



JON M.
HUNTSMAN
SCHOOL OF BUSINESS

Department of Economics and Finance Working Paper Series

Optimal Irrational Behavior

James Feigenbaum
Frank N. Caliendo
Emin Gahramanov
DEFWP2009-01

Department of Economics and Finance
Huntsman School of Business
Utah State University
3565 Old Main Hill
Logan, UT 84322-3565

Optimal Irrational Behavior

James Feigenbaum^{*†} Frank N. Caliendo
University of Pittsburgh Utah State University

Emin Gahramanov
Deakin University

November 27, 2008

Abstract

Contrary to the usual presumption that welfare is maximized if consumers behave rationally, we show in a two-period overlapping generations model that there always exists a rule of thumb that can weakly improve upon the lifecycle/permanent-income rule in general equilibrium with irrational households. The market-clearing mechanism introduces a pecuniary externality that individual rational households do not consider when making decisions, but a publically shared rule of thumb can exploit this effect. For typical calibrations, the improvement of the welfare of irrational households is robust to the introduction of rational agents. Generalizing to a more realistic lifecycle model, we find in particular that the Save More TomorrowTM (SMarT) Plan can confer higher lifetime utility than the permanent-income rule in general equilibrium.

JEL Classification: C61, D11, E21

Keywords: consumption, saving, coordination, lifecycle/permanent-income hypothesis, SMarT Plan, general equilibrium, rules of thumb, pecuniary externality

^{*}Corresponding author: jfeigen@pitt.edu. <http://www.pitt.edu/~jfeigen/>

[†]The authors would like to thank Shlomo Benartzi, Basudeb Biswas, Dave DeJong, Scott Findley, Dwight Israelsen, Geng Li, Hyeon Park, Marla Ripoll, and Pedro Silos for their suggestions and comments. We also appreciate feedback we received from seminar participants at SUNY Buffalo and the University of New Mexico.

1 Introduction

The dominant paradigm in economics for at least the last three decades has been to assume that people behave rationally, meaning they individually maximize their utility. Nevertheless, if we test whether people follow the basic lifecycle/permanent-income (LCPI) consumption rule that comes out of the benchmark preference specification, then empirically they often fall short.¹ According to the 1996 Health and Retirement Study, 33% of retirees have no private savings, a fact that cannot be reconciled with the notion that households smooth their consumption, given that Social Security only provides about a third of income prior to retirement (Huang and Caliendo (2007)). Indeed, the preponderance of empirical evidence suggests that a large fraction of households lack either the ability or the inclination to figure out their optimal consumption. According to the 1997 Retirement Confidence Survey, only 36% of respondents said they had even attempted to calculate how much saving they need to maintain their desired lifestyle after retirement (Yakoboski and Dickemper (1997)). Likewise, Lusardi and Tufano (2008) found that a majority of respondents to a market-research survey could not correctly answer basic questions about debt and interest. If households are not maximizing utility, then presumably they are following simple rules of thumb. For example, Benartzi and Thaler (2001) have found that a noticeable fraction of households naively diversify by allocating funds evenly across all available investment options without regard for the statistical dependence of their returns. The question addressed by this paper is, if irrational households follow rules of thumb, then what is the optimal rule of thumb for them to adopt?

Remarkably, we find that, in an overlapping generations (OLG) environment, the optimal rule of thumb actually confers higher lifetime utility than the LCPI rule. This could not happen in partial equilibrium, where, by definition, rational behavior maximizes lifetime utility given exogenous prices. But in general equilibrium, markets exhibit a pecuniary externality (McKean (1958), Prest and Turvey (1965)) in that household behavior determines prices. For example, if everyone saves more, the capital stock will rise. This should increase the marginal product of labor and wages, which increases each household's lifetime wealth.² In competitive markets where no one has any pricing power, ratio-

¹If we allow for frictions such as borrowing constraints or uninsurable idiosyncratic income risk, rational consumers ought to save even more than is dictated by the LCPI rule (Carroll and Kimball (2001), Feigenbaum (2008), Skinner (1988)).

²There is also a change in interest rates, but this has an ambiguous effect on lifetime utility.

nal households maximize their own utility assuming their actions do not affect prices, even though collectively their actions do affect prices. Thus individually rational households cannot exploit the pecuniary externality. But a publically shared rule of thumb can act as a focal point to coordinate behavior so the gains from the pecuniary externality are harnessed, much as a rule to always do what the woman wants (or what the man wants) can avoid coordination failures in the Battle of the Sexes game.³ Although there is the possibility of a coordination failure, we emphasize that the only friction in our environment is the intrinsic friction of an OLG model that there are infinitely many agents who can only trade with contemporaneous agents.⁴ The pecuniary externality is not really an externality per se, though its effects are mathematically similar to a technological externality.

At first blush, our result may appear to contradict the Welfare Theorems.⁵ The reason why there is no inconsistency is because the optimal rule of thumb does not generate a Pareto optimal allocation. While we can frame the problem of determining the optimal rule of thumb as a social planner's problem, it is not the same problem that a Pareto social planner would solve. In the standard Pareto problem, a social planner is given an initial capital stock and needs to solve for the optimal allocation through the rest of time, maximizing a social welfare function that is a weighted sum of the utilities of all generations who are alive to begin with or will be born thereafter. Our social planner solves for the consumption rule that maximizes lifetime utility in the steady state. The pecuniary externality does not play a role in the Pareto problem because at every decision point the existing capital stock is either fixed by initial conditions or determined at a previous decision point. In contrast, our problem focuses on the steady state, so the existing capital stock is a choice variable.

In the context of an OLG model, choosing a rule of thumb to maximize lifetime utility in the steady state is the natural analog of what we would do to compute a rational steady-state equilibrium, which is also to maximize lifetime utility, albeit with steady-state prices determined outside this optimization problem. Economists have rightly been wary of the notion that we can make

³Heifetz, Shannon, and Spiegel (2007) have similarly shown that strategic externalities in a game-theoretic environment can lead to improvements in expected utility if agents do not maximize their own expected utility.

⁴More precisely, Geanakoplos (2008) argues that OLG models differ from their finite-horizon Arrow-Debreu counterparts because of a lack of market clearing at infinity.

⁵The Welfare Theorems are not applicable to OLG models where dynamic inefficiencies arise, but this requires extreme parameters. For plausibly calibrated OLG models, the Welfare Theorems do apply.

welfare comparisons of different policies based only on what happens in the steady state. Within the scope of the present paper, this issue is not relevant to the positive question of how the economy would fare in the long run under different rules of thumb.⁶

Nevertheless, our finding that the optimal rule of thumb outperforms the LCPI rule creates a strong impetus to reconsider the merits of employing steady-state lifetime utility as a social welfare function. Indeed, there is much controversy about whether Pareto social welfare functions are appropriate for analyzing policy in OLG models. With an infinite number of generations, a Pareto social planner must discount future generations in order to obtain a finite social welfare function, a practice that Ramsey (1928) decried as “ethically indefensible and arises merely from the weakness of the imagination” in some of the very first work on the subject. Using steady-state lifetime utility as the social welfare function has the advantage that the social planner weights all generations the same and uses the household’s own welfare function as the basis for the social welfare function, just as is done for an infinitely-lived representative agent.⁷ In effect, we are setting the generational discount rate to zero, which many economists would now advocate.⁸

The conventional criticism of this approach is that it completely disregards the costs of transitioning to the steady state. A social planner with dictatorial powers who only cares about long-run utility will confiscate resources from early generations to build up the golden-rule capital stock that maximizes steady-state consumption, after which he will redistribute the fruits of this investment to later generations. Many economists, including ourselves, would argue that imposing tremendous sacrifices on early generations to achieve a utopian future is unacceptable (Arrow (1999)). However, in our setup, the “social planner” is not a dictator but just a mathematical device. No one has any power to redistribute wealth. All transfers of goods have to occur through trade in markets. If the entire population adopts the optimal rule of thumb, there will be an initial utility loss due to an unfavorable change in factor prices, but the

⁶We are following a similar methodology as Imrohoroglu et al (2003), who considered the lifetime utility in the steady state under different Social Security arrangements, and other papers along the same lines.

⁷Unless we explicitly introduce the utility of the next generation (and this utility is discounted relative to the present generation’s) into a household’s utility function, a finitely-lived rational household will not maximize an objective function that looks like a Pareto social welfare function.

⁸See Portney and Weyant (1999) for a collection of papers arguing the pros and cons of zero discounting.

market puts a check on the size of this loss. Moreover, a utility loss would also occur if an initially irrational population that saves too little begins to behave rationally and adopts the LCPI rule.

For explicative purposes, we begin with a two-period OLG model *à la* Diamond (1965). In this simple example, there is only one decision for the household to make: how much of his initial income should he consume or save. Thus the LCPI rule is nested within a one-parameter family of possible consumption rules. We show that the LCPI rule coincides with the optimal consumption rule only for knife-edge parameterizations where the decentralized market equilibrium also coincides with a golden-rule steady state. If the LCPI rule confers a dynamically efficient market equilibrium then the optimal rule is to save more (though generally not as much as in the golden-rule equilibrium). Conversely, if the LCPI rule confers a dynamically inefficient market equilibrium, then the optimal rule is, naturally, to save less. In either case, the optimal rule of thumb weakly improves upon the LCPI rule.

Since no one has any power to force agents to follow the rule of thumb, one might be concerned that this improvement will disappear if there are some rational agents who can work out the LCPI consumption rule that maximizes their own individual utility. In the two-period model, if everyone follows the optimal rule of thumb, they will all earn higher utility than they would if they all follow the LCPI rule. However, if one agent deviates by acting rationally and following the LCPI rule, he can gain even higher utility than those who adhere to the optimal rule of thumb. In essence, a rational agent could act as a free rider, taking advantage of the high wages engendered by the high saving of his more socially conscious brethren while enjoying a smoother consumption path.⁹ We consider the experiment where a fraction of agents are rational and the remainder follow the optimal rule of thumb (derived assuming that everyone follows the rule). For plausible calibrations of the model, a huge influx of rational agents is needed to create a situation where the irrational followers do worse than they would if everyone followed the LCPI rule. For the baseline model, the necessary fraction is two thirds. Thus even if some agents do behave rationally, this need not destroy the result that everyone can be made better off than they would be in the rational market equilibrium.

⁹Interestingly though, the irrational agents end up accumulating more financial wealth. While the rational agents exploit the irrational agents in terms of utility, they do not crowd the irrational agents out of the market as often happens in models where rational agents interact with irrational agents. See, for example, Blume and Easley (2000).

Next, as a concrete application of our idea, we consider a family of rules that have actually received popular interest for practical implementation. Save More TomorrowTM is a plan proposed by Thaler and Benartzi (2004) in which employees commit today to save some fraction of future wage increases. The plan is also referred to as escalated saving. Several leading investment institutions in the U.S. have given a pool of workers numbering into the millions the option of participating in such a plan, and the response has been quite positive.¹⁰

Extending the two-period model to a continuous-time OLG model, we find under standard calibrations that the SMarT economy dominates the permanent-income economy for a wide range of SMarT Plan saving rates. For example, if we consider a general-equilibrium economy consisting of consumers who follow the very simple rule of thumb of just saving 3.5% of wage income (just as in the pilot implementations of the SMarT Plan at a U.S. manufacturing firm), then they would do worse than in the corresponding general equilibrium where they follow the LCPI rule. However, if the rule-of-thumb consumers participate in a SMarT Plan, committing to save between 55% to 85% of future wage increases, then lifetime utility will exceed the level reached in the permanent-income economy.¹¹

A second implication of our general-equilibrium result relates to the common concern that the SMarT Plan is so powerful that it may actually induce too much saving. This is a valid concern in partial equilibrium. Because the SMarT Plan deviates from a smooth consumption path, if households save too much they can actually do worse than if they do not save at all (Findley and Caliendo (2007)). But in general equilibrium, the pecuniary externality becomes relevant. If a large fraction of the population saves more, this will increase wages, so consumers may be better off even if the SMarT Plan produces a consumption profile that is far from a smooth permanent-income path. For typical calibrations of the model, this wealth effect usually dominates. If we suppose that households have a base saving rate of 3.5% then any SMarT Plan that increases saving out of future wage increases will increase lifetime utility relative to the 3.5% rule of thumb.¹²

¹⁰In the 2006 Retirement Confidence Survey, 65% of workers said they would like to participate in a SMarT Plan (Helman et al (2006)). As evidence of the widespread political support, the recent Pension Protection Act of 2006 creates incentives for firms with 401(k) plans to make the default setting a SMarT Plan, meaning workers would have to actively opt out to not participate (Moore (2006)).

¹¹Such saving rates may seem preposterously high, but in actual practice people did commit to such rates in the pilot program (Thaler and Benartzi (2004)).

¹²We have not done this experiment for SMarT Plan saving rates greater than 100%, although in practice some individuals may opt for this.

Thus while the presumption in macroeconomics has been that consumers allocate resources most efficiently across the lifecycle if they follow the LCPI rule, in general this is only true if the entire population is rational. If some agents are irrational, they can be given consumption rules that, if followed by the irrational agents, will make everyone better off, albeit with even larger gains for those agents who are rational. Broadly speaking, our findings should perhaps not be too surprising since microeconomists have identified many models, such as the Prisoner's Dilemma, with rational equilibria that are Pareto dominated by equilibria involving irrational behavior. Macroeconomists have typically assumed that such considerations do not matter in macro models absent frictions, but in an overlapping generations context there are benefits to coordinating behavior across generations.

The paper is summarized as follows. In Section 2, we show in the two-period model that there is always a rule of thumb that can weakly improve upon the LCPI rule in general equilibrium, and we consider some experiments where the economy transitions between different consumption rules. In Section 3, we present quantitative results for the implementation of the SMarT Plan in a more realistic lifecycle model. We conclude in Section 4 with some discussion of the larger implications of these results.

2 A Simple Two-Period Example

2.1 The Basic Model

The pecuniary externality that allows us to improve upon the LCPI rule can best be understood in terms of a two-period OLG model. In each period, a continuum of agents of unit measure is born and lives with certainty for two periods. Consumers value allocations over the lifecycle according to the preferences

$$U(c_0, c_1) = u(c_0) + \beta u(c_1), \quad (1)$$

where the discount factor $\beta > 0$, $u(\cdot)$ is a strictly increasing, strictly concave, twice-differentiable utility function, and c_t is consumption at age t where $t = 0, 1$.

A consumer at age t has an endowment of e_t efficiency units of labor that he supplies inelastically to the market, for which he is compensated at the real wage w . The consumer can invest in capital K at age 0, for which he is compensated

at the (gross) rate of return R . Thus the consumer faces the budget constraints

$$c_0 + K = we_0 \quad (2)$$

$$c_1 = we_1 + RK. \quad (3)$$

A rational consumer will choose c_0 , c_1 , and K to maximize (1) subject to (2) and (3). This can be reduced to the problem of choosing the saving K to maximize

$$L_d(K|R, w) = u(we_0 - K) + \beta u(we_1 + RK), \quad (4)$$

where d indicates this is the allocation problem for a decentralized economy. However, we do not assume that all the agents in this economy are rational. While all consumers have preferences described by (1), some or all of them may not be capable of solving the problem of maximizing $L_d(K)$.

On the production side of the economy, we assume there is a competitive firm with the production function

$$Y = F(K, N), \quad (5)$$

where K is the capital stock and N is the supply of labor as measured in efficiency units. In equilibrium, the aggregate capital stock K must equal the saving of young agents and the labor supply must satisfy

$$N = e_0 + e_1. \quad (6)$$

We assume that F exhibits constant returns to scale, so the production function can be rewritten

$$F(K, N) = f\left(\frac{K}{N}\right)N, \quad (7)$$

where f is a strictly increasing, strictly concave, twice-differentiable function. Factor prices are determined by marginal principles under perfect competition. Capital depreciates at the rate $\delta \in [0, 1]$, so the gross rate of return on capital (net of depreciation) is

$$R(K) = F_K(K, N) + 1 - \delta = f'\left(\frac{K}{e_0 + e_1}\right) + 1 - \delta \quad (8)$$

while the real wage is

$$w(K) = F_N(K, N) = f\left(\frac{K}{e_0 + e_1}\right) - \frac{K}{e_0 + e_1} f'\left(\frac{K}{e_0 + e_1}\right). \quad (9)$$

In this simple example we focus on steady-state no-growth equilibria, which we define as follows. A *generalized steady-state market equilibrium* consists of a consumption rule (c_0, c_1) , a capital stock K , and factor prices (R, w) such that (i) the consumer's budget constraints (2)-(3) are satisfied and (ii) the firm's profit-maximizing conditions (8)-(9) are satisfied.¹³ If, in addition, K maximizes $L_d(K|R, w)$, we will refer to the equilibrium as a *strict steady-state market equilibrium*, but in general we forgo this rationality condition. We consider (c_0, c_1, K, R, w) to be a generalized equilibrium as long as these variables are feasible and consistent with market-clearing. Instead of assuming that consumers independently choose to follow a consumption rule that maximizes their utility, we take the view here that the bulk of households follow consumption rules inculcated by parental or societal authorities.¹⁴ We can then ask what is the optimal consumption rule to disseminate across the population.

2.2 Optimality in the Steady State

The generalized steady-state market equilibrium concept allows us to straightforwardly define what we mean by an optimal rule of thumb for consumption and saving decisions. We can cast this problem as a social planner's problem. However, this problem differs from the usual Pareto social planner's problem since we require more than that allocations be technologically feasible. Because our households continue to interact through markets, a rule of thumb must produce an allocation that satisfies (2)-(3) and (8)-(9). Thus our social planner must choose an allocation that is both technologically feasible and supportable by markets.

Note that the social planner should only be viewed as a convenient fiction. In particular, the social planner has no actual power. We could imagine that the social planner advises households what rule of thumb to use, but we do not model the process by which households acquire their rule of thumb, so we could

¹³These conditions are easily generalized to allow for a balanced-growth equilibrium as in Section 3.

¹⁴Allen and Carroll (2001) have argued there simply is not enough time for individual consumers to adopt the LCPI rule (or even a linear approximation to it) within one lifecycle in a model with endogenous learning.

just as well imagine that it is the result of some process of natural selection.

In this simple two-period OLG model, the conditions that define a generalized market equilibrium express c_0 , c_1 , R , and w as functions of the capital stock K . Thus the social planner's problem reduces to choosing K to maximize

$$L_s(K) = u(w(K)e_0 - K) + \beta u(w(K)e_1 + R(K)K). \quad (10)$$

The distinction between (4) and (10) is that the market-clearing conditions are imposed after the decentralized household maximizes (4) whereas in (10) the market-clearing conditions are included as constraints on the optimization.¹⁵

Although both problems involve the social planner maximizing the utility of his charges, we impose different initial conditions than would arise in the Pareto problem that economists usually pose to a social planner. In the Pareto problem, the social planner considers a dynamic environment where the initial capital stock K_0 is specified. The social planner then chooses an allocation of goods across all agents that maximizes a linear combination of the utility functions of each agent. In contrast, when solving the steady-state problem, the social planner assumes the initial capital stock is such that the economy begins already in a steady state.¹⁶

The two approaches are better suited for setting policy in different environments. In an infinite-horizon representative-agent model, the Pareto problem consists of maximizing the agent's utility. In order to define lifetime utility for an infinitely-lived agent, he must have a positive discount rate. The golden-rule steady state is defined as the steady state that maximizes aggregate consumption in each period given the production function, and it is characterized by a return on capital equal to the growth rate of the economy, which in the present context is zero. Since the agent values future consumption less than current consumption, a policy that achieves the golden-rule consumption in every period will not maximize his utility. Suppose the economy starts with the golden-rule capital stock. Then the agent would do better to move some of his consumption forward in time, consuming more than the golden-rule value early on. In the long run, the utility-maximizing policy will converge to a steady state but one

¹⁵Households that follow the rule of thumb that maximizes (10) are not individually rational according to traditional terminology, which is why we classify them as irrational, but one could in some sense view them as being hyperrational.

¹⁶The solution to the steady-state problem is similar in spirit to Phelps' (1966) notion of a "commanding consumption path", i.e. a feasible consumption path that gives higher consumption some of the time and never gives less consumption than all other feasible consumption paths, irrespective of the initial state.

where consumption is less than its golden-rule value. Since there is only the one agent, it is quite sensible for a Pareto social planner to advise this dynamic policy.

On the other hand, the Pareto problem has its drawbacks for setting policy with an infinite collection of finitely-lived consumers, as in our OLG model. In order to define the Pareto problem, the social planner must first choose an arbitrary weighting scheme that determines the relative value of the utility of different generations. Although an egalitarian social planner (or economist) might prefer to weight the utility of all agents the same, to obtain a finite Pareto social-welfare function he must choose weights for each generation that converge to zero as the generational birthdate goes to infinity.

Arguments for why it is inappropriate to use different weights usually rest on ethical considerations. However, in the extreme case where consumers live only one period so a planner is necessary to ensure there is a continuing capital stock, it is easy to demonstrate why a Pareto social-welfare function would produce absurd results. Since each household consumes only one good, a policy that maximizes consumption of this good while also ensuring there is a constant stream of this good to be enjoyed by all generations would seem quite reasonable. This policy will achieve a golden-rule steady state, but it is not the solution to any Pareto problem. In solving a Pareto problem, the social planner must give less weight to the utility of future generations, so an allocation that solves such a problem must converge to a steady state where each cohort's consumption is less than the golden-rule value.

Where households live more than one period so markets can function, the downside of the golden-rule steady state is that achieving it requires tremendous sacrifice of the starting population.¹⁷ By constraining our households to continue operating through markets, we limit the burden imposed on any one cohort even while maximizing steady-state utility. We also avoid the problem of how to equip the social planner with the dictatorial power to enforce his will. Our rule of thumb can be viewed as nothing more than a coordination mechanism and does not require any new political technology.

Of course, the Pareto problem has two major selling points. First, an allocation that solves the Pareto problem has the interpretation that it is efficient under Pareto's definition that there is no way to improve upon one agent's utility without making someone worse off. Second, under conditions specified

¹⁷This will be shown explicitly below in 2.5.

by the Welfare Theorems, such an allocation can be achieved by rational agents participating in a market. These are both appealing results. Since rational agents will choose to deviate from optimal steady-state allocations, it is not clear to what extent the first result holds for optimal steady-state allocations and the second definitely does not hold.¹⁸ However, if we dispense with the assumption that all agents are rational, the second result is less relevant.

Let us define K_s to be the capital stock in the optimal generalized market equilibrium (solved for by the fictional social planner), K_d the capital stock for the (decentralized) strict market equilibrium, and K_g the golden-rule capital stock, defined by $R(K_g) = 1$. For the case of the two-period model presented above, the following proposition, proved in Appendix A, addresses the question of whether the LCPI consumption rule is optimal in the steady-state market equilibrium.

Proposition 1 *The optimal consumption rule consistent with a generalized market equilibrium coincides with a strict market equilibrium only if $K_s = K_d = K_g$. If the functional form of f and u is such that L_s is a strictly concave function, then if the strict market equilibrium is dynamically efficient (i.e. $K_d \leq K_g$) $K_s \geq K_d$, with equality only if $K_d = K_g$. If the strict market equilibrium is dynamically inefficient, the opposite inequalities hold.*

If we specialize to the case where the production function is Cobb-Douglas with share of capital strictly between 0 and 1, the condition that L_s is strictly concave will be satisfied.¹⁹ Thus the LCPI consumption rule will only be optimal for the knife-edge set of parameters where the strict steady-state market equilibrium is the golden-rule steady state.²⁰

The intuition behind Proposition 1 can best be understood by comparing the present model to one where there is a single cohort that lives for two generations with a given initial capital stock. In this single-cohort model, we will get the

¹⁸ Allocations do exist in OLG environments that are Pareto efficient but do not optimize a Pareto social welfare function. As an example, for parameterizations where the rational steady-state competitive equilibrium coincides with the golden-rule allocation, the economy is producing the maximum possible aggregate consumption, so no one can be given more consumption without taking consumption away from someone else, and everyone optimally allocates his own consumption. But this allocation will not solve any Pareto social planner's problem.

¹⁹ See Appendix B.

²⁰ A dictatorial social planner maximizing steady-state utility will set the capital stock to the golden-rule capital stock and impose a consumption allocation equivalent to the LCPI rule with $R = R(K_g) = 1$. Note that a dictatorial planner with a Pareto social welfare function will also impose a consumption allocation equivalent to the LCPI rule, though with $R(K) > 1$, so the optimal rule of thumb will generally not be Pareto optimal.

same consumption allocation in each of three regimes: (i) we assume that agents individually maximize their lifetime utility subject to market-determined budget constraints while assuming prices, which are determined by market-clearing conditions, are fixed; (ii) a social planner maximizes the cohort's lifetime utility subject to the technological feasibility constraint; (iii) a social planner maximizes the cohort's lifetime utility subject to market-determined budget constraints while accounting for the dependence of prices on the allocation. Regimes i and ii correspond respectively to the usual decentralized and usual social planner's problem, and the Welfare Theorems guarantee these problems have the same solution. Regime iii is the variation on the planner's problem that we introduce here. Its equivalence to the other regimes arises because the initial capital stock is exogenous. Only in Regime iii will the social planner exploit the pecuniary externality (since there are no factor prices in Regime ii), but with a constant-returns-to-scale production function, derivatives of factor prices with respect to the second-period capital stock cancel out. Since the initial capital stock is not a choice variable, any influence it might have on factor prices is irrelevant. Thus the solution to the problem in Regime iii is the same as in Regime ii and thereby Regime i.

On the other hand, if we solve for a steady state of the OLG model, these three problems will, in general, have three different answers. The driving distinction between the OLG and single-cohort environments is that in a steady state of an OLG model the capital stock is endogenous. If we require agents to trade through markets, the effect of the capital stock on factor prices does matter. Assuming the strict-market equilibrium is dynamically efficient, then Regime iii will have a larger optimal capital stock than Regime i because the optimization in Regime iii does take into account the pecuniary externality.²¹

2.3 Calibrating the Example

To demonstrate that Proposition 1 is not just an abstract result but can actually matter for typical calibrations of the macroeconomy, let us consider a numerical example. If we assume $u(c) = \ln c$ and a Cobb-Douglas production function with depreciation, the two-period model has only five free parameters: the share of capital α , the discount factor β , the depreciation rate δ , and the income endowments e_0 and e_1 for young and old workers respectively. We set the

²¹Unencumbered by markets, Regime ii will have the highest capital stock of the three, but the absence of markets allows for more punishing transition dynamics as we discuss in 2.5.

technology parameters to typical values from the literature: $\alpha = 1/3$ and $\delta_{an} = 0.10$. Likewise, we set $\beta_{an} = 0.96$.²² It is easier to obtain equilibria with bizarre properties, such as dynamic inefficiency, if old agents have no income endowment since young consumers will have to save at any interest rate, so we normalize e_0 to 1 and set $e_1 = 1/3$, imagining that the first period of life covers ages 25 to 55 and the second period covers the working years 55 to 65 as well as a retirement period.

In this two-period model, the determination of the capital stock K invested by young workers fully specifies the consumption rule, but it is more natural to characterize the saving rule in terms of the saving rate

$$s(K) = \frac{K}{w(K)e_0} \quad (11)$$

of young workers out of their labor income. Let s_d denote the saving rate of the strict steady-state market equilibrium, i.e. for the decentralized equilibrium where consumers make their choices individually (and rationally) to maximize the decentralized objective (4). Let s^* denote the optimal saving rate that produces the capital stock that maximizes the social planner's objective (10). Fig. 1 shows lifetime utility U as a function of the saving rate s (i.e. $U(s) = L_s(K(s))$, where $K(s)$ is the inverse of $s(K)$ as defined by (11)). The horizontal line is the lifetime utility that would be obtained under the decentralized equilibrium.

For our baseline parameters, the decentralized saving rate is $s_d = 16\%$. The optimal saving rate is $s^* = 39\%$. Fig. 1 demonstrates the optimality of these saving rates for their respective problems. The thick red curve plots lifetime utility as a function of the saving rate in general equilibrium. This is the objective $L_s(K(s))$ of the social planner and is maximized at s^* . The thin blue curve plots lifetime utility as a function of the saving rate in a partial equilibrium where prices are held fixed at their decentralized equilibrium values. This is the objective $L_d(sw(K(s_d))e_0)$ of a rational household in the decentralized equilibrium and is maximized at s_d .

There is a wide range of saving rates between $s_d = 16\%$ and 65% that give rise to equilibria where the consumer is better off than in the decentralized equilibrium. Since we do not have any growth of population or technology in

²² A period is viewed as lasting for $T = 30$ years, so the per-period discount factor is $\beta = \beta_{an}^T$ and the per-period depreciation rate is $\delta = 1 - (1 - \delta_{an})^T$. All time-dimensionable quantities specified in the text will be given in annual terms.

this simple model, a generalized market equilibrium will be dynamically efficient if the net interest rate r is nonnegative. Fig. 2 shows the equilibrium interest rate (per annum) as a function of the saving rate s and for the decentralized equilibrium. Not surprisingly, because of diminishing returns to capital, the interest rate is a decreasing function of the saving rate. Note that the golden-rule saving rate such that $r = 0$ is $s_{gr} = 69\%$, so for this example all saving rates that improve upon the decentralized equilibrium also lead to dynamically efficient equilibria, including the optimal saving rate s^* .

It is difficult to interpret how much of an improvement the change in lifetime welfare from $U = -2.29$ for the decentralized equilibrium to $U = -2.05$ for the optimal generalized equilibrium actually confers. However, Fig. 2 also demonstrates that the variation in interest rates between these two equilibria is quite large, for r decreases from 4.78% per annum to 1.83%. We see similar variation in other important macroeconomic observables. Fig. 3 shows the variation of aggregate output Y as a function of the saving rate. Going from the decentralized equilibrium to the optimal saving rate, output increases by 50%. Thus if consumers coordinate on a higher saving rate than is accorded by the LCPI rule, the economy can achieve a remarkable degree of greater prosperity. Even if consumers only deviate from the LCPI rule by a modest amount, say increasing their saving rate to 18%, output would increase by 5%.

Nevertheless, while the consumers in these generalized market equilibria are better off than they would be in the decentralized equilibrium, as individuals they could increase their lifetime utility even more if they maximize their own utility. The dashed purple curve in Fig. 1 plots lifetime utility as a function of the saving rate in a partial equilibrium where prices are held fixed at their values from the optimal generalized equilibrium. This corresponds to $L_d(sw(K(s^*))e_0)$, which would be the objective function of a rational household if everyone else follows the optimal rule of thumb. It is maximized at 7.5%, the saving rate prescribed by the LCPI rule, for which the rate of consumption growth c_1/c_0 equals βR . If households save $s^* = 39\%$, their consumption will increase by 65% from the first period to the second, but $\beta R = 0.50$ in per-period terms. Because the interest rate is so low, the consumer could increase his welfare from -2.05 to -1.86 if he saves only 7.5%, enjoying the bulk of his consumption while young and letting his consumption drop by 50% when he is old. If only a couple agents deviate in this way, this should not affect aggregate quantities. But if a measurable fraction of the population behaves rationally, this will have a negative impact on the equilibrium. So next we consider what




Figure 1: Lifetime utility in the baseline two-period model as a function of the saving rate s in general equilibrium and in partial equilibrium, holding the prices fixed at both the decentralized equilibrium prices and the social planner's equilibrium prices. The dotted line corresponds to lifetime utility in the decentralized equilibrium.




Figure 2: Net interest rate r (per annum) in the baseline two-period model as a function of the saving rate s and for the decentralized equilibrium. The interest rate for the optimal saving rate s^* is represented by a large dot.




Figure 3: Aggregate output Y in the baseline two-period model as a function of the saving rate s and for the decentralized equilibrium. Output for the optimal saving rate s^* is represented by a large dot.

happens if rational agents invade the population.

2.4 An Influx of Rational Agents

Now suppose that a fraction μ of the population is rational and chooses its saving

$$k_r(R, w) = \frac{w}{1 + \beta} \left[\beta e_0 - \frac{e_1}{R} \right] \quad (12)$$

to maximize $L_d(k_r | R, w)$ given the market-determined R and w . The remaining fraction $1 - \mu$ follows the optimal saving rate s^* that maximizes $L_s(K(s))$ and chooses its saving according to

$$k_i(w) = s^* w e_0 \quad (13)$$

given w .²³ The aggregate capital stock will be

$$K = \mu k_r + (1 - \mu) k_i.$$

²³We could also consider the case of a smarter social planner who takes into account the fact that some agents will deviate from the rule of thumb. This fine tuning would improve welfare even more for the irrational agents, but it would require information that may be difficult to obtain.

A generalized market equilibrium will then be determined by the capital stock K that solves the market-clearing condition for capital

$$K = \mu k_r(R(K), w(K)) + (1 - \mu)k_i(w(K)).$$

As a concrete example, let us consider what happens if we inject rational agents into the model under the calibration of 2.3. Fig. 4 graphs the lifetime utility of both rational and irrational agents as a function of the fraction of rational agents μ . As we discussed above, the rational agents always do better than the irrational agents and, moreover, always do better than in the decentralized equilibrium. Indeed, the rational agents are essentially taking advantage of the irrational agents, who through their high saving allow the economy to enjoy high wages. However, unlike in most models where rational and irrational agents coexist, while the rational agents end up with higher lifetime utility they do not end up with higher wealth. Fig. 5 shows how the individual saving of the two types of agents varies with μ . In all cases, the irrational agent saves substantially more than the rational agent.²⁴

Nevertheless, notice that while the lifetime utility of both rational and irrational agents decline with μ , they both do so nearly linearly. A small influx of rational agents does not destroy the result that the rule of thumb can improve upon the decentralized equilibrium. Indeed, as long as less than two thirds of the population is rational, both the rational and irrational agents continue to enjoy greater lifetime utility than they could attain in the decentralized equilibrium where everyone is rational.²⁵

Fig. 6 shows how aggregate output varies with the fraction of rational agents μ . For the baseline parameters, output is a monotonic, nearly linear function of μ . Thus while rational agents diminish the material prosperity that the social planner could achieve if everyone followed his advice, they only do so in proportion to μ . Macroeconomic quantities behave continuously in the limit as $\mu \rightarrow 0$ and are not especially sensitive to small changes in μ .

2.5 Transition Dynamics

Since a common criticism of steady-state optimality criteria (Arrow (1999)) is that they lead the social planner to immiserate early cohorts in order to

²⁴Of course, in this model there is no advantage to having more wealth.

²⁵Interestingly, the mean utility across the population is higher than the decentralized utility for all values of μ .




Figure 4: Lifetime utility for both rational and irrational agents in the baseline two-period model as a function of the fraction of rational agents μ . Lifetime utility for the decentralized equilibrium is the dashed line.




Figure 5: Saving (in absolute terms) by a rational agent k_r and by an irrational agent k_i as functions of the fraction of rational agents μ for the two-period baseline model.




Figure 6: Aggregate output Y in the baseline two-period model as a function of the fraction of rational agents μ .

build up the optimal capital stock, we conclude our discussion of the two-period model by considering what happens in various transitional experiments. As in the Solow (1956) model, if consumers follow a given consumption rule and if there is a unique steady-state equilibrium consistent with this rule, the economy will converge to this steady state from most initial capital stocks. We can therefore consider experiments where the economy starts in either the strict market equilibrium or a steady state market equilibrium with a saving rate even less than the LCPI rate, and we end in the steady state of one of the three regimes discussed at the end of 2.2.

For the remainder of this section, we assume there is no income in old age, so $e_1 = 0$, and we can normalize $e_0 = 1$. This assumption greatly simplifies the transition dynamics under the LCPI rule, which only depends on current income. Given an initial capital stock K_0 , we then define a generalized dynamic market equilibrium as a sequence of consumption rules $\{c_{0,t}, c_{1,t}\}_{t=0}^{\infty}$, a sequence of factor prices $\{R_t, w_t\}_{t=0}^{\infty}$, and a saving rule $K_{t+1}(w_t)$ such that (i) the budget constraints

$$\begin{aligned} c_{0,t} + K_{t+1} &= w_t \\ c_{1,t} &= R_t K_t \end{aligned}$$

are satisfied for all t , and (ii) the profit-maximizing conditions

$$\begin{aligned} w_t &= w(K_t) \\ R_t &= R(K_t) \end{aligned}$$

are satisfied for all t .

We then consider two types of generalized dynamic market equilibria. Given factor prices R_{t+1} (which turns out to be irrelevant under our assumptions) and w_t faced by an individual born at time t , the saving rule that maximizes (1) is the rational, LCPI saving rule

$$K_{t+1}^r = \frac{\beta}{1 + \beta} w_t, \quad (14)$$

which also happens to be a constant saving rule.²⁶ Meanwhile, if we have irrational consumers who follow the rule of thumb that is optimal in the steady state, the saving rule will be

$$K_{t+1}^i = s^* w_t, \quad (15)$$

where s^* is the optimal saving rate in the steady state. We will also consider the transition to a golden-rule steady state, but this cannot be described by a generalized dynamic market equilibrium.

Aside from setting $e_1 = 0$, we consider the same calibration as in 2.3. The optimal saving rate remains $s^* = 39\%$ and confers steady-state lifetime utility of -2.06. In this special case, the LCPI rule is also a constant saving rate rule with $s^r = \beta/(1 + \beta) = 23\%$ and confers steady-state lifetime utility of -2.18. Since most evidence suggests present-day households save less than is dictated by the LCPI rule, we also consider a steady-state market equilibrium with a very low saving rate of $s = 3.5\%$, which confers lifetime utility of -3.15.

First, suppose that the economy is in the strict market steady state for $t < 0$, so consumers are initially rational. At $t = 0$, an unrestricted social planner comes to power and decides to bring the economy as quickly as possible to a golden-rule steady state. We assume he does this by forcing the cohort born at $t = 0$ to save the golden-rule capital stock.²⁷ The time path of lifetime utility

²⁶If $e_1 > 0$, we would have to account for the present value of future income $w(K_{t+1})e_1/R(K_{t+1})$, and we would not obtain an analytic expression relating K_{t+1} to K_t .

²⁷There are actually infinitely many transition paths to the golden-rule steady state. All require some generations to pay a utility cost, but the social planner is indifferent between

in this experiment is the thick dashed line shown in Fig. 7. From $t = 1$ onward, the economy will be in the golden-rule steady state, and consumers will enjoy lifetime utility of -1.89. However, the cohort born at $t = 0$ pays a high cost for the prosperity of its progeny because its lifetime utility is -3.50, which is lower even than the lifetime utility under the 3.5% saving rule and considerably lower than what it would achieve if it continued to follow the LCPI rule. Thus, in order to achieve the golden-rule steady state, the unrestricted social planner must essentially enslave the first cohort.

In contrast, if households adopt the optimal rule of thumb, the first generation will only reduce its lifetime utility from -2.18 to -2.36. Lifetime utility goes down because the higher capital stock results in a lower interest rate. Every generation thereafter earns higher wages and enjoys lifetime utility higher than the old steady-state value. In the long run, lifetime utility converges to the new steady-state value of -2.06, as is shown in Fig. 7.

We see a similar pattern if the economy starts in the steady state with the 3.5% saving rule. The first generation has to reduce its lifetime utility to -3.38 (which is still above the level required to achieve the golden-rule steady state), but everyone is better off than in the initial steady state from there on out.

Note that the utility loss suffered by the first generation is not a property of following a rule of thumb.²⁸ If households start out by irrationally saving 3.5% and then, at $t = 0$, become rational, the first generation will still reduce its lifetime utility to -3.19, again because of the resulting decrease in interest rates.²⁹ While it is true that adopting the optimal saving rule will involve some pain for early adopters, the same would also be true if economists succeeded at their current objective of teaching more households to behave rationally.³⁰

3 SMarT Plan in General Equilibrium

While the two-period OLG model is simple enough to allow us to prove the abstract result of Proposition 1, which establishes that a consumption rule must

these paths.

²⁸We do not consider the transition from the 3.5% steady state to the golden-rule steady state because the initial output and capital are too low to build up the golden-rule capital stock in one step.

²⁹If the initial saving rate is unrealistically low, it is possible that implementing the LCPI rule can be done without any utility loss by the first generation.

³⁰It would be possible to spread this pain over several generations by gradually ramping up the prescribed saving rate over time.




Figure 7: Lifetime utility as a function of cohort birth date in four transitional experiments: starting from a steady state with a 3.5% saving rule and going to the LCPI rule and the optimal saving rule; and starting from the LCPI rule and going to the optimal saving rule and the golden-rule steady state.

exist that improves upon the LCPI rule in general equilibrium, many results about OLG models are specific to the case where agents only live for two periods. In particular, since agents from different cohorts can only ever meet once, it is not possible for them to engage in intertemporal trade, which can force behaviors that would not otherwise be optimal. To show that the results of the previous section only depend on the OLG aspect of the model and not the special assumption that agents live for two periods, we now consider the opposite extreme of a continuous-time OLG model. This greatly enlarges the space of possible consumption rules since consumers now face multiple saving decisions. We focus on the one-dimensional family of rules that fall under the heading of Thaler and Benartzi's (2004) Save More Tomorrow or SMarT Plan, which policymakers actually have encouraged people to adopt.³¹

³¹The SMarT plan consumption rule, as presented in 3.1, is actually characterized by two parameters: the SMarT plan saving rate and the base saving rate. However, we would view the base saving rate as an exogenous parameter since it describes how consumers would behave if left to their own devices. Only the SMarT plan saving rate is a policy instrument under the control of our fictional social planner.

3.1 The Continuous-Time Lifecycle Model

First, we describe the environment faced by an individual agent “born” at time $\tau \in \mathbf{R}$, which is a generalization of the model in Bullard and Feigenbaum (2007). The agent starts work at τ , retires at time $\tau + T$, and passes away at time $\tau + \bar{T}$, where $0 < T < \bar{T}$. During his working life, the agent earns the real wage $w(t)$, which grows at the rate x that corresponds to the rate of technological growth for the economy. The agent can borrow or save at the rate of return r .³² Let $c(t, \tau)$ denote the flow of consumption at t of an agent born at τ and $k(t, \tau)$ denote the corresponding stock of savings. Then the agent faces the differential budget constraint

$$\frac{dk(t, \tau)}{dt} = \begin{cases} rk(t, \tau) + w(t) - c(t, \tau) & t \in [\tau, \tau + T] \\ rk(t, \tau) - c(t, \tau) & t \in [\tau + T, \tau + \bar{T}] \end{cases} \quad (16)$$

along with the boundary conditions

$$k(\tau, \tau) = k(\tau + \bar{T}, \tau) = 0. \quad (17)$$

Analogous to the two-period model, we assume that the agent values consumption allocations according to the utility function

$$U(\tau) \equiv \int_0^{\bar{T}} e^{-\rho t} u(c(\tau + t, \tau); \sigma) dt, \quad (18)$$

where

$$u(c; \sigma) = \begin{cases} \frac{1}{1-\sigma} c^{1-\sigma} & \sigma \neq 1 \\ \ln c & \sigma = 1 \end{cases}$$

for $\sigma > 0$. However, the consumer is not able to compute the consumption rule that maximizes (18) subject to (16) and (17). Instead, during the working life the agent would, absent any direction, follow the rule of thumb of saving a constant fraction of income. After retirement, we assume the agent’s savings account pays out the annuity value of the account,

$$a(\tau) = \frac{\int_{\tau}^{\tau+T} (w(t) - c(t, \tau)) e^{r(\tau+T-t)} dt}{\int_{\tau+T}^{\tau+\bar{T}} e^{r(\tau+T-t)} dt}, \quad (19)$$

³²Both $w(t)$ and r are market-determined as we describe in 3.2.

that confers a perfectly smooth flow of consumption while running savings down to zero at the time of death $\tau + \bar{T}$. Thus the basic consumption rule is

$$c_b(t, \tau|s) = \begin{cases} (1-s)w(t) & t \in [\tau, \tau + T] \\ a(\tau) & t \in [\tau + T, \tau + \bar{T}] \end{cases} .$$

Now suppose the social planner implements a SMarT Plan that requires agents to save at a higher rate $\gamma \geq s$ from any wage increases at $t \geq \tau$. Thus under the SMarT Plan, the consumption rule is

$$c_S(t, \tau|s, \gamma) = \begin{cases} w(t) - (sw(\tau) + \gamma[w(t) - w(\tau)]) & t \in [\tau, \tau + T] \\ a(\tau) & t \in [\tau + T, \tau + \bar{T}] \end{cases} .$$

Note that this nests the basic consumption rule with $\gamma = s$. The path of the individual's savings account is

$$k(t, \tau) = \int_{\tau}^t \{sw(\tau) + \gamma[w(u) - w(\tau)]\} e^{r(t-u)} du \quad t \in [\tau, \tau + T] \quad (20)$$

$$k(t, \tau) = \int_t^{\tau + \bar{T}} a(\tau) e^{r(t-u)} du \quad t \in [\tau + T, \tau + \bar{T}] . \quad (21)$$

3.2 The Continuous-Time OLG Model

In partial equilibrium, the LCPI rule is necessarily strictly superior to any other rule since it is the rule obeyed by rational agents.³³ However, in an OLG model, market interactions between agents of different cohorts can cause the LCPI rule to confer strictly less lifetime utility than other consumption rules in general equilibrium. Thus we now consider what happens if we mix agents born at different times but still described by the continuous-time lifecycle model.

At each instant t a new cohort of population $N(t)$ is born and the oldest cohort, born at $t - \bar{T}$ dies. The size of each successive cohort grows at rate n (and hence the total population also grows at rate n), so the size of a cohort born at time t is $N(t) = e^{nt}$, where we normalize the cohort population at $t = 0$ to one.

We endogenize factor prices by incorporating a production sector that employs capital and labor as factors. In equilibrium, the capital stock $K(t)$ is

³³Findley and Caliendo (2007) have studied the welfare effects of participation in a SMarT plan under partial equilibrium.

equated to the aggregate sum of savings of consumers

$$K(t) = \int_{t-\bar{T}}^t N(t)e^{n(\tau-t)}k(t, \tau)d\tau. \quad (22)$$

Capital also depreciates at the rate δ . Unretired consumers inelastically supply one unit of labor, so the labor supply is

$$L(t) = \int_{t-T}^t N(t)e^{n(\tau-t)}d\tau = L(0)e^{nt}. \quad (23)$$

Firms behave competitively and have a constant returns to scale production function

$$Y(t) = K(t)^\alpha [A(t)L(t)]^{1-\alpha},$$

where $A(t) = e^{xt}$ is a labor-augmenting technology factor that grows at the constant rate x . Factor prices are determined competitively:

$$r(t) = \alpha \left(\frac{K(t)}{A(t)L(t)} \right)^{\alpha-1} - \delta \quad (24)$$

$$w(t) = (1 - \alpha) \left(\frac{K(t)}{A(t)L(t)} \right)^\alpha A(t). \quad (25)$$

For this model, a generalized steady-state market equilibrium is characterized by a consumption rule $c(t, \tau) = c_0(t - \tau)e^{xt}$ in which consumption depends only on the household's age $t - \tau$ (after accounting for technological growth), a saving rule $k(t, \tau) = k_0(t - \tau)e^{xt}$, a time-independent interest rate r , and a wage function $w(t) = w_0e^{xt}$ such that (i) $c(t, \tau)$ and $k(t, \tau)$ satisfy (16) and (17) given r and $w(t)$ and (ii) r and $w(t)$ satisfy (24) and (25) given $k(t, \tau)$. For comparison with the LCPI rule, we will also consider strict steady-state market equilibria, which satisfy the additional condition that $c(t, \tau)$ and $k(t, \tau)$ maximize (18) under the constraints (16) and (17).

3.3 The Optimal SMarT Plan Saving Rate

Following the principle of optimality in the steady state, let us consider the optimal SMarT Plan saving rate γ if we calibrate the model using parameter values typical of the literature on general-equilibrium lifecycle models. The rate of population growth n is set to 1% to match the U.S. experience over the last few decades. Following Bullard and Feigenbaum (2007), the rate of technology

growth x is set to 1.56%. The technology parameters are set to common values of $\alpha = 0.35$ for the share of capital and $\delta = 8\%$ for the depreciation rate. The worklife T is 40 years and the lifespan $\bar{T}$ is 55 years, which if we assume that the working life begins at age 25 is consistent with an average lifespan of roughly 80 years. Finally, we set the preference parameters to $\sigma = 2$ for the inverse elasticity of intertemporal substitution and $\rho = 2\%$ for the discount rate, which are both near the middle of the broad range of values that researchers have considered for these difficult-to-pin-down parameters.

Figure 8 plots three curves. The first is a flat line (black) set equal to the level of lifetime utility in the strict steady-state market equilibrium where the LCPI rule is obtained by maximizing $U(\tau)$ with our calibration of σ and ρ . The other two curves depict lifetime utility in steady-state market equilibria where consumers participate in a SMarT Plan as a function of the SMarT saving rate γ . For the dashed, blue curve, we assume consumers have a default saving rate of $s = 7\%$ whereas for the solid, red curve $s = 3.5\%$. We see that the optimal SMarT rate is near 60 percent when $s = 7\%$ and is near 70 percent when $s = 3.5\%$.

It is clear that the result proved in Section 2 is not academic. It can be generalized to more realistic and empirically relevant models with consumption rules that are straightforward enough that consumers could practically implement them. Economy-wide participation in a program such as the SMarT Plan can push an economy in general equilibrium from a level of utility that is initially below the permanent-income utility (i.e. when $\gamma = s$) to a new level that is even higher than the permanent-income utility.

Figure 9 is exactly the same as Figure 8 except that the discount rate ρ has been lowered from 2% to 0%. Comparing Figures 8 and 9, we find that the optimal SMarT rate is not terribly sensitive to the discount rate. However, in this case we find that the LCPI rule achieves superior lifetime utility than the SMarT Plan for any value of the SMarT saving rate γ . This does not, however, imply that there is no consumption rule that can outperform the LCPI rule in general equilibrium since we have only considered a one-dimensional subspace of the entire space of feasible consumption rules. Indeed, for the high-saving economy, the optimal SMarT Plan comes very close to matching the utility of the LCPI rule, so it is likely that a small perturbation of the SMarT Plan rule could outperform the permanent-income economy.

One criticism of the SMarT Plan is that if consumers opt to choose a SMarT saving rate γ that is too high then they may actually do worse than if they stick




Figure 8: Lifetime utility for three types of economies with $\sigma = 2$ and $\rho = 0.02$. Thin black line is the permanent-income rule. Blue dashed and red solid lines are rule-of-thumb economies respectively with high (7%) and low (3.5%) base saving rates.




Figure 9: Lifetime utility for three types of economies with $\sigma = 2$ and $\rho = 0$. Thin black line is the permanent-income rule. Blue dashed and red solid lines are rule-of-thumb economies respectively with high (7%) and low (3.5%) base saving rates.

to the default rule of thumb of saving a small fraction of their income. In partial equilibrium, Findley and Caliendo (2007) have shown that this criticism may be warranted. Here though we see that if all consumers adopt the same SMarT saving rate then in general equilibrium the SMarT Plan will outperform the default rule for any $\gamma > s$.

To show the importance of the pecuniary externality for arriving at these results, we repeat the $\rho = 0.02$ experiment in partial equilibrium, keeping factor prices constant at their general-equilibrium values for the permanent-income economy. The results are depicted in Figure 10. The dashed blue line shows lifetime utility for a SMarT Plan partial-equilibrium economy with a base saving rate $s = 3.5\%$, plotted as a function of the SMarT saving rate γ . That is, we fix factor prices at their values from the permanent-income economy and we do not allow these factor prices to change as we adjust the SMarT rate. For comparison, the straight black line and solid red line are the same as in Figure 8. The former is lifetime utility for the permanent-income economy, which, by construction, has not changed. The solid red line is the level of lifetime utility in a SMarT Plan economy with general-equilibrium effects included. It is clear that in partial equilibrium the rule-of-thumb economy can never dominate the permanent-income one, but it is also interesting that even in partial equilibrium SMarT Plan participation can take lifetime utility *almost* to the permanent-income level.³⁴

One concern that needs to be addressed is whether these steady-state market equilibria that dominate their permanent-income counterparts might have too much capital. That is, are they dynamically inefficient? Is it feasible to achieve the same aggregate consumption with less saving? In Figure 11 we plot the equilibrium rates of return for our baseline economies. The solid black line is the equilibrium rate of return from the permanent-income economy. The dotted line is the rate of growth of aggregate income $n + x$. An economy will be dynamically inefficient if the rate of return is less than this growth rate. Clearly, the permanent-income economy is not dynamically inefficient. The blue and red lines give the equilibrium interest rates for SMarT Plan economies as a function of the SMarT saving rate. The solid red line corresponds to a base saving rate of $s = 3.5\%$ and the dashed blue line to $s = 7\%$. The large dots mark the point on the graph for each SMarT Plan economy where the SMarT saving rate is optimal. Both of these points are above the dotted line, indicating these

³⁴This last result is not surprising, however, given the partial equilibrium treatment of SMarT Plans in Findley and Caliendo (2007).




Figure 10: Lifetime utility for $\sigma = 2$, $\rho = 0.02$, and base saving rate $s = 3.5\%$. Thin black line is for permanent-income consumer. Blue dashed and red solid lines are rule-of-thumb economies respectively in partial and general equilibrium.

equilibria are also dynamically efficient. At even higher SMarT Plan saving rates, we do obtain dynamically inefficient equilibria. Notice, however, that some of these dynamically inefficient equilibria still deliver higher utility than the permanent-income economy.³⁵

While we have established that a social planner can improve upon a permanent-income economy with respect to lifetime utility if he convinces households to adopt a SMarT Plan with a sufficiently high saving rate, utility is not observable, so to close this section let us consider how sensitive macroeconomic observables are to the SMarT Plan saving rate. Fig. 12 plots steady state values of equilibrium output Y as a function of the saving rate γ both with a base saving rate s of 3.5% (red solid line) and of 7% (blue dashed line). For comparison, the strict-market equilibrium value is represented by the thin black line. The effect on output of going from the LCPI rule to the optimal SMarT plan is not as pronounced as in the corresponding experiment for the two-period model,

³⁵It is possible in this model for the equilibrium with the optimal SMarT rate to be dynamically inefficient while also delivering a level of lifetime utility that dominates the permanent-income rule. For example, if the base saving rate is 15 percent, the optimal SMarT rate is near 40 percent, and the level of utility exceeds the permanent-income level but the rate of return in the rule-of-thumb economy is below the growth rate of the economy. Thus, some dynamically inefficient parameterizations can dominate dynamically efficient ones.




Figure 11: Equilibrium interest rate for $\sigma = 2$, $\rho = 0.02$. Thin solid black line is for permanent-income economy. Blue dashed and red solid lines are rule-of-thumb economies for high (7%) and low (3.5%) base saving rates. Big circles correspond to optimal SMarT saving rates. Thin dashed black line is the exogenous growth rate (population plus technology).

pictured in Fig. 3. Here, output only increases from the permanent-income economy by 7% for the two calibrations. It is very plausible that the two-period model exaggerates the possible gains from exploiting the pecuniary externality, which is why serious investigation of lifecycle models requires finer periods. Of course, the optimal parameterization of the SMarT Plan is not the optimal consumption rule. Additional gains could be obtained if the social planner varied the consumption rule along other dimensions.

4 Conclusion

Contrary to conventional intuition, we have shown it is possible in a general-equilibrium OLG model for consumers to attain higher lifetime utility than they would obtain if they all behaved rationally. This is still possible even if a large fraction of agents, in some cases even a majority, do behave rationally. Moreover, this is not just an abstract result that involves peculiar rules of thumb




Figure 12: Equilibrium aggregate output Y for $\sigma = 2$, $\rho = 0.02$. Thin black line is for permanent-income economy. Blue dashed and red solid lines are rule-of-thumb economies for high (7%) and low (3.5%) base saving rates. Big circles correspond to optimal SMarT saving rates.

that no one in their right mind would consider. On the contrary, we can improve upon the LCPI rule simply by implementing the Save More TomorrowTM Plan with a sufficiently high rate of saving out of future raises.

Our results imply the logical underpinnings of the rational paradigm are quite wobbly. Most arguments for assuming agents are rational start from the notion that humans will evolve toward optimal behavior (Friedman (1953)). However, natural selection will seek an optimal steady state for the species, not an optimal solution to the Pareto problem. Once a generation procreates, its own well-being is irrelevant to the propagation of the species. Indeed, nature does put a heavy burden that has been likened to a prison sentence on human parents during the gestation and upbringing of their children. It is difficult to imagine a Pareto social planner designing us to invest so much time and effort into an activity that is not essential to our own survival.

If a large fraction of households do follow rules of thumb, then a policy instrument which political leaders need to be mindful of is their influence over these rules of thumb. Cross-country comparisons reveal that American households often save a smaller fraction of their income than their counterparts abroad. This

fact is hard to account for in rational-choice models without invoking ad hoc explanations such as that preferences differ across countries, which is difficult to test. However, it is much easier to observe what rules of thumb governments advocate. If we consider the history of the United States, during World War II citizens on the home front were repeatedly exhorted to buy war bonds. In contrast, during the recent Middle-East conflicts, Americans were urged to spend, spend, spend. There has not been a conscious effort by policymakers in the United States to encourage people to coordinate on rules of thumb that involve high saving. While there have been various attempts to enact policies that encourage saving such as the introduction of IRAs, at the same time these are often undermined by other policies such as high taxes on capital. Rather than focusing so much on the positive question of how consumers make consumption-saving decisions, perhaps economists should take a closer look at the issue of how to guide consumers toward increasing their saving.

Much of the benefit from the optimal rule of thumb comes from the high wages earned by young workers as a consequence of the large capital stock produced by their forebears, but the first generation to save more will not enjoy this benefit and, moreover, their additional saving will act to reduce the return on their saving. Note that this embodiment of Keynes' (1936) paradox of thrift is not just a property of irrational consumption rules. If a large number of consumers is spurred to start saving more for any reason, the pioneers will suffer. This is true even if consumers adopt the LCPI rule. Nevertheless, this need not preclude educators and policymakers from urging the population to save more. Saving is difficult precisely because it involves sacrifice, and again we do see parents making huge sacrifices so their children will have the opportunity to earn higher wages. A more realistic model is required to quantify how much trailblazing savers must give up to put the economy on a path to an optimal steady state and to determine whether the government might redistribute the gains from this saving so that every generation can be made better off.

The optimal consumption rules that we identify here involve considerable deferment of consumption to the future and large deviations from the LCPI rule. In practice, it may not be possible to convince people to move so far away from rational behavior. For example, in many bounded-rationality models the size of deviations from rationality are controlled by decision-making and information-processing costs. Even if we cannot achieve the optimal steady-state equilibrium, it would still be advantageous to urge consumers to choose rules of thumb that perturb away from the LCPI rule in a direction that increases

the aggregate capital stock, as opposed to rules of thumb that squander societal resources.

Another question for future research is whether there are equilibria where fully rational agents follow the optimal rule of thumb. The OLG model considered here resembles the Prisoner's Dilemma in that the Nash equilibrium for both games is dominated by strategy combinations that improve the utility of all players. In the infinitely-repeated Prisoner's Dilemma, Pareto superior allocations can be supported by strategies that punish deviation from these allocations. The OLG model is not an infinitely-repeated game, so the Folk Theorem does not apply. However, agents only face a temptation to deviate from the optimal rule of thumb when they are young, so if agents are sufficiently patient one would expect it should be possible to compel rational agents to adhere to this consumption rule with the threat of punishment when the agent is old.

A Proof of Optimal Two-Period Consumption Rule

Given w and R ,

$$L'_d(K|R, w) = -u'(we_0 - K) + \beta Ru'(we_1 + RK),$$

and

$$L''_d(K|R, w) = u''(we_0 - K) + \beta R^2 u''(we_1 + RK) < 0$$

since we have assumed u is strictly concave. Thus L_d is strictly concave, and

$$K_d = \arg \max L_d(K|R, w) \tag{26}$$

is unique, satisfying

$$u'(we_0 - K_d) = \beta Ru'(we_1 + RK_d). \tag{27}$$

In particular, K_d must satisfy (27) in a generalized market equilibrium where $w = w(K_d)$ and $R = R(K_d)$.

Since we have assumed u and f are twice-differentiable, L_s is differentiable. Thus a maximum of the social planner's objective (10) must satisfy $L'_s(K) = 0$, where

$$\begin{aligned} L'_s(K) &= (w'(K)e_0 - 1)u'(w(K)e_0 - K) \\ &\quad + \beta(w'(K)e_1 + R'(K)K + R(K))u'(w(K)e_1 + R(K)K). \end{aligned} \quad (28)$$

Substituting (27) into (28),

$$\begin{aligned} L'_s(K_d) &= \beta [R(K_d)(w'(K_d)e_0 - 1) + w'(K_d)e_1 + R'(K_d)K_d + R(K_d)] \\ &\quad \times u'(w(K_d)e_1 + R(K_d)K_d) \\ &= \beta [w'(K_d)(R(K_d)e_0 + e_1) + R'(K_d)K_d] u'(w(K_d)e_1 + R(K_d)K_d). \end{aligned}$$

Using (8) and (9),

$$R'(K) = \frac{1}{e_0 + e_1} f'' \left(\frac{K}{e_0 + e_1} \right)$$

and

$$\begin{aligned} w'(K) &= \frac{1}{e_0 + e_1} f' \left(\frac{K}{e_0 + e_1} \right) - \frac{1}{e_0 + e_1} f' \left(\frac{K}{e_0 + e_1} \right) \\ &\quad - \frac{K}{(e_0 + e_1)^2} f'' \left(\frac{K}{e_0 + e_1} \right) \\ &= - \frac{K}{(e_0 + e_1)^2} f'' \left(\frac{K}{e_0 + e_1} \right) > 0. \end{aligned}$$

Thus

$$\begin{aligned} (e_0 + e_1)w'(K) + R'(K)K &= \frac{K}{e_0 + e_1} f'' \left(\frac{K}{e_0 + e_1} \right) - \frac{K}{e_0 + e_1} f'' \left(\frac{K}{e_0 + e_1} \right) \\ &= 0. \end{aligned}$$

Therefore

$$L'_s(K_d) = \beta(R(K_d) - 1)w'(K_d)e_0 u'(w(K_d)e_1 + R(K_d)K_d).$$

Since u and w are strictly increasing, $L'_s(K_d)$ has the same sign as $R(K_d) - 1$. K_d can only be the optimal saving if $R(K_d) = 1$. Since there is no population or technological growth in this simple model, a net interest rate of zero characterizes the golden-rule steady state. Thus the optimal steady state

is a strict market equilibrium only if the strict market equilibrium coincides with the golden-rule steady state. More generally, if we assume that L_s is strictly concave, then there is a unique solution to the social planner's problem. If the strict market equilibrium is dynamically efficient, $R(K_d) \geq 1$, so $L'_s(K_d) \geq 0$. Thus if the strict market equilibrium is not the golden-rule steady state, the optimal capital stock will be greater than K_d . Conversely, if the strict market equilibrium is dynamically inefficient, the optimal capital stock will be less than K_d .

B Concavity of Social Planner's Objective

Suppose that the production function is Cobb-Douglas with share of capital $\alpha \in (0, 1)$, so $f(k) = k^\alpha$. Then

$$\begin{aligned} w(K) &= (1 - \alpha) \left(\frac{K}{N} \right)^\alpha \\ R(K) &= \alpha \left(\frac{K}{N} \right)^{\alpha-1} + 1 - \delta, \end{aligned}$$

and the derivative of the social planner's objective function simplifies to

$$L_s(K) = u(w(K)e_0 - K) + \beta u(w(K)e_1 + R(K)K).$$

$$\begin{aligned} L'_s(K) &= (w'(K)e_0 - 1)u'(w(K)e_0 - K) \\ &\quad + \beta[w'(K)e_1 + R'(K)K + R(K)]u'(w(K)e_1 + R(K)K) \end{aligned}$$

$$\begin{aligned} L''_s(K) &= w''(K)e_0u'(w(K)e_0 - K) + (w'(K)e_0 - 1)^2u''(w(K)e_0 - K) \\ &\quad + \beta[w'(K)e_1 + R'(K)K + R(K)]^2u''(w(K)e_1 + R(K)K) \\ &\quad + \beta[w''(K)e_1 + R''(K)K + 2R'(K)]u'(w(K)e_1 + R(K)K) \end{aligned}$$

Since

$$\begin{aligned} w'(K) &= \frac{\alpha(1-\alpha)}{N} \left(\frac{K}{N}\right)^{\alpha-1} \\ w''(K) &= -\frac{\alpha(1-\alpha)^2}{N^2} \left(\frac{K}{N}\right)^{\alpha-2} < 0 \end{aligned}$$

and

$$\begin{aligned} R'(K) &= \frac{\alpha(\alpha-1)}{N} \left(\frac{K}{N}\right)^{\alpha-2} \\ R''(K) &= \frac{\alpha(\alpha-1)(\alpha-2)}{N^2} \left(\frac{K}{N}\right)^{\alpha-3} \end{aligned}$$

$$\begin{aligned} R''(K)K + 2R'(K) &= \frac{2\alpha(\alpha-1)}{N} \left(\frac{K}{N}\right)^{\alpha-2} + \frac{\alpha(\alpha-1)(\alpha-2)}{N} \left(\frac{K}{N}\right)^{\alpha-2} \\ &= \frac{\alpha^2(\alpha-1)}{N} \left(\frac{K}{N}\right)^{\alpha-2} < 0, \end{aligned}$$

we have that $L''_s(K) < 0$, so the social planner's objective is strictly concave.

References

- [1] Allen, Todd W. and Christopher D. Carroll (2001), Individual Learning about Consumption. *Macroeconomic Dynamics* 5: 255-271.
- [2] Arrow, Kenneth J. (1999), "Discounting, Morality, and Gaming," in *Discounting and Intergenerational Equity*, ed. by Paul R. Portney and John P. Weyant (Harvard University Press: Cambridge, MA) (Resources for the Future: Washington, DC).
- [3] Benartzi, Shlomo and Richard H. Thaler, (2001), "Naive Diversification Strategies in Defined Contribution Saving Plans," *American Economic Review* 91: 79-98.
- [4] Blume, Lawrence and David Easley (2000), If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. Cornell University Working Paper.

- [5] Bullard, James and James Feigenbaum (2007), A Leisurely Reading of the Life-Cycle Consumption Data. *Journal of Monetary Economics* 54: 2305-2320.
- [6] Carroll, Christopher D. and Miles S. Kimball (2001), Liquidity Constraints and Precautionary Saving. Working Paper.
- [7] Diamond, Peter A. (1965), National Debt in a Neoclassical Growth Model. *American Economic Review* 55: 1126-1150.
- [8] Feigenbaum, James (2008), Information Shocks and Precautionary Saving. *Journal of Economic Dynamics and Control* 32: 3917-3938.
- [9] Findley, T. Scott and Frank N. Caliendo (2007), Does it Pay To Be SMarT? Forthcoming in *Journal of Pension Economics and Finance*.
- [10] Friedman, Milton, (1953), *Essays in Positive Economics* (University of Chicago Press: Chicago).
- [11] Geanakoplos, John, (2008), "Overlapping Generations Models of General Equilibrium," Cowles Foundation Discussion Paper 1663.
- [12] Heifetz, Aviad, Chris Shannon, and Yossi Spiegel (2007), "What to Maximize if You Must," *Journal of Economic Theory* 133: 31-57.
- [13] Helman, Ruth, Craig Copeland, and Jack VanDerhei (2006), Will More of Us Be Working Forever? EBRI Issue Brief 292.
- [14] Huang, Kevin and Frank N. Caliendo, (2007), Rationalizing Seven Consumption-Saving Puzzles in a Unified Framework. Working Paper. Available online at <http://fcaliendo.com/>
- [15] İmrohoroglu, Ayşe, Selahattin İmrohoroglu, and Douglas H. Joines, (2003), "Time-Inconsistent Preferences and Social Security," *Quarterly Journal of Economics* 118: 745-784.
- [16] Keynes, John Maynard (1965 [1936]), *The General Theory of Employment, Interest, and Money* (Harbinger: New York).
- [17] Lusardi, Annamaria and Peter Tufano (2008), Debt Literacy, Financial Experience, and Overindebtedness. Working Paper.

- [18] McKean, Roland N. (1958), *Efficiency in Government through Systems Analysis: With Emphasis on Water Resources Development* (John Wiley: New York).
- [19] Moore, Matt (2006), Automatic 401(k)s: A Win for Workers. National Center for Policy Analysis, Brief Analysis No. 567.
- [20] Phelps, Edmund S. (1966), *Golden Rules of Economic Growth: Studies of Efficient and Optimal Investment* (W. W. Norton: New York).
- [21] Portney, Paul R. and John P. Weyant ed. (1999), *Discounting and Intergenerational Equity* (Resources for the Future: Washington, DC).
- [22] Prest, A. R. and R. Turvey (1965), Cost-Benefit Analysis: A Survey. *Economic Journal* 75: 683-735.
- [23] Ramsey, Frank P. (1928), A Mathematical Theory of Saving. *Economic Journal* 38: 543-559.
- [24] Skinner, Jonathan (1988), Risky Income, Life Cycle Consumption, and Precautionary Savings. *Journal of Monetary Economics* 22: 237-255.
- [25] Solow, Robert M. (1956), A Contribution to the Theory of Economic Growth. *Quarterly Journal of Economics* 70: 65-94.
- [26] Thaler, Richard H. and Shlomo Benartzi (2004), Save More TomorrowTM: Using Behavioral Economics to Increase Employee Saving. *Journal of Political Economy* 112(1), 164-187.
- [27] Yakoboski, Paul and Jennifer Dickemper (1997), Increased Saving but Little Planning: Results of the 1997 Retirement Confidence Survey. EBRI Issue Brief 191.