLG framework establishes, and this is its analytical force, that money is possible without the state. This is a stronger claim than is sometimes recognised. Pre-Samuelsonian intuitions grounded the value of fiat money in legal-tender laws, in tax acceptability, or, in metallist accounts, in residual convertibility into a real commodity. Samuelson dispenses with all of these. The valuation of money in his model is supported by nothing beyond the expectation that future cohorts will accept it on the same terms. Once the argument is taken seriously, the conceptual space for a privately issued, decentralised, intrinsically worthless monetary asset opens. Bitcoin occupies this room exactly: it is not an exotic special case but the textbook configuration sanctioned by the OLG framework.

6.1.1 The Informal Arguments Translated

The public case for Bitcoin rarely cites Samuelson by name. The dominant rhetorical registers are libertarian and technological: "digital gold", "sound money", "freedom money", "the separation of money and state". Yet each of these formulations, when translated into the analytical vocabulary of monetary theory, restates a corollary of Samuelson's result.

The "digital gold" framing identifies Bitcoin as a non-sovereign store of value whose supply discipline is institutionally guaranteed; the Samuelsonian counterpart is the fixed-stock token that secures positive purchasing power across generations through a pre-committed supply rule. The "free-market money" framing, articulated most explicitly in Hayek's later work on the denationalisation of money, argues that monetary equilibrium can emerge without a sovereign monopoly (F. Hayek, 1976); Samuelson's 1958 model is, on the most generous reading, a formal demonstration of this. The "separation of money and state" argument is the political reading of the same structural claim. In each case, what is presented as an ideological commitment rests on a wager that the conditions of the OLG framework are sufficient to deliver a monetary equilibrium in practice. This translation is not intended to flatter public arguments or to attribute to their proponents a sophistication they do not always claim. It is intended only to identify their analytical anchor: when the most rigorous version of the case for Bitcoin is reconstructed, it does not invoke any feature of the cryptocurrency that the OLG framework would deny to viable money.

The steelman case can be stated compactly. Bitcoin is an intrinsically valueless, decentralised, fixed-supply asset whose value, if positive, is sustained by collective, forward-looking acceptance. Samuelson's framework establishes that such an asset can, in principle, support a Pareto-improving monetary equilibrium without state

backing. Bitcoin, therefore, satisfies the formal preconditions of Samuelson money. To deny this would require either rejecting the OLG framework itself or imposing a criterion on Bitcoin that Samuelson's analysis does not impose.

This is not yet a verdict on whether Bitcoin functions as money in the world. The OLG result is a possibility theorem, not an existence proof: it establishes that a monetary equilibrium of the relevant form can be obtained, not that any particular candidate asset will find itself in that equilibrium rather than in a non-monetary one. The remainder of the chapter takes up precisely that distinction. Before doing so, it is worth noting what the steelman has achieved: an honest reading of Samuelson grants Bitcoin the formal status its proponents claim for it. The question that remains is whether the formal status is enough.

6.2 Medium of Exchange

The case made in the previous section establishes that Bitcoin satisfies the formal preconditions of Samuelson money. The cryptocurrency is now subjected to the three classical functional tests, beginning with the function most central to money's practical role and most decisively absent from the OLG framework. Samuelson's model assumes frictionless exchange: trade between cohorts is costless, instantaneous, and unmediated by any technology whose properties might constrain settlement. The medium-of-exchange function is the dimension on which the model is silent by construction. It is also the dimension on which Bitcoin's claim to monetary status fails most decisively.

Bitcoin's settlement architecture imposes costs absent from any modern payment infrastructure. New transactions are confirmed roughly every 10 minutes, and standard security practice requires multiple confirmations before a payment is treated as final. Transaction fees, set by congestion-driven auction rather than by a posted schedule, have varied dramatically: an average of approximately \$60 per transaction in April 2021 and \$1.80 in February 2022 (Alvarez, Argente, and Patten, 2023). Withdrawing in local currency through a Bitcoin ATM typically costs between 5 and 20 per cent of the transaction value, with an average of around 8.5 per cent. None of these properties is fatal in principle; none is comparable to the throughput, latency, or fee structure of established payment systems. The protocol's block-size constraint caps system-wide throughput at roughly 7 transactions per second, compared with tens of thousands sustained by major card networks (Croman et al., 2016). The friction is structural, and the empirical question is whether it is binding in practice.

6.2.1 The El Salvador Natural Experiment

The strongest test of that question is provided by El Salvador's Bitcoin Law of September 2021. The reform combined essentially every instrument a monetary authority can deploy to facilitate an asset's adoption. Bitcoin was granted legal tender status, with businesses required to accept it. The state launched the Chivo Wallet, a national digital payments platform supporting both bitcoin and dollar transactions. Each citizen who downloaded the application received a \$30 bitcoin bonus, equivalent to 0.7 per cent of GDP per capita and redeemable only after first being transmitted within the system. Conversions between bitcoin and dollars were subsidised through fee waivers, and petrol purchases routed through Chivo received an additional discount. Alvarez, Argente, and Patten (2023) characterise this as a maximal-intensity adoption push, conducted in a country with an unusually favourable institutional setting: a dollarised economy, a young and digitally connected population, and remittance flows among the highest in the hemisphere.

The outcomes are unambiguous. 68 per cent of the population became aware of the Chivo Wallet, but the volume of substantive engagement collapsed almost immediately. 40% of all lifetime downloads occurred in the launch month, with virtually none thereafter. Among adults who downloaded the application, fewer than half continued to use it after spending the \$30 incentive: approximately 20 per cent of the adult population, and most of this residual usage was denominated in dollars rather than bitcoin. Fewer than ten per cent of adults continued to transact in bitcoin in any form. The median user reported no monthly ATM withdrawals and no monthly bitcoin payments, against an international benchmark of 1.3–1.4 daily transactions per person across payment methods. The remittance channel, which the legal-tender status was explicitly intended to capture, stabilised at 1.6 per cent of total remittance flows by February 2022 and did not subsequently recover (Alvarez, Argente, and Patten, 2023).

The episode is informative because it constitutes the strongest possible push for adoption: legal-tender status, mandatory acceptance, direct subsidies, and a state-sponsored payment platform together exhaust the toolkit available to a government committed to bitcoin adoption. The population-level rejection that followed cannot be attributed to insufficient incentives, regulatory hostility, or coordination failures of the kind a sovereign push is designed to overcome. It must instead be located in the underlying frictions of the payment technology itself: settlement times, fee volatility, custodial mistrust, and the absence of complementarities with existing economic activity.

This is one of the signals that the OLG framework, while necessary for understanding the structural possibility of fiat-style money, is not sufficient to predict which candidate monies will succeed. The framework abstracts from precisely the dimension on which Bitcoin's failure is decisive. Samuelson's silence on the medium-of-exchange function is not a defect of the model; it is the price of its analytical parsimony. But it is also a warning: satisfying the OLG preconditions tells us nothing about whether an asset can, in practice, intermediate the trades that monetary equilibrium presupposes.

6.3 Store of Value

The store-of-value function is the dimension on which the OLG framework speaks most directly. Samuelson's argument hinges precisely on the ability of the fiat token to carry purchasing power across the agent's lifecycle: a unit of currency saved in working life must be redeemable, in real terms, against the consumption goods of the next generation. The biological-rate result is, at its core, a claim about price behaviour: under stationary or balanced-growth conditions, the price level of the monetary asset evolves predictably with population dynamics. If Bitcoin is to be read as Samuelson money, its empirical price behaviour ought to be consistent with this prediction. The data suggest otherwise.

Computing the annualised standard deviation of daily log returns over the five-year window from 1 January 2021 to 31 December 2025 (a horizon that spans a full Bitcoin market cycle, including the 2021 rally, the 2022 contraction, and the ETF-driven recovery of 2024) yields a realised volatility of approximately 58.4 per cent for Bitcoin, against 15.6 per cent for gold (proxied by GLD), 15.9 per cent for global equities (ACWI), and 7.5 per cent for the euro-dollar exchange rate.¹ Bitcoin's volatility thus runs at roughly 3.7 times that of gold and global equities, and at nearly eight times that of a major fiat currency. The amplitude is structural: it has compressed since the early-period figures Yermack reported, but its order of magnitude relative to the relevant monetary benchmarks has not changed. Corbet et al. (2018) document that cryptocurrency volatility remained significantly higher than that of other asset classes throughout the 2013–2017 sample, and Cheah and Fry (2015) present formal evidence that Bitcoin's price dynamics are consistent with speculative bubbles rather than with a stable monetary equilibrium.

The implication for an OLG agent is direct. In Samuelson's model, the young accept the fiat token in exchange for present output, expecting it to retain its value

¹ Author's calculations using daily adjusted close prices from Yahoo Finance. Annualisation factors of $\sqrt{365}$ for Bitcoin and $\sqrt{252}$ for the other series.

and, in the growing-population case, to accrete purchasing power by the time they retire. The reliability of this transfer is what permits the token to clear the intergenerational savings problem. An asset whose price can halve or double over a single period cannot reliably perform this function. The variance of the real return swamps the mean: even if Bitcoin's expected long-run appreciation were positive, the distribution of outcomes for any particular cohort attempting to transfer consumption from working life to retirement would be dominated by realisation risk rather than by the equilibrium return predicted by the framework. The simulation reported in this work recovered Samuelson's biological-rate equilibrium under both rational and damped-adaptive expectations precisely because the price path was anchored by the structural mechanism described by the model. Bitcoin's empirical price path does not exhibit this kind of anchoring.

A second feature of the data complicates the OLG reading from a different direction. Makarov and Schoar (2021) document, using transaction-level blockchain data, that the top 10,000 entities, roughly 0.01 per cent of holders, control approximately 5 million bitcoins, or about a quarter of the circulating supply. The benchmark comparison is again instructive: the top one per cent of US shareholders hold approximately fifty per cent of the equity market, and the top ten per cent of US households hold roughly eighty per cent of total wealth. On a per-holder basis, Bitcoin's ownership concentration substantially exceeds both. Mining capacity displays a parallel pattern: roughly fifty miners, the top 0.1 per cent, control about half of the network's hash rate.

The economic implication is that the assumption of a competitive market (implicit in any application of the OLG framework to Bitcoin) is empirically strained. Makarov and Schoar (2021) further document that whale exits are systematically correlated with positive deviations of the Bitcoin price from its trend: large holders sell into rallies, and the magnitude of these exits is large enough to materially affect the price path. In a Samuelson equilibrium, by contrast, the price level is determined impersonally by the ratio of the fixed money stock to the aggregate real savings demand of the generations alive at each date. No individual agent holds a sufficient share of the asset to move the price in equilibrium. The mechanism the OLG framework describes is a property of the aggregate; Bitcoin's price formation, on the available evidence, is not.

Taken together, the volatility and concentration data support a specific theoretical reading. Samuelson's framework establishes that an intrinsically valueless, fixed-supply asset can sustain a stable monetary equilibrium under the conditions specified by the model; it does not establish that any particular asset will occupy that

equilibrium rather than another. It is, in this sense, a possibility result and not a selection result: the equilibrium it isolates is one element of the bubbly set characterised in Chapter 4, the alternative to which is the bubbleless equilibrium in which the asset commands no positive price.

Bitcoin satisfies the preconditions for occupying some element of this set, but its empirical price dynamics are not those of the anchored equilibrium that the clearing condition selects. Its realised variance and the trading behaviour of its largest holders are inconsistent with the stable, demographically disciplined price path that anchoring implies; they are also inconsistent with the off-anchor paths of Tirole's deterministic continuum, on which the bubble fades in per-capita terms rather than oscillating. The amplitude and narrative-sensitivity that Cheah and Fry (2015), Corbet et al. (2018), and a substantial subsequent literature document correspond, within Tirole's framework, to the stochastic-bubble case set aside in Chapter 4. The OLG framework establishes that the anchored equilibrium is possible; it contains no mechanism that would deliver a given asset into it rather than into the rest of the bubbly set, and Bitcoin's record indicates that satisfying the formal preconditions is consistent with occupying a quite different region of it.

6.4 Unit of Account

On the third function, the failure is categorical rather than partial. Even merchants who nominally accept Bitcoin price their goods in fiat and convert at the point of sale (Yermack, 2015): the listed price is denominated in dollars or euros, the BTC equivalent is computed using a live exchange rate, and settlement is typically followed by immediate conversion back into fiat. The mechanism reveals what the headline acceptance figures conceal, namely, that the quasi-totality of "Bitcoin commerce" consists of fiat-denominated transactions that merely settle on the Bitcoin network. No meaningful volume of wages, debts, rents, or forward contracts is denominated natively in BTC. Where prices are nominally posted in BTC at all, the units required are inconveniently small: a coffee or a paperback is quoted in the fourth or fifth decimal place, a regime that consumer pricing heuristics are not equipped to handle and that even mainstream accounting software does not natively support (Yermack, 2015). The closest historical analogue is the mirror-image case of high-inflation economies, where consumers learn to ignore trailing zeros; ignoring leading zeros has no comparable precedent. Where Bitcoin is used at all, it operates as a transient settlement rail riding atop fiat-denominated commerce: closer in function to a payment instrument than to a numéraire.

The observation has a sharper theoretical edge than it first appears. Samuelson's

economy contains a single homogeneous good, so the unit-of-account function is, by construction, trivially satisfied: there is nothing to denominate against. The OLG framework is consequently silent on whether a candidate's money can serve as a denomination convention in a multi-good economy. It tells us when a fiat token can support an intertemporal equilibrium; it does not tell us when such a token will be adopted as the linguistic unit through which prices, contracts, and accounts are expressed. The silence is structural rather than incidental; e.g. Diamond's OLG extension inherits it intact, since it preserves Samuelson's single-good simplification and modifies only the asset side of the model.

This is therefore not a dimension on which Bitcoin fails the OLG criterion. It is a dimension on which OLG has nothing to say, yet Bitcoin nevertheless fails. Bullmann, Klemm, and Pinna (2019) makes the analogous point from the opposite direction: even stablecoins, designed precisely to cure the volatility that hampers Bitcoin's other monetary functions, must specify a "currency of reference" against which their stability is assessed, thereby conceding the unit-of-account role to the fiat money they nominally substitute. They further note that protecting the unit-of-account function of the underlying currency is a regulatory concern distinct from price stability, confirming, on the policy side, that denomination is treated as a separate institutional fact. It is not derivable from the equilibrium logic that establishes monetary value in the first place; it must be furnished externally, by the convention the OLG model leaves unmodelled.

The chapter set out to test the strongest version of the claim that Bitcoin is Samuelson money, and it granted that claim where it was due. Reduced to its structural features (fixed supply, decentralised issuance, intrinsic worthlessness, and a positive value sustainable only through the forward-looking acceptance of successive cohorts), Bitcoin meets the preconditions the formal account requires, and the steelman was conceded in full rather than as a courtesy. What it could not survive was the imposition of the three classical functions. The asset fails prohibitively as a medium of exchange, unreliably as a store of value, and completely as a unit of account: distinct breakdowns rather than points on a common scale, and qualitative rather than marginal on the first and third.

The verdict is therefore internal to Bitcoin and unambiguous: whatever conceptual room a private, fixed-supply token might in principle occupy, this particular asset does not occupy the monetary equilibrium its proponents describe. The features that earn it the formal preconditions are real, and so is its failure to convert any of them into a working monetary function; the distance between the two shows no sign of closing across the data assembled here. What that distance ultimately implies:

whether it bears against the possibility result established earlier or simply falls outside its scope, belongs to the concluding chapter. It is enough here to record the narrower finding: the asset that satisfies Samuelson's preconditions in the abstract performs none of money's offices in practice, and its standing, on the evidence, is that of a speculative holding rather than a circulating money.

Conclusions

The thesis set out to ask what role an intrinsically valueless asset can play as a store of value in an overlapping-generations economy, noting from the outset that the answer matters in two registers at once. The first is formal: can such an asset be coherently priced at all? The second is empirical: which assets, if any, actually inhabit the equilibrium the theory describes? The chapters that followed settled the first register and sharpened the second. The formal register is no longer Samuelson's result standing alone, it is the single proposition that Chapter 2,3 and 4 assemble in stages and that Section 4.5 gathers in one place: Samuelson, Diamond, and Tirole are not three neighbouring results but three faces of one phenomenon.

That proposition is the formal payoff of the thesis, and it is this: money in Samuelson, rolled-over public debt in Diamond, and the pure bubble in Tirole are three names for a single object: an intrinsically valueless claim that commands a positive price because each generation rationally expects the next to buy it, that can be priced precisely when the economy is dynamically inefficient ($\bar{r} < m$), and that, where it exists, raises welfare by performing the young-to-old transfer that the market alone cannot arrange. They differ only in dress: Samuelson's asset lives in pure exchange and is called money; Diamond's lives in a production economy and is issued by a state; Tirole's lives in the same economy and is issued by no one. What Tirole contributes is generality. The phenomenon is not a peculiarity of money, of capital-less economies, or of fiscal policy; it is a property of overlapping-generations economies as such, switched on and off by a single inequality. Three results that appear to share a formula are, in the end, one result.

The simulation chapter does not extend this result; it illustrates it. A decentralised, log-utility population converges on the biological-rate equilibrium under both rational and damped-adaptive expectations, and almost continues to do so when the nominal money stock is repeatedly doubled. This is the mechanism made visible, not an independent proof of it; the analytics already carry the claim. With the formal register settled, what remains is the question the framework was never built to answer, not whether an intrinsically worthless asset can be money, but which one is. That question is empirical, and Bitcoin is its natural laboratory: a privately issued, protocol-fixed, decentralised token whose proponents claim, in almost these terms, that it occupies the very equilibrium Section 4.5 identifies. It is the empirical test of the unified result, and the Bitcoin chapter takes it up (Chapter 6).

Subjected to that test, Bitcoin yields a verdict in two parts. The first the framework affirms rather than concedes: Bitcoin meets Samuelson's specifications cleanly: a valueless entry on a distributed ledger, carrying no claim on real resources and no issuer's liability; its terminal stock is fixed by protocol at twenty-one million units, a

fixed ceiling no sovereign currency has matched; its issuance is decentralised; and its value, where positive, can rest only on the forward-looking acceptance of successive cohorts. These are the conditions under which Samuelson's result holds, and, on a strict reading, they hold. The framework grants Bitcoin its formal monetary status as theory, not as a reluctant empirical concession. The verdict is non-trivial: it vindicates the structural insight that money needs neither intrinsic value nor sovereign backing to be coherent.

What the formal status cannot survive is the imposition of money's three classical functions. As a medium of exchange, Bitcoin is prohibitive; as a store of value, it is unreliable; as a unit of account, it fails categorically.

Placed within Tirole's framework, Bitcoin's case is one of possibility rather than selection. It satisfies the preconditions to occupy some element of the bubbly set, but its realised variance and narrative sensitivity place it in the stochastic-bubble region set aside in that chapter, rather than the anchored, demographically disciplined equilibrium the clearing condition selects. The preconditions are necessary; they are nowhere near sufficient.

The two verdicts do not conflict; they answer different questions, and Bitcoin is interesting for sitting at their intersection. The OLG framework identifies the necessary condition for fiat-style money: a self-sustaining structure of intergenerational belief, capable of holding positive purchasing power for an asset with no other source of value. That condition is structural, admitting formal demonstration. The three functions identify the conditions sufficient in practice: low frictions, a price path anchored enough that intertemporal transfers do not dissolve into realisation risk, and an external convention of denomination. These are institutional and behavioural, observable but not derivable inside the model. The organising finding is that the dimensions on which Bitcoin breaks down are exactly the dimensions the framework abstracts away: frictionless exchange is assumed, the price path is anchored by the model's own mechanism rather than selected, and the unit of account is made trivial by the single good. This gap does not discredit the framework; it is where the framework falls silent by construction, and that silence is informative. The possibility theorem is true but thin: vindicated, yet carrying no predictive content about which asset is selected into the equilibrium it describes.

The point is not only analytical, and El Salvador is where the world tested it. In September 2021, the state mounted what remains the maximum push any sovereign can make towards an asset's adoption. Awareness became near-universal; substantive use collapsed at once. This rejection cannot be charged to weak incentives,

regulatory hostility, or coordination failure: the sovereign push existed precisely to defeat them. The cause can only lie in the asset's failure to become an object of common belief.

Four years on, the experiment ran in reverse, and it is the reversal that turns the thesis towards its conclusion. In January 2025, as a prior action for a forty-month Extended Fund Facility worth about US\$1.4 billion, the Legislative Assembly amended the Bitcoin Law: acceptance was made voluntary, the obligation on merchants repealed, tax payment in bitcoin ended, with the state's participation in Chivo to be gradually unwound (International Monetary Fund, 2025a). Bitcoin lost its statutory designation as a currency, and with it the footing on which prices could be expressed in it, a footing that had never in practice displaced the dollar, which the original law itself kept as the sole currency of account (International Monetary Fund, 2025b). The sovereign thus withdrew in law the numeraire function the asset had never organically held, leaving what the market had already settled: conversion to and from dollars, not denomination in bitcoin. The experiment had now run in both directions: maximal backing, then institutional retreat, and at neither end did organic acceptance appear. The El Salvador experiment shows us that institutions can compel acceptance and bankroll it, but they cannot manufacture belief.

This returns the thesis to where it began. Chapter 1 described money as a collective fiction with real effects: a social fact in Durkheim's sense, sustained not by substance but by the web of trust and regulation that binds it; relational rather than material in Simmel; and in Weber dependent, however abstract it becomes, on a shared faith in its validity. Samuelson made the recursive-belief half of that description formally tractable. The Bitcoin chapter exposed the other half and showed why no model of comparable parsimony can absorb it without ceasing to do the work the formal model exists to do.

It is worth saying plainly, and in my own voice, what I take the thesis to have shown. Money is coherent without intrinsic value, that much is settled, in Samuelson's terms, and Bitcoin's formal credentials confirm it. But money cannot do without belief. The recursive expectation that others will go on accepting it is not one property among many that a monetary asset happens to possess; it is the property on which all the others depend. Strip a token of intrinsic worth, and it can still be money. Strip it of the shared conviction that it will be accepted tomorrow, and no balance sheet, statute, or protocol will hold its value up. This is the conviction these conclusions exist to record.

Bitcoin is the proof of this in the negative, and El Salvador sharpens it. The cryp-

to currency assembled every structural precondition the theory asks of money, and in El Salvador it acquired, on top of these, the full coercive and fiscal weight of a state. The one thing the apparatus could not supply was the population's belief, and without it the asset never became money. Faith does not matter alongside scarcity, decentralisation, and backing; it is the binding constraint, and the others are not. Backing is neither necessary (Samuelson's token has none, whether reserve or sovereign) nor sufficient, as a state that mustered every instrument of the sovereign kind discovered.

This invites the obvious objection: if all money is belief, why is the dollar stable and Bitcoin not? Because institutions do not substitute for faith; they coordinate it. Legal tender, taxes, the central bank's balance sheet: none is itself the source of value, and none can summon belief by decree. Even the central-bank assets Chapter 1 treated as the dollar's anchor work by making its promise credible enough to coordinate on, not by grounding the value they sustain. What they do is serve as focal points, letting the dispersed beliefs of millions converge on one equilibrium and hold there, among the infinity the fundamentals leave open. Tirole's closing conditions name the requirements exactly: durability, scarcity, and common beliefs. A protocol supplies the first two and not the third. Bitcoin's code can make an asset durable and scarce, but it cannot make the asset an object of common belief.

The anchor is worth making concrete, because the dollar's stability is not magic. Taxes create a recurring, non-discretionary demand for the currency, a floor under its acceptance renewed every fiscal period. Legal tender settles the focal-point problem by convention backed by law, telling everyone which equilibrium to coordinate on. And the credibility of the central bank's balance sheet, together with markets deep enough that no single holder moves the price, keeps price formation impersonal, the real-world face of the model's impersonal clearing. The unit's value is set by the aggregate of savers, not by any single holder: the opposite of Bitcoin's whale-driven path. These are exactly the coordinating devices Bitcoin's design dispenses with. Bitcoin realises Samuelson's 'no central authority' more fully than any fiat currency, and it is for that very reason that it sits outside the anchored equilibrium. The purity is the problem.

The difference between the dollar and Bitcoin is therefore one of kind, not degree: the dollar possesses the coordinating institution that Bitcoin, by construction, lacks. But kind is not permanence, and it would be a misreading to take the dollar's robustness as an escape. Because even the densest institutions coordinate belief rather than replace it, a fully institutionalised fiat currency remains, in the last analysis, a structure of faith, and the historical record is consistent. Currencies have

died of trust crises with their legal scaffolding intact: hyperinflations in which legal-tender laws stayed on the books while the money turned worthless, and the bank runs Chapter 1 reads, after Diamond and Dybvig, as crises of collective confidence rather than failures of the underlying assets. The dollar is robust not because it stands outside the logic that governs Bitcoin, but because its coordination is dense, credible, and continually renewed. Same logic; a different, better-defended position on it.

The argument has driven toward a division between two questions that are easily run together. Whether an intrinsically worthless asset can be money is settled: it can, in Samuelson's terms, and the synthesis of the three models settles it generally rather than as a special case. Whether any particular worthless asset will become money is a different question entirely (institutional, behavioural, conventional) and one the framework was never built to answer. The error in either direction is to mistake one question for the other: to read the possibility result as a prediction, or to read the empirical failure as a refutation of the possibility. Bitcoin's value to the argument is that it holds the two apart, and shows that taking either seriously means taking the other seriously too.

I should be equally plain about the limit, which Tirole stated first, and with more modesty than the debate has since shown. Everything turns on which world describes ours. Where new cohorts arrive without end, perfectly rational agents may pay a positive price for nothing at all; in a world of infinitely-lived traders, or of generations bound to one another by bequests, they cannot, and economics has not settled which of these describes ours. The verdict reached here is therefore conditional on the overlapping-generations account: a claim about what follows if that account is right, not a proof that it is. I take that conditionality to be a limit of the framework, not a defect of the argument, the boundary of what a model of this kind can claim, not a flaw inside the claim.

That boundary also marks where the work points next. The synthesis already flagged the natural extension, Tirole wrote before any of it existed, so the suggestion is the thesis's own: assets engineered to satisfy his conditions, cryptocurrencies durable and scarce by protocol, non-fungible tokens manufacturing uniqueness by design, with the common-belief condition outsourced to code. The open question Bitcoin sharpens is whether engineered scarcity can ever bootstrap that third condition into being, or whether the social coordinator remains irreducible, furnished by convention or not at all. Bitcoin's record points to the latter. But the experiment is young.

There is, in the end, something close to magic in this. Money asks nothing of the

object, no metal, no yield, no sovereign promise, and everything of us: of the shared, continually renewed willingness to believe that it will be worth tomorrow what it is worth today. That willingness is at once the whole of its fragility and the whole of its power.

Bibliography

- Alvarez, Fernando, David Argente, and Diana Van Patten (2023). “Are cryptocurrencies currencies? Bitcoin as legal tender in El Salvador”. In: *Science* 382.6677.
- Balasko, Yves and Karl Shell (1980). “The Overlapping-Generations Model, I: The Case of Pure Exchange without Money”. In: *Journal of Economic Theory* 23, pp. 281–306.
- Bullmann, Dirk, Jonas Klemm, and Andrea Pinna (2019). *In search for stability in crypto-assets: are stablecoins the solution?* Occasional Paper Series 230. European Central Bank.
- Cass, David (1972). “On Capital Overaccumulation in the Aggregative, Neoclassical Model of Economic Growth: A Complete Characterization”. In: *Journal of Economic Theory* 4, pp. 200–223.
- Cheah, Eng-Tuck and John Fry (2015). “Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin”. In: *Economics Letters* 130, pp. 32–36.
- Coleman, James (1990). *Foundations of Social Theory*.
- Corbet, Shaen et al. (2018). “Exploring the dynamic relationships between cryptocurrencies and other financial assets”. In: *Economics Letters* 165.C, pp. 28–34.
- Croman, Kyle et al. (2016). “On Scaling Decentralized Blockchains”. In: *Financial Cryptography and Data Security*. Ed. by Jeremy Clark et al. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 106–125.
- Deflem, Mathieu (2003). “The Sociology of the Sociology of Money: Simmel and the Contemporary Battle of the Classics”. In: *Journal of Classical Sociology* 3.1.
- Diamond, Douglas W. and Philip H. Dybvig (1983). “Bank Runs, Deposit Insurance, and Liquidity”. In: *Journal of Political Economy* 91.3.
- Diamond, Peter A. (1965). “National Debt in a Neoclassical Growth Model”. In: *The American Economic Review* 55.5.
- Gorton, Gary (2017). “The History and Economics of Safe Assets”. In: *Annual Review of Economics* 9.
- Graeber, David (2011). *Debt: The First 5000 Years*.
- Hayek, Friedrich (1976). *Denationalisation of Money*. Institute of Economic Affairs.
- Hayek, Friedrich August von (1939). *Freedom and the economic system*. University of Chicago Press.
- International Monetary Fund (2026). *World Economic Outlook, April 2026*. [Data set]. URL: <https://data.imf.org/en/datasets/IMF.RES:WEO>.
- International Monetary Fund, Western Hemisphere Dept. (2025a). *El Salvador: Request for Extended Arrangement Under the Extended Fund Facility-Press Release; Staff Report; and Statement by the Executive Director for El Salvador*. USA. DOI: 10.5089/9798229003483.002. URL: <https://www.elibrary.imf.org/view/journals/002/2025/058/002.2025.issue-058-en.xml>.

- International Monetary Fund, Western Hemisphere Dept. (2025b). *El Salvador: Selected Issues*. USA. DOI: 10.5089/9798229005104.002. URL: <https://www.elibrary.imf.org/view/journals/002/2025/068/002.2025.issue-068-en.xml>.
- Kiyotaki, Nobuhiro and Randall Wright (1989). “On Money as a Medium of Exchange”. In: *Journal of Political Economy* 97.4.
- Ljungqvist, Lars and Thomas J. Sargent (2018). *Recursive Macroeconomic Theory - Fourth edition*.
- Makarov, Igor and Antoinette Schoar (Oct. 2021). *Blockchain Analysis of the Bitcoin Market*. Tech. rep. 29396. National Bureau of Economic Research.
- McLeay, Michael, Amar Radia, and Ryland Thomas (2014). “Money creation in the modern economy”. In: *Bank of England Quarterly Bulletin* 54.1.
- Samuelson, Paul A. (1958). “An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money”. In: *Journal of Political Economy* 66.6.
- Shell, Karl (1971). “Notes on the Economics of Infinity”. In: *Journal of Political Economy* 79.5.
- Simmel, Georg (1900). *The Philosophy of Money*.
- Solow, Robert M. (1956). “A Contribution to the Theory of Economic Growth”. In: *The Quarterly Journal of Economics* 70.1.
- Tirole, Jean (1982). “On the Possibility of Speculation under Rational Expectations”. In: *Econometrica* 50.5.
- (1985). “Asset Bubbles and Overlapping Generations”. In: *Econometrica* 53.6.
- Wallace, Neil (1980). “The Overlapping-Generations Model of Fiat Money”. In: *Models of Monetary Economies*. Ed. by John H. Kareken and Neil Wallace. Minneapolis: Federal Reserve Bank of Minneapolis, pp. 49–82.
- Yermack, David (2015). “Is Bitcoin a Real Currency? An Economic Appraisal”. In: *Handbook of Digital Currency*. Elsevier, pp. 31–43.

Appendix A

Listing 1: Unified Monte Carlo Simulation

```
1 suppressPackageStartupMessages({
2   library(future)
3   library(future.apply)
4   library(parallelly)
5 })
6 # The four Results-chapter scenarios are obtained by setting
   the two switches below:
7 #   GRATE           : 0   -> stationary population
8 #                   ((1 + 0.005)^25) - 1   -> growing population
   at 0.05% annual
9 #   USE_SHOCKS      : FALSE -> no monetary shocks
10 #                   TRUE  -> recurrent helicopter-drop shocks
11 GRATE <- 0
12 USE_SHOCKS <- FALSE
13 # Shock schedule and set-up
14 SHOCK_EVERY <- 14   # generations between shocks
15 N_SHOCKS    <- 4    # number of shocks
16 SHOCK_MULT  <- 2    # money-supply multiplier per shock
17 #Number of generations
18 N_GENS <- 70
19 #Number of Monte Carlo repetitions
20 N_REPS <- as.integer(10000)
21 #Number of agents
22 DEFAULT_N_AGENTS <- if (GRATE == 0) 100000L else 500L
23 N_AGENTS <- as.integer(as.character(DEFAULT_N_AGENTS))
24 #Number of cores for parallelism
25 N_WORKERS <- as.integer(availableCores()-4)
26 #Population and simulation settings
27 SHAPE <- list(a = 7, b = 3)
28 ADAPT_WEIGHT <- 0.1
29 BASE_SEED <- 1958
30 #Output setup
31 library(this.path)
32 current_dir <- this.path::this.dir()
33 source(file.path(current_dir, "coreSim.R"))
34 REGIME_LABEL <- if (GRATE == 0) "Stationary" else "Growing"
35 SHOCK_LABEL  <- if (USE_SHOCKS) "MonShocks" else "NoShocks"
36 OUT_DIR <- file.path(
37   current_dir,
38   sprintf("%smonteCarlo%s_results", REGIME_LABEL, SHOCK_LABEL
```

```

    )
39 )
40 dir.create(OUT_DIR, showWarnings = FALSE, recursive = TRUE)
41 #Helper function for one Monte Carlo repetition
42 oneRep <- function(rep_id,
43                     mode,
44                     shape,
45                     nAgents,
46                     gRate,
47                     nGens,
48                     adaptWeight,
49                     baseSeed,
50                     useShocks,
51                     shockEvery,
52                     nShocks,
53                     shockMult) {
54   set.seed(baseSeed + rep_id)
55   g <- tryCatch(
56     run1Economy(
57       shape,
58       nAgents,
59       1,
60       gRate,
61       nGens,
62       expectMode = mode,
63       adaptWeight = adaptWeight,
64       useShocks = useShocks,
65       shockEvery = shockEvery,
66       nShocks = nShocks,
67       shockMult = shockMult
68     ),
69     error = function(e) {
70       message(sprintf("Rep %d (%s) failed: %s",
71                       rep_id, mode, conditionMessage(e)))
72       NULL
73     }
74   )
75   if (is.null(g)) {
76     N <- nGens + 1
77     nan_vec <- rep(NA_real_, N)
78     return(list(

```

```

79     failed                = TRUE,
80     rep_id                = rep_id,
81     P                     = nan_vec,
82     r_annual              = nan_vec,
83     aggRealSav            = nan_vec,
84     moneyYoung            = nan_vec,
85     moneyMid              = nan_vec,
86     meanBeta              = nan_vec,
87     sdBeta                = nan_vec,
88     nYoung                = nan_vec,
89     nMid                  = nan_vec,
90     nOld                  = nan_vec,
91     R_exp                 = nan_vec,
92     M                     = nan_vec,
93     borrowerShareFinal    = NA_real_,
94     meanBetaFinal         = NA_real_,
95     sdBetaFinal           = NA_real_
96   ))
97 }
98 youngFinal <- g$finalPop$age == 0L
99 list(
100   failed                = FALSE,
101   rep_id                = rep_id,
102   P                     = g$results$P,
103   r_annual              = g$results$r_annual,
104   aggRealSav            = g$results$aggRealSav,
105   moneyYoung            = g$results$moneyYoung,
106   moneyMid              = g$results$moneyMid,
107   meanBeta              = g$results$meanBeta,
108   sdBeta                = g$results$sdBeta,
109   nYoung                = g$results$nYoung,
110   nMid                  = g$results$nMid,
111   nOld                  = g$results$nOld,
112   R_exp                 = g$results$R_exp,
113   M                     = g$results$M,
114   borrowerShareFinal    = mean(g$finalPop$money[youngFinal] <
115     0),
116   meanBetaFinal         = mean(g$finalPop$discFact),
117   sdBetaFinal           = sd(g$finalPop$discFact)
118 )

```

```

119 #Monte Carlo Simulation
120 runMC <- function(mode, label) {
121   cat(sprintf("[%s] Starting %s mode: %d reps on %d workers
122     ...\n", format(Sys.time(), "%H:%M:%S"), label, N_REPS, N
123     _WORKERS))
124   t0 <- Sys.time()
125   results <- future_lapply(
126     seq_len(N_REPS),
127     oneRep,
128     gRate = GRATE,
129     mode = mode,
130     shape = SHAPE,
131     nAgents = N_AGENTS,
132     nGens = N_GENS,
133     adaptWeight = ADAPT_WEIGHT,
134     baseSeed = BASE_SEED,
135     useShocks = USE_SHOCKS,
136     shockEvery = SHOCK_EVERY,
137     nShocks = N_SHOCKS,
138     shockMult = SHOCK_MULT,
139     future.seed = NULL,
140     future.scheduling = 2.0
141   )
142   elapsed <- difftime(Sys.time(), t0, units = "mins")
143   cat(sprintf("[%s] %s done in %.2f min\n", format(Sys.time()
144     , "%H:%M:%S"), label, as.numeric(elapsed)))
145   failed <- vapply(results, function(x) isTRUE(x$failed),
146     logical(1))
147   nFailed <- sum(failed)
148   if (nFailed > 0) {
149     failed_ids <- vapply(results[failed], '[', integer(1), "
150       rep_id")
151     cat(sprintf("[%s] WARNING: %d / %d reps failed in %s mode
152       . IDs: %s\n",
153         format(Sys.time(), "%H:%M:%S"),
154         nFailed, N_REPS, label,
155         paste(failed_ids, collapse = ", ")))
156   } else {
157     failed_ids <- integer(0)
158   }
159   N <- N_GENS + 1

```

```

154 packMat <- function(field) {
155   m <- vapply(results, function(x) x[[field]], numeric(N))
156   storage.mode(m) <- "double"
157   m
158 }
159 packScalar <- function(field) {
160   vapply(results, function(x) x[[field]], numeric(1))
161 }
162 out <- list(
163   mode = mode,
164   nReps = N_REPS,
165   nGenerations = N_GENS,
166   nAgents = N_AGENTS,
167   shape = SHAPE,
168   adaptWeight = ADAPT_WEIGHT,
169   gRate = GRATE,
170   useShocks = USE_SHOCKS,
171   shockEvery = SHOCK_EVERY,
172   nShocks = N_SHOCKS,
173   shockMult = SHOCK_MULT,
174   nRepsFailed = nFailed,
175   failedRepIds = failed_ids,
176   nRepsOk = N_REPS - nFailed,
177   # Period x rep matrices
178   P = packMat("P"),
179   r_annual = packMat("r_annual"),
180   aggRealSav = packMat("aggRealSav"),
181   meanBeta = packMat("meanBeta"),
182   sdBeta = packMat("sdBeta"),
183   nYoung = packMat("nYoung"),
184   nMid = packMat("nMid"),
185   nOld = packMat("nOld"),
186   R_exp = packMat("R_exp"),
187   moneyYoung = packMat("moneyYoung"),
188   moneyMid = packMat("moneyMid"),
189   M = packMat("M"),
190   borrowerShareFinal = packScalar("borrowerShareFinal"),
191   meanBetaFinal = packScalar("meanBetaFinal"),
192   sdBetaFinal = packScalar("sdBetaFinal"),
193   elapsed_min = as.numeric(elapsed)
194 )

```

```

195   outFile <- file.path(OUT_DIR, sprintf("mc_%s.rds", label))
196   saveRDS(out, outFile, compress = "xz")
197   cat(sprintf(
198     "[%s] Saved %s (%.1f MB)\n",
199     format(Sys.time(), "%H:%M:%S"),
200     outFile,
201     file.info(outFile)$size / 1e6
202   ))
203   invisible(out)
204 }
205 # Allow for parallelism both on Windows and Unix
206 if (.Platform$OS.type == "unix") {
207   plan(multicore, workers = N_WORKERS)
208 } else {
209   plan(multisession, workers = N_WORKERS)
210 }
211 cat(sprintf(
212   "[%s] Backend: %s with %d workers\n",
213   format(Sys.time(), "%H:%M:%S"),
214   class(plan())[1], N_WORKERS
215 ))
216 cat(sprintf(
217   "[%s] Scenario: %s population, shocks %s\n",
218   format(Sys.time(), "%H:%M:%S"),
219   tolower(REGIME_LABEL), if (USE_SHOCKS) "ON" else "OFF"
220 ))
221 cat(sprintf(
222   "[%s] Detected logical cores: %d\n",
223   format(Sys.time(), "%H:%M:%S"),
224   availableCores()
225 ))
226 # Run both modes
227 mc_rat <- runMC("rational", "rational")
228 mc_ada <- runMC("adaptive", "adaptive")
229 # Summary
230 cat("\n===== SUMMARY =====\n")
231 for (mc in list(mc_rat, mc_ada)) {
232   burnIn <- 2
233   idx <- burnIn:(N_GENS + 1)
234   meanR <- mean(rowMeans(mc$r_annual[idx, , drop = FALSE], na
     .rm = TRUE))

```

```

235   cat(sprintf(
236     "%-9s mode | mean r_annual (post burn-in) = %+.3e |
        runtime = %.2f min\n",
237     mc$mode, meanR, mc$elapsed_min
238   ))
239 }
240 plan(sequential)
241 cat(sprintf(
242   "[%s] All done. Files in %s/\n",
243   format(Sys.time(), "%H:%M:%S"), OUT_DIR
244 ))

```

Listing 2: Unified Core OLG Simulation

```

1  # Core OLG consumption-loan simulation (Samuelson, 1958)
2  # Population set-up
3  pop <- list(
4    age      = integer(0),
5    discFact = numeric(0),
6    money    = numeric(0)
7  )
8  # Initialise population at t=0
9  stGen <- function(shape, nAgents, gRate) {
10    wts <- c(1, 1 / (1 + gRate), 1 / (1 + gRate)^2)
11    wts <- wts / sum(wts)
12    counts <- round(nAgents * wts)
13    counts[1] <- nAgents - sum(counts[2:3])
14    list(
15      age      = sample(rep(0:2, counts)),
16      discFact = rbeta(nAgents, shape$a, shape$b),
17      money    = numeric(nAgents)
18    )
19  }
20  # OLG dynamics
21  nGen <- function(shape, pop, gRate, targetFloat) {
22    # Kill the old
23    keep <- pop$age != 2L
24    # Age the survivors
25    age2 <- pop$age[keep] + 1L
26    # Carry on discount factor and money balances
27    disc2 <- pop$discFact[keep]
28    mny2 <- pop$money[keep]
29    # Grow the population, accounting for rounding errors

```

```

30 targetSize <- round(targetFloat)
31 numNewBorns <- max(1L, targetSize - length(age2))
32 # Mate only middle aged individuals, and avoid self
   fertilization
33 poolIdx <- which(age2 == 1L)
34 nP <- length(poolIdx)
35 i1 <- poolIdx[sample.int(nP, numNewBorns, replace = TRUE)]
36 shift <- sample.int(nP - 1, numNewBorns, replace = TRUE)
37 i2 <- poolIdx[((match(i1, poolIdx) - 1L + shift) %% nP) + 1
   L]
38 # Create new individuls
39 randComp <- rbeta(numNewBorns, shape$a, shape$b)
40 newDisc <- disc2[i1] * 1 / 4 + disc2[i2] * 1 / 4 + randComp
   * 1 / 2
41 # Redefine population
42 list(
43   age      = c(age2, integer(numNewBorns)),
44   discFact = c(disc2, newDisc),
45   money     = c(mny2, numeric(numNewBorns))
46 )
47 }
48 # Market clearing dynamic
49 clearMarket <- function(pop, w, M, R_exp) {
50   young <- pop$age == 0L
51   mid <- pop$age == 1L
52   old <- pop$age == 2L
53   b <- pop$discFact
54   b2 <- b * b
55   lifetime_wealth_y <- w + w / R_exp
56   c0 <- lifetime_wealth_y / (1 + b + b2)
57   s_young <- w - c0
58   A <- sum(s_young[young])
59   B <- sum((b[mid] / (1 + b[mid])) * w)
60   C <- sum((b[mid] / (1 + b[mid])) * pop$money[mid])
61   # Clear the savings market
62   P_t <- (M - C) / (A + B)
63   pop$money[old] <- 0
64   pop$money[young] <- s_young[young] * P_t
65   pop$money[mid] <- (b[mid] / (1 + b[mid])) * (pop$money[mid]
   + w * P_t)
66   stopifnot(all(pop$money[pop$age == 1L] >= -1e-8))

```

```

67   stopifnot(abs(sum(pop$money) - M) < 1e-6)
68   list(pop = pop, P = P_t)
69 }
70 # Run one economy from initialization to final generation
71 run1Economy <- function(shape,
72                          nAgents,
73                          wAnnual,
74                          gRate,
75                          nGenerations,
76                          M = NULL,
77                          expecMode = "rational",
78                          adaptWeight = 0.1,
79                          useShocks = FALSE,
80                          shockEvery = 14,
81                          nShocks = 4,
82                          shockMult = 2) {
83   # Helicopter money shock schedule (applied only when
84     useShocks = TRUE)
85   shockPeriods <- shockEvery * seq_len(nShocks)
86   # Turn annual wage into period wage
87   w <- wAnnual * 25
88   # Initialize a coherent money supply, can also be exogenous
89   if (is.null(M)) M <- nAgents * w * 0.5
90   # Create population at t=0
91   pop <- stGen(shape, nAgents, gRate)
92   # Endow all fiat money to old generation
93   oldMask <- pop$age == 2L
94   pop$money[oldMask] <- M / sum(oldMask)
95   targetFloat <- as.numeric(nAgents)
96   # Shift periods by 1, so that t=0 equals to starting
97     generation
98   N <- nGenerations + 1
99   # Initilise results table
100   results <- data.frame(
101     period      = 0:nGenerations,
102     nAgents     = integer(N),
103     nYoung     = integer(N),
104     nMid       = integer(N),
105     nOld       = integer(N),
106     meanBeta   = numeric(N),
107     sdBeta     = numeric(N),

```

```

106     P                = numeric(N),
107     logP             = numeric(N),
108     R_real           = numeric(N),
109     r_annual         = numeric(N),
110     aggRealSav       = numeric(N),
111     moneyYoung       = numeric(N),
112     moneyMid         = numeric(N),
113     R_exp            = numeric(N),
114     M                = numeric(N)
115 )
116 R_ss <- 1 + gRate
117 R_exp <- R_ss
118 # Main loop
119 for (i in 1:N) {
120     # Set expectations before clearing the market
121     if (expectMode == "rational") {
122         R_exp <- R_ss
123     }
124     out <- clearMarket(pop, w, M, R_exp)
125     pop <- out$pop
126     results$nAgents[i] <- length(pop$age)
127     results$nYoung[i] <- sum(pop$age == 0L)
128     results$nMid[i] <- sum(pop$age == 1L)
129     results$nOld[i] <- sum(pop$age == 2L)
130     results$meanBeta[i] <- mean(pop$discFact)
131     results$sdBeta[i] <- sd(pop$discFact)
132     results$P[i] <- out$P
133     results$logP[i] <- log(out$P)
134     results$aggRealSav[i] <- M / out$P
135     results$moneyYoung[i] <- sum(pop$money[pop$age == 0L])
136     results$moneyMid[i] <- sum(pop$money[pop$age == 1L])
137     results$R_exp[i] <- R_exp
138     results$M[i] <- M
139     # Helicopter-drop shock
140     if (useShocks && (i - 1) %in% shockPeriods) {
141         recipients <- pop$age != 2L
142         dropAmount <- M * (shockMult - 1)
143         perAgent <- dropAmount / sum(recipients)
144         pop$money[recipients] <- pop$money[recipients] +
            perAgent
145         M <- M * shockMult

```

```

146   }
147   if (i >= 2) {
148     results$R_real[i] <- results$P[i - 1] / results$P[i]
149     results$r_annual[i] <- results$R_real[i]^(1 / 25) - 1
150     if (expectMode == "adaptive") {
151       R_exp <- (1 - adaptWeight) * R_exp + adaptWeight *
         results$R_real[i]
152     }
153   } else {
154     results$R_real[i] <- NA
155     results$r_annual[i] <- NA
156   }
157   if (i < N) {
158     targetFloat <- targetFloat * (1 + gRate)
159     pop <- nGen(shape, pop, gRate, targetFloat)
160   }
161 }
162 list(results = results, finalPop = pop, params = list(
163   shape = shape, nAgents = nAgents, wAnnual = wAnnual,
164   gRate = gRate, nGenerations = nGenerations, M = M,
165   expectMode = expectMode, adaptWeight = adaptWeight,
166   useShocks = useShocks, shockEvery = shockEvery,
167   nShocks = nShocks, shockMult = shockMult
168 ))
169 }

```

Dichiarazione del ricorso a strumenti di Intelligenza Artificiale

Il sottoscritto *Mattia Cogo*, matricola numero *902265*, iscritto presso Ca' Foscari al corso di Laurea in *Economics Markets, and Finance*, autore dell'elaborato finale *When nothing is worth something - Intrinsically worthless assets in OLG frameworks*, relatore *Ch. Prof. Pietro Enrico Dino Dindo*, consapevole delle sanzioni previste dall'art. 76 del Testo unico, D.P.R. 28/12/2000 n.445, e della decadenza dei benefici prevista dall'art. 75 del medesimo Testo unico in caso di dichiarazioni false o mendaci, sotto la propria personale responsabilità ai sensi degli artt. 46 e 47 del D.P.R. 28/12/2000 n. 445. Dichiaro di aver fatto ricorso a tecnologie di intelligenza artificiale generativa per supportare vari aspetti della ricerca e della scrittura del suo elaborato finale, in particolare, di aver utilizzato gli strumenti: *Gemini 3.1 Pro* - *Giugno 2026*; *Claude Opus 4.8* - *Giugno 2026*; *Grammarly* - *Giugno 2026*, al fine di: 1. revisione e ottimizzazione del testo per migliorarne la chiarezza e coesione in lingua inglese; 2. correzione di bozze e controllo ortografico; 3. supporto al formattazione in LaTeX; 4. assistenza alla scrittura di codice in R; sezioni e capitoli interessati: l'interezza del testo per i fini numero 1,2,3; il capitolo 5 per il fine numero 4. Il sottoscritto ha condotto in modo indipendente tutta la ricerca e la revisione della letteratura, consultando direttamente fonti primarie e secondarie. Il contenuto analitico, le argomentazioni e le conclusioni della tesi sono interamente opera dell'autore, gli strumenti di intelligenza artificiale sono stati utilizzati esclusivamente a supporto. Tutte le revisioni effettuate con l'ausilio dell'IA sono state esaminate criticamente e approvate dall'autore.

Dichiaro inoltre che tutto il materiale generato dall'IA è stato attentamente verificato, rivisto e rielaborato dall'autore, che si assume la piena responsabilità per la qualità, l'accuratezza e l'integrità scientifica del lavoro.



Grazie per l'attenzione.
Thank you for your attention.

