

Paternalism and Pseudo-Rationality*

Itzik Fadlon, University of California, San Diego and NBER

David Laibson, Harvard University and NBER

August 17, 2017

Abstract

Resource allocations are jointly determined by the actions of social planners and households. In this paper we highlight the distinction between planner optimization and household optimization. We show that planner optimization is a substitute for household optimization and that this is true even when there are information asymmetries, so that households know more about their preferences than planners. Our analysis illustrates the scope for mis-attribution in economic analysis. Are seemingly optimal allocations caused by optimizing households, or are such allocations caused by planners who paternalistically influence myopic and passive households? We show that widely studied allocative optimality conditions that are implied by household optimization also arise in an economy with a rational planner who uses tools such as default savings and Social Security to influence the choices of non-optimizing households. Many classical optimization conditions do not resolve the question of household optimization. Pseudo-rationality arises when rational planners elicit approximately optimal behavior from non-optimizing households.

*We acknowledge very helpful feedback from Youssef Benzarti, Raj Chetty, David Cutler, Emir Kamenica, and Scott Kominers. We thank Edward Kong for excellent research assistance. We are grateful for funding from the Eric Mindich fund for Research in the Foundations of Human Behavior.

1 Introduction

Resource allocations are jointly determined by many actors, including households, firms, and governments. In this paper, we study the interaction between households and a well-intentioned government, which we will refer to as the social planner.¹ We analyze the case in which (i) the social planner is a (benevolent) rational utilitarian, (ii) households are heterogeneous in their degree of rationality, and (iii) the social planner has some scope to intervene – e.g., default savings and mandatory savings, which mirror institutions like 401(k) auto-enrollment and Social Security.² We show that in equilibrium, planner optimization is a (close) substitute for household optimization. This is true even when there are information asymmetries, so that households know more about their preferences than planners.

Our analysis illustrates a potential mis-attribution in economic analysis. Are seemingly optimal allocations caused by optimizing households, or are such allocations caused by planners who paternalistically influence household behavior? We find that some widely studied classical tests for household optimization, which rely on the Euler equation and its variants, do not identify the causal source of the (average) allocative optimality. Equilibrium properties that are implied by household optimization are also implied by planner optimization. The actions of rational planners may result in approximately optimal allocations, by causing the Euler equation to hold on average in a population of non-optimizing households, including myopic and passive households.

To demonstrate these mechanisms, we study a classic life-cycle model in which agents earn labor income during working life and can save a fraction of their earnings for retirement consumption. The planner has two policy levers: mandatory retirement savings (similar to Social Security or defined benefit pensions), and voluntary retirement savings with a default savings rate (similar to 401(k) accounts).³ We consider an economy with three types of

¹Naturally, many governments are not (even approximately) well-intentioned. They are not trying to maximize the well-being of their citizens. These cases can also be studied using the framework in this paper, but we leave those cases for future work.

²Defined benefit (DB) pensions are another example of a mandatory savings scheme. DB pensions have fallen in popularity in the U.S., though they remain commonplace in other developed economies.

³Later in the paper we show that this restriction to only two policy levers is without loss of generality. This generalization occurs in Section 6.

households: optimizing households, who behave optimally throughout their life-cycle, myopic households, who opt out of the default and consume their *entire* disposable income in each period, and passive households, who accept the planner’s default and consume their residual income. We include myopes – an extreme type – to emphasize that our results hold even when the deck is stacked against social efficiency. Additionally, we allow for agents to have privately observed preference parameters of arbitrary structure. Our planner jointly chooses a default level of savings within the system of voluntary savings accounts and designs a Social Security system in order to maximize total social welfare, taking into account the behavior of optimizing households and the suboptimal behavior of myopic and passive households.

We show that, in equilibrium, average marginal utility before retirement is equal to average marginal utility after retirement for *any* distribution of optimizing, myopic, and passive households. Such marginal utility smoothing applies to all concave utility functions with a general class of taste shocks. Accordingly, Euler equations, which characterize an economy with only optimizing households, also arise in an economy with institutions that are optimally designed by a planner for non-optimizing households. As such, aggregate Euler equations do not differentiate between household optimization and planner optimization.

Marginal utility smoothing is closely related to consumption smoothing, which is among the most common tests for household optimization. The precise form of consumption smoothing depends on the curvature of the utility function and the structure of taste shocks. We show that when average consumption smoothing arises in a world comprised only of optimizing households, it also arises in our economy with an optimizing planner and any distribution of optimizing, myopic, and passive households. Moreover, in some leading cases consumption smoothing does not arise with optimizing households, but does arise with an optimizing planner and non-optimizing households. These results imply that consumption smoothing is a more robust property of the model *without* optimizing households, in contrast to the common view.

Although we show that the average Euler equation and average consumption smoothing are not generally diagnostic of household optimization, we demonstrate that other types of

economic information – both from the cross section and from time series – identify the extent of household optimization. For example, bunching (i.e., a mass point) at the default savings level identifies passive households. In the time series, only passive households change their consumption when the default savings rate changes. Such tests are not confounded by the equilibrium link between household and planner optimization. As such, they provide ways to overcome the mis-attribution problems that we highlight. We provide a range of methods for identifying the mass of optimizing, myopic, and passive households.

Our choice of the particular framework for illustrating our arguments – namely, savings over the life-cycle – is motivated by the seemingly contradictory findings of the recent research on household savings. Some papers find evidence that is consistent with optimal savings while others highlight savings anomalies. See Skinner (2007) and Poterba (2014) for reviews of this literature and Section 7 of our paper for a specific discussion of how our model reconciles the different sets of findings. For example, our model implies that household-level sub-optimization, arising from passivity and myopia, will be partially offset by paternalistic policies, like defaults and Social Security. Hence, the economy that our model describes will simultaneously feature both household-level mistakes and many characteristics of "aggregate" optimality. In equilibrium, the Euler equation is satisfied on average across all households, though it is not satisfied for some, or potentially even all, individual households.

Our paper is also related to the behavioral economics literature on optimal paternalism. In our setting the planner chooses policies that dramatically improve the welfare of myopic and passive agents, while relatively weakly distorting the choices and welfare of rational agents (whose behavior would be optimal under laissez faire policies). The social desirability of policies that disproportionately affect non-rational consumers was first highlighted by Camerer, Issacharoff, Loewenstein, O'Donoghue, and Rabin (2003) in their paper on asymmetric paternalism. Our analysis incorporates defaults, which preserve freedom of choice while still influencing behavior. Sunstein and Thaler (2003) and Thaler and Sunstein (2003, 2008) refer to such choice-preserving nudges as libertarian paternalism. Our analytic framework also includes mandates (like Social Security), which fall outside the domain of libertarian

paternalism. Our paper derives socially optimal paternalistic policies for households with behavioral biases, which follows a line of related papers: Diamond (1977), Kotlikoff, Spivak, and Summers (1982), Feldstein (1985), Gruber and Köszegi (2001), Choi et al. (2003), O’Donoghue and Rabin (2003, 2006), Carroll et al. (2009), Fang and Silverman (2009), Bubb and Pildes (2014), Alcott and Taubinsky (2015), Bubb, Corrigan and Warren (2015), Chetty (2015), Moser and de Souza e Silva (2015), Farhi and Gabaix (2015), Bubb and Warren (2016), and Lockwood (2017). Another closely related literature studies self-binding policies, in other words, self-directed paternalism: e.g., Ashraf, Karlan and Yin (2006), Gine, Karlan, and Zinman (2010), Bryan, Karlan, and Nelson (2010), Amador, Werning, and Angeletos (2006), Augenblick, Niederle, and Sprenger (2015), Kaur, Kremer and Mullainathan (2015), and Schilbach (2015).

The remainder of the paper is organized as follows. In Section 2, we describe our baseline model in which the population is comprised of three types of households: optimizers, myopes, and passives. We describe the general (privately observed) taste shocks that these households experience. We also describe the policy levers available to the government. Finally, we characterize the equilibrium behavior of the households (holding fixed the behavior of the social planner). Section 3 derives the equilibrium behavior of the social planner. In this section we show that the classical aggregate Euler equation holds, regardless of the proportions of optimizers, myopes, and passives (as long as the social planner is a rational utilitarian). In Section 4, we study several special cases – multiplicative taste shocks, quadratic utility, and constant relative risk aversion. Using these special cases and our earlier Euler equation results, we derive equilibrium consumption dynamics, which are related to but distinct from the dynamics of marginal utility. We show that consumption smoothing is a robust feature of our model. Indeed, the forces that produce consumption smoothing only grow stronger as the fraction of the population that is myopic and passive increases. In Section 5, we discuss the issue of identification – how can the distribution of optimizing, myopic, and passive households be identified? In Section 6, we derive an isomorphism between our analysis and the associated mechanism design problem. Specifically, we show that the equilibrium of our model, with a

restricted policy space using only default saving and mandatory saving, exactly matches the equilibrium that arises when the government’s policy tools are maximally generalized and the problem is treated as a mechanism design problem. Section 7 explains why our model reconciles several tensions in the empirical consumption literature. Our framework explains why a population that is comprised of a mix of optimizing, myopic, and passive households will have relatively smooth average consumption dynamics, although significant deviations from optimality will arise at the level of individual households, including sensitivity to default savings. Section 8 concludes.

2 Baseline Model

Setup. We consider a two-period model ($t = \{1, 2\}$), where period 1 is working life and period 2 is retirement. Real output during working life is $y_1 = y$ and real output during retirement is 0. Real consumption is denoted c_1 and c_2 . We assume that the real interest rate, r , is fixed and let $R = 1 + r$. The life-time budget constraint household i faces is

$$c_1 + \frac{c_2}{R} \leq y.$$

Preferences. Consider a household with consumption $\{c_1, c_2\}$ and a general taste shifter, vector θ . Total life-time utility is given by

$$U(c_1, c_2; \theta) = u_1(c_1; \theta) + \delta u_2(c_2; \theta),$$

where δ is a discount factor.⁴ The taste parameter $\theta \in \Theta$ varies across households and is independently drawn from a common cumulative distribution function $F(\theta)$ over Θ with a probability density function $f(\theta)$.⁵ We assume that θ is known to the households at time 0,

⁴The case in which δ is heterogeneous across households is embedded in our model. To illustrate this point, a special case of our model is

$$U(c_1, c_2; \theta) = u(c_1) + \widehat{\delta} u(c_2),$$

where $\widehat{\delta} = \delta\theta$.

⁵This common density assumption for θ can be relaxed. As long as we maintain the assumption that the

but not known to the government. Finally, we assume that $u'_t(\cdot; \theta) > 0$ and $u''_t(\cdot; \theta) < 0$ for all θ and for $t = \{1, 2\}$.

To simplify notation, we suppress the i index, and simply refer to households by their taste shock θ .

Institutions. There are two kinds of institutions: a voluntary savings account and a forced retirement savings account. We assume that the planner sets a default level of savings in the voluntary savings account, $s_D \geq 0$. The households are able to opt-out of this default at zero cost. In addition, the planner sets forced retirement savings (a minimum level of savings from labor income) of $s_F \geq 0$, which is deposited into the forced savings account during working life.⁶

In section 6, we show that the institutional restrictions described in the previous paragraph are made without loss of generality. In other words, even if the planner solves a completely general mechanism design problem (for this economy), the solution is the same as the solution that we characterize for the institutionally restricted model.

Household types. There are three types of households (indexed on a continuum from 0 to 1): *Optimizers*, *Myopes*, and *Passives*. We explain each of these in turn.

Optimizers (notated O) choose the optimal level of consumption in all time periods, taking into account their private information about θ and the institutional constraints that they face. Optimizing households may not be able to achieve their first best allocation if their optimal level of savings is lower than the forced level of savings. Formally, optimizing households choose the life-time consumption path $\{c_1, c_2\}$ that solves

$$\max_{\{c_1, c_2\}} u_1(c_1; \theta) + \delta u_2(c_2; \theta)$$

government holds rational expectations, we can allow households to have preference parameters that are drawn from (ex-ante) heterogeneous distributions. To simplify notation, we adopt the assumption of a common iid CDF, $F(\theta)$.

⁶This is essentially what has been adopted by Australia, Israel, and Singapore, and has similarities to Social Security in the US.

subject to two constraints

$$c_1 + \frac{c_2}{R} \leq y,$$

$$c_1 \leq y - s_F \equiv \bar{c}_1.$$

The first constraint is the budget constraint. The second constraint is the period-one liquidity constraint because s_F is the level of mandatory savings – i.e., the minimum level of savings.

To characterize the equilibrium behavior of optimizers, $\{c_1(\theta), c_2(\theta)\}$, we first consider the unconstrained problem – that is, we focus on optimizing households that are unconstrained during their working life. In this case, the standard Euler equation holds for each household:

$$u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta).$$

By contrast, for the constrained optimizing households $c_1(\theta) = y - s_F = \bar{c}_1$ and $c_2(\theta) = R \times s_F$.

This completes our discussion of optimizers. We now turn to the second and third types of households.

Myopes (notated M) opt out of the default and consume as much as possible in every period. Hence, myopes are constrained only by the forced savings, so that they consume $c_1^M = y - s_F = \bar{c}_1$ and $c_2^M = R \times s_F$, which is the same as the constrained optimizers.

Passives (notated P) accept the default and consume the residual income flow. That is, for them $c_1^P = y - s_F - s_D \equiv c_1^D$ and $c_2^P = R \times (s_F + s_D)$.

The rational planner's problem. The shares of optimizing, myopic, and passive households are μ_O , μ_M , and μ_P , respectively, where $0 \leq \mu_O, \mu_M, \mu_P \leq 1$ and $\mu_O + \mu_M + \mu_P = 1$. We denote the distribution of these "decision" types by $\mu \equiv (\mu_O, \mu_M, \mu_P)$. The utilitarian social planner's objective is to choose the policy tools $\{s_D, s_F\}$ that maximize total utility. Note that any pair of values $\{s_D, s_F\}$ generates equilibrium values for period-one consumption by optimizers, $c_1(\theta)$, myopes, c_1^M , and passives, c_1^P . Because of the household budget constraint, these consumption levels imply $c_2(\theta) = R \times (y - c_1(\theta))$, $c_2^M = R \times (y - c_1^M)$, and $c_2^P =$

$R \times (y - c_1^P)$. Accordingly, the planner chooses $\{s_D, s_F\}$ to maximize

$$\begin{aligned} W \equiv & \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta)] dF(\theta) \\ & + \mu_M \int_{\Theta} [u_1(c_1^M; \theta) + \delta u_2(c_2^M; \theta)] dF(\theta) \\ & + \mu_P \int_{\Theta} [u_1(c_1^P; \theta) + \delta u_2(c_2^P; \theta)] dF(\theta). \end{aligned} \quad (1)$$

We will make frequent use of the expectation operator, $E[\cdot]$, which is the expectation taken over the entire population. Specifically, for any random variable $x(\theta)$, this expectation operator integrates jointly over decision types (optimizers, myopes, and passives) and over taste shocks (θ):

$$E[x(\theta)] \equiv \mu_O \int_{\Theta} x^O(\theta) dF(\theta) + \mu_P \int_{\Theta} x^M(\theta) dF(\theta) + \mu_M \int_{\Theta} x^P(\theta) dF(\theta).$$

3 Equilibrium with Optimal Institutions

We begin by analyzing the basic Euler equation. First, consider the benchmark of an economy in which *all* households are optimizers, so that $u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta)$ for each household. In principle, if an econometrician knew each household's value of θ , then it would be possible to test this equation at the household level. However, θ , which is a taste shock, is not observable. Moreover, in practice most variables are measured with noise, which prevents equations from holding exactly. Accordingly, empirical analysis tends to focus on whether the Euler equation is satisfied on average, that is, whether:

$$E[u'_1(c_1(\theta); \theta)] = \delta R E[u'_2(c_2(\theta); \theta)].$$

Naturally, this holds in the benchmark economy of optimizing households. Our first proposition proves that this last equation *also* holds in our economy, in which optimizing households can represent any fraction of the economy (i.e., $\mu_O \in [0, 1]$).

We now characterize the equilibrium allocation in our economy with a *rational* social planner. We prove that some basic equilibrium features that are commonly attributable to household optimization also appear in our partially passive and partially myopic economy as a result of the planner's intervention.

Proposition 1 *Assume a rational planner. Then for any distribution of optimizing, myopic, and passive households, a classical Euler equation will hold on average in the population:*

$$E [u'_1(c_1; \theta)] = \delta RE [u'_2(c_2; \theta)] . \quad (2)$$

Proposition 1 establishes that the aggregate Euler equation (2) holds for any mass vector μ characterizing the fraction of optimizing, myopic, and passive households. The results in the rest of the paper all have the property that a classical optimality condition holds on average in the population regardless of the fraction of optimizing households.

The proof of this first proposition uses three steps that correspond to the following three lemmas.

Lemma 1 ($c_1^D < \bar{c}_1$) *At the planner's optimum, the default consumption in period 1 is strictly less than maximal consumption in period 1:*

$$c_1^D \equiv y - s_F - s_D < y - s_F \equiv \bar{c}_1.$$

Equivalently, $s_D > 0$.

The social planner uses the sum $s_F + s_D$ to pin down total savings for passives. The social planner recognizes that mandatory savings, s_F , has a negative externality on optimizers, some of whom are constrained by s_F . Note that s_D has no negative externality because it only affects the choices of passives. Accordingly, the social planner sets the default savings rate, s_D , strictly greater than zero, which implies that myopes save less than passives: $s_F < s_F + s_D$. This first lemma ($s_D > 0$) is proved in the appendix.

Our second lemma describes a change in variables for the planner's optimization problem.

Lemma 2 (Change of Variables) *With the institutions we have assumed, the planner's problem is isomorphic to jointly choosing the optimal level of two variables:*

- (i) $c_1^D \equiv y - s_F - s_D$ (default consumption in period 1),
 - (ii) $\bar{c}_1 \equiv y - s_F$ (maximal consumption in period 1, given the level of mandatory savings),
- subject to the constraint that $c_1^D \leq \bar{c}_1$ (implied by the original constraint $s_D \geq 0$).

To see this, let $\Gamma(\bar{c}_1) \subset \Theta$ denote the set of θ values that would induce an optimizer to be strictly constrained if period-one consumption were bounded above by \bar{c}_1 . Then, we can re-write the planner's optimization problem (1) as choosing c_1^D and \bar{c}_1 to maximize

$$W \equiv \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(R(y - c_1(\theta)); \theta)] dF(\theta) \quad (3)$$

$$+ \mu_M \int_{\Theta} [u_1(\bar{c}_1; \theta) + \delta u_2(R(y - \bar{c}_1); \theta)] dF(\theta)$$

$$+ \mu_P \int_{\Theta} [u_1(c_1^D; \theta) + \delta u_2(R(y - c_1^D); \theta)] dF(\theta), \quad (4)$$

subject to the constraint

$$c_1^D \leq \bar{c}_1, \quad (5)$$

where $c_1(\theta)$ solves $u_1'(c_1(\theta); \theta) = \delta R u_2'(R(y - c_1(\theta)); \theta)$ for $\theta \notin \Gamma(\bar{c}_1)$ and $c_1(\theta) = \bar{c}_1$ for $\theta \in \Gamma(\bar{c}_1)$.

Lemma 3 (Euler Equation) *If social welfare is maximized with respect to c_1^D and \bar{c}_1 , then the associated first order conditions imply*

$$E [u_1'(c_1; \theta)] = \delta R E [u_2'(c_2; \theta)].$$

Proof. Lemma 1 establishes that at an optimum $\bar{c}_1 > c_1^D$. In other words, myopes consume *strictly* more than passives at an optimum. Because $\bar{c}_1 > c_1^D$, (local) perturbations of c_1^D do not affect the socially optimal value of \bar{c}_1 , and vice versa. Accordingly, the constraint in Lemma 2 (equation (5)) can be ignored when taking first order conditions. Exploiting

the change of variables in Lemma 2 and the irrelevance of the constraint, we know that at the optimum

$$\frac{\partial W}{\partial c_1^D} = \frac{\partial W}{\partial \bar{c}_1} = 0.$$

Note that perturbations of c_1^D will only affect passives, a property we will exploit in the next paragraph. Likewise, perturbations of \bar{c}_1 will only affect myopes and (some) optimizers, which we will also exploit in the next paragraph.

Recall that $c_1^P = c_1^D$ and $c_2^P = R(y - c_1^D)$. The planner's choice of c_1^D establishes an average Euler equation for passives. Specifically, $\frac{\partial W}{\partial c_1^D} = 0$ implies that

$$\int_{\Theta} u'_1(c_1^P; \theta) dF(\theta) = \delta R \int_{\Theta} u'_2(c_2^P; \theta) dF(\theta) \quad (6)$$

$$\text{where } c_2^P = R \times (y - c_1^P).$$

The first order condition for \bar{c}_1 is

$$\begin{aligned} \int_{\Theta} \left[\mu_O u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} + \mu_M u'_1(c_1^M; \theta) \frac{dc_1^M(\theta)}{d\bar{c}_1} \right] dF(\theta) \\ + \delta \int_{\Theta} \left[\mu_O u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} + \mu_M u'_2(c_2^M; \theta) \frac{dc_2^M}{d\bar{c}_1} \right] dF(\theta) = 0. \quad (7) \end{aligned}$$

For myopes, $\frac{dc_1^M}{d\bar{c}_1} R + \frac{dc_2^M}{d\bar{c}_1} = 0$. This is a direct consequence of the household budget constraint.

Moreover, $\frac{dc_1^M}{d\bar{c}_1} = 1$ so that $\frac{dc_2^M}{d\bar{c}_1} = -R$.

For optimizers $\frac{dc_1(\theta)}{d\bar{c}_1} \in \{0, 1\}$, where $\frac{dc_1(\theta)}{d\bar{c}_1} = 1$ iff $\theta \in \Gamma(\bar{c}_1)$.⁷ This implies that

$$\begin{aligned} \int_{\theta \in \Theta} u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} dF(\theta) &= \int_{\theta \notin \Gamma(\bar{c}_1)} [u'_1(c_1(\theta); \theta) \times 0] dF(\theta) + \int_{\theta \in \Gamma(\bar{c}_1)} [u'_1(\bar{c}_1; \theta) \times 1] dF(\theta) \\ &= \int_{\theta \in \Gamma(\bar{c}_1)} u'_1(c_1(\theta); \theta) dF(\theta). \end{aligned}$$

⁷This follows from the following argument. For unconstrained households whose choice of $c_1(\theta)$ is strictly interior to their choice set, we have $\frac{dc_1(\theta)}{d\bar{c}_1} = 0$. For households who are strictly constrained, we have $\frac{dc_1(\theta)}{d\bar{c}_1} = 1$. Lastly, for households who are weakly constrained, we have $\frac{dc_1(\theta)}{d\bar{c}_1} = 0$.

Likewise, $\frac{dc_2(\theta)}{d\bar{c}_1} \in \{0, -R\}$, where $\frac{dc_2(\theta)}{d\bar{c}_1} = -R$ iff $\theta \in \Gamma(\bar{c}_1)$. This implies that

$$\begin{aligned} \int_{\theta \in \Theta} u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} dF(\theta) &= \int_{\theta \notin \Gamma(\bar{c}_1)} [u'_2(c_2(\theta); \theta) \times 0] dF(\theta) + \int_{\theta \in \Gamma(\bar{c}_1)} [u'_2(\bar{c}_2; \theta) \times (-R)] dF(\theta) \\ &= -R \int_{\theta \in \Gamma(\bar{c}_1)} u'_2(c_2(\theta); \theta) dF(\theta). \end{aligned}$$

Equation (7) therefore reduces to:

$$\begin{aligned} \int_{\theta \in \Gamma(\bar{c}_1)} \mu_O u'_1(c_1(\theta); \theta) dF(\theta) + \int_{\theta \in \Theta} \mu_M u'_1(c_1^M; \theta) dF(\theta) \\ = \delta R \int_{\theta \in \Gamma(\bar{c}_1)} \mu_O u'_2(c_2(\theta); \theta) dF(\theta) + \delta R \int_{\theta \in \Theta} \mu_M u'_2(c_2^M; \theta) dF(\theta). \end{aligned} \quad (8)$$

We also have for all $\theta \notin \Gamma(\bar{c}_1)$:

$$u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta). \quad (9)$$

Combining equations (6), (8), and (9), we have

$$E[u'_1(c_1; \theta)] = \delta R E[u'_2(c_2; \theta)]. \blacksquare$$

Proposition 1 follows immediately by combining Lemmas 1, 2, and 3.

In summary, Proposition 1 implies that a standard classical Euler equation characterizes equilibrium allocations in the economy, regardless of the proportions of optimizers, myopes, and passives. Note that the Euler equation does not hold at the level of the individual household. Rather, it holds in expectation across all households. Some households consume too little (optimizers and myopes who have a taste shifter, θ , that would imply an optimal level of $c_1 > \bar{c}_1$, and passives who have a taste shifter, θ , that would imply an optimal level of $c_1 > c_1^D$). Some households consume too much (myopes who have a taste shifter, θ , that would imply an optimal level of $c_1 < \bar{c}_1$, and passives who have a taste shifter, θ , that would

imply an optimal level of $c_1 < c_1^D$). On *average*, the Euler equation is satisfied.

3.1 When could an econometrician reject the Euler equation?

Proposition 1 establishes that an Euler equation is satisfied on *average* in the economy that we study despite the existence of non-optimizing households. However, as we have explained, an analogous Euler equation will not be satisfied household-by-household (though it will be satisfied for unconstrained optimizers). An econometrician with full information – i.e., an econometrician who knows household preferences, including each household’s taste-shifter, θ – would be able to test the Euler equation household-by-household,

$$\begin{aligned} u'_1(c_1; \theta) &= \delta R u'_2(c_2; \theta) && \text{for } c_1 < \bar{c}_1 \\ u'_1(c_1; \theta) &\geq \delta R u'_2(c_2; \theta) && \text{for } c_1 = \bar{c}_1, \end{aligned}$$

and would find that it is not universally satisfied. More generally, the econometrician will find that the Euler equation holds when she uses the same information partition that the planner has and the econometrician will find that the Euler equation does not hold when she uses a finer information partition than the planner has.

4 Special Cases

We now study the special case in which household life-time utility is given by

$$u_1(c_1; \theta) + \delta u_2(c_2; \theta) = u(c_1) + \delta \theta u(c_2).$$

This is a commonly studied case of multiplicative taste shocks (c.f., Atkeson and Lucas 1992; Amador, Werning, and Angeletos 2006; and Beshears et al. 2015). Without loss of generality and to simplify notation, we assume that $\delta R = 1$. We maintain these assumptions throughout this section.

Using these restrictions, we first study the mean of the ratio of marginal utilities before

and after retirement, $E \left[\frac{u'(c_1)}{u'(c_2)} \right]$. Note that in a fully optimizing economy, for any given value of θ we have

$$\frac{u'(c_1(\theta))}{u'(c_2(\theta))} = \theta.$$

Accordingly, averaging across all households yields

$$E \left[\frac{u'(c_1)}{u'(c_2)} \right] = E[\theta].$$

When $E[\theta] = 1$ (a natural benchmark), marginal utility is smoothed such that the mean ratio of marginal utilities equals one, i.e., $E \left[\frac{u'(c_1)}{u'(c_2)} \right] = 1$.

In the next proposition we prove that these relationships hold in our economy as well.

Proposition 2 (Multiplicative Taste Shocks) *Assume a rational planner. Then for any distribution of optimizing, myopic, and passive households, a classical Euler equation ratio will hold on average in the population:*

$$E \left[\frac{u'(c_1)}{u'(c_2)} \right] = E[\theta].$$

Proof. Equation (6) implies that $\int_{\Theta} \left(-\frac{u'(c_1^P)}{u'(c_2^P)} + \theta \right) dF(\theta) = 0$. Since $c_2(\theta) = R(y - \bar{c}_1)$ for $\theta \in \Gamma(\bar{c}_1)$ and $c_2^M = R(y - \bar{c}_1)$, we can also re-write equation (8) as $\int_{\theta \in \Gamma(\bar{c}_1)} \mu_O \left(-\frac{u'(c_1(\theta))}{u'(c_2(\theta))} + \theta \right) dF(\theta) + \int_{\Theta} \mu_M \left(-\frac{u'(c_1^M)}{u'(c_2^M)} + \theta \right) dF(\theta) = 0$. Lastly, for $\theta \notin \Gamma(\bar{c}_1)$, $u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta)$ implies that $-\frac{u'(c_1(\theta))}{u'(c_2(\theta))} + \theta = 0$ and hence $\int_{\theta \notin \Gamma(\bar{c}_1)} \mu_O \left(-\frac{u'(c_1(\theta))}{u'(c_2(\theta))} + \theta \right) dF(\theta) = 0$. Put together, we have that

$$\begin{aligned} E \left[\frac{u'(c_1)}{u'(c_2)} \right] &= \mu_O \left[\int_{\theta \in \Gamma(\bar{c}_1)} \frac{u'(c_1(\theta))}{u'(c_2(\theta))} dF(\theta) + \int_{\theta \notin \Gamma(\bar{c}_1)} \frac{u'(c_1(\theta))}{u'(c_2(\theta))} dF(\theta) \right] \\ &\quad + \mu_P \int_{\Theta} \frac{u'(c_1^P)}{u'(c_2^P)} dF(\theta) + \mu_M \int_{\Theta} \frac{u'(c_1^M)}{u'(c_2^M)} dF(\theta) = E[\theta]. \end{aligned}$$

This completes the proof. ■

4.1 Consumption Smoothing

The degree of consumption smoothing is widely used as an indicator of household optimization. The mapping from marginal-utility smoothing to consumption smoothing depends on the curvature of the utility function. In this subsection, we work out this mapping for quadratic utility and constant relative risk aversion.

For quadratic utility we get exact results. When there are no taste shocks ($\theta = 1$ for all households), we show that $E[c_2 - c_1] = 0$ *regardless* of the distribution of optimizing, myopic, and passive households. On the other hand, when there are household-level taste shocks but no taste shocks on average ($E[\theta] = 1$), we show that $E[c_2 - c_1] = 0$ *only if* there are *no* optimizing households. Surprisingly, the presence of optimizing households causes average consumption to fall at retirement: $E[c_2 - c_1] < 0$.

For the case of constant relative risk aversion, we get analogous results studying the growth rate of consumption at retirement (and using a first-order approximation).

Together, these results imply that *consumption* smoothing is a more robust property of the model *without* optimizing households. In other words, there are leading preference and taste-shock cases in which consumption smoothing arises without optimizing households but does not arise with optimizing households (and not vice versa).

We now introduce our notation. We study the cases of: (i) quadratic utility, i.e., $u(c) = c - \frac{b}{2}c^2$, where $b > 0$, and (ii) constant relative risk aversion, i.e., $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, where $\gamma > 0$.⁸ For each case, we analyze consumption moments under two assumptions: first, when $\theta = 1$ for every household as implicitly assumed in most papers, and, second, when there are no taste shocks on average at retirement, i.e., when $E[\theta] = 1$. The results are summarized in the following lemma.

Lemma 4 (Parametric Utility Functions) *Assume any distribution of optimizing, myopic, and passive households.*

⁸As $\gamma \rightarrow 1$, this function converges to $\ln(c)$.

(A) Quadratic utility implies:

1. With no taste shocks at retirement ($\theta = 1$ for all households),

$$E[c_2 - c_1] = 0;$$

2. With no taste shocks on average at retirement ($E[\theta] = 1$),

$$E[c_2 - c_1] = -cov(\theta, c_2) \leq 0.$$

(B) Constant relative risk aversion utility implies (to a first-order approximation):

1. With no taste shocks at retirement ($\theta = 1$ for all households),

$$E\left[\frac{c_2 - c_1}{c_1}\right] \cong 0;$$

2. With no taste shocks on average at retirement ($E[\theta] = 1$),

$$E\left[\frac{c_2 - c_1}{c_1}\right] \cong -cov\left(\theta, \frac{c_2}{c_1}\right) \leq 0.$$

Proof. (A) Plugging $u(c) = c - \frac{b}{2}c^2$ into Proposition 1 implies that $E[c_1 - \theta c_2] = E[1 - \theta]/b$. Assuming that $\theta = 1$ for every household yields $E[c_2 - c_1] = 0$, and assuming that $E[\theta] = 1$ yields $E[c_2 - c_1] = -cov(\theta, c_2)$.

(B) Similar to the proof of Proposition 2, one can show that $E\left[\frac{\theta u'(c_2)}{u'(c_1)}\right] = 1$. This is because equation (6) implies that $\int_{\Theta} \left(-1 + \frac{\theta u'(c_2^F)}{u'(c_1^F)}\right) dF(\theta_i) = 0$, equation (8) combined

with $c_1(\theta) = \bar{c}_1$ for $\theta \in \Gamma(\bar{c}_1)$ and $c_1^M = \bar{c}_1$ implies that $\int_{\theta \in \Gamma(\bar{c}_1)} \mu_O \left(-1 + \frac{\theta u'(c_2(\theta))}{u'(c_1(\theta))}\right) dF(\theta) +$

$\int_{\Theta} \mu_M \left(-1 + \frac{\theta u'(c_2^M)}{u'(c_1^M)}\right) dF(\theta) = 0$, and $u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta)$ for $\theta \notin \Gamma(\bar{c}_1)$ implies

that $\frac{\theta u'(c_2(\theta))}{u'(c_1(\theta))} = 1$ and hence $\int_{\theta \notin \Gamma(\bar{c}_1)} \mu_O \left(-1 + \frac{\theta u'(c_2(\theta))}{u'(c_1(\theta))}\right) dF(\theta) = 0$, which together yield

$E\left[\frac{\theta u'(c_2)}{u'(c_1)}\right] = 1$. Linearization of $u'(c_2)$ around $u'(c_1)$ in the latter equality results in $E\left[-\theta \gamma \frac{c_2 - c_1}{c_1}\right] \cong$

$1 - E[\theta]$. Assuming that $\theta = 1$ for every household yields $E\left[\frac{c_2 - c_1}{c_1}\right] \cong 0$, and assuming that $E[\theta] = 1$ yields $E\left[\frac{c_2 - c_1}{c_1}\right] \cong -cov\left(\theta, \frac{c_2}{c_1}\right)$. ■

Lemma 4 demonstrates that when there are no taste shocks, our economy is characterized by mean consumption smoothing for any distribution of optimizing, myopic, and passive households. Therefore, in this commonly studied case, consumption smoothing is not a diagnostic test for household optimization. In addition, with no *average* taste shocks, the covariance expressions in Lemma 4 are weakly positive and proportional to the mass of optimizing households. This introduces some counter-intuitive properties. First, the average change in consumption (entering retirement) will be exactly zero only when the fraction of optimizing households is zero (since $cov(\theta, c_2) = cov\left(\theta, \frac{c_2}{c_1}\right) = 0$ if and only if $\mu_O = 0$). Second, with a positive share of optimizing households, the consumption drop at retirement *increases* (since $\mu_O > 0$ implies that $cov(\theta, c_2) > 0$ and $cov\left(\theta, \frac{c_2}{c_1}\right) > 0$). That is, not only are consumption smoothing tests not conclusive for analyzing household optimization, but in the commonly studied case of $E[\theta] = 1$, average consumption smoothing is actually inconsistent with the null hypothesis of household optimization.

5 Identification of Optimizers, Myopes, and Passives

While the average Euler equation and average consumption smoothing cannot generally reveal the distribution of household types, tests that rely on other moments of households' economic behavior do reveal the mass of optimizing and non-optimizing households.

Consider the cross-sectional distribution of savings. We characterize the distribution of overall savings assuming that the distribution of taste shocks, $F(\theta)$, is continuous.

The first type of test for household optimization relies on the cross section. One indicator is bunching (excess mass) around the combined levels of forced and default savings, $s_F + s_D$. A discrete jump in the savings' cumulative distribution function at this point identifies the mass of passive households.

Another type of test employs quasi-experiments and relies on behavioral responses to

policy variations,⁹ which can be measured by changes in the distribution of savings (or by household-level elasticities). First, if there are passive households, a change in the default level of savings from s_D to s'_D will engender (new) bunching around $s_F + s'_D$. A different measure for testing the same hypothesis is the change in average (or overall) savings. If all households are optimizers, then a change in defaults will keep the distribution of savings, and hence average savings, unchanged. Any change in average savings would reject the hypothesis of perfect household optimization, and in our world, would reveal the share of passive households.

Second, changes to the level of forced savings can shed light on the presence and fraction of myopic households and of optimizing households. To see this, consider a change in the level of mandatory savings and evaluate the change in the mass of households that continue to save no more than the mandatory savings level, s_F :

$$\frac{d\Pr(s = s_F)}{ds_F} = \frac{d}{ds_F} [\mu_O \times \Pr(\theta \in \Gamma(\bar{c}_1)) + \mu_M] = \mu_O \times \frac{d\Pr(\theta \in \Gamma(\bar{c}_1))}{ds_F}.$$

This *change* in the mass of households at the forced savings level is positive if and only if there are optimizing households on the boundary. Indeed, if the observer knows the utility function and the distribution of θ , $F(\theta)$, then the derivative $\frac{d\Pr(\theta \in \Gamma(\bar{c}_1))}{ds_F}$ can be calculated. Hence, observing $\frac{d\Pr(s=s_F)}{ds_F}$ enables the econometrician to calculate the mass of optimizing agents, μ_0 . Once μ_0 is known, the econometrician can also calculate the mass of myopes:

$$\mu_M = \Pr(s = s_F) - \mu_O \times \Pr(\theta \in \Gamma(\bar{c}_1)).$$

A related global test can be developed by studying average savings. Consider a decrease in the forced savings level. If this engenders a one-for-one decline in average economy-wide savings, then all households are either myopic or passive. However, if the decline is less than one-for-one, then there is a strictly positive measure of optimizing households: parametric assumptions on the utility function (and the distribution of taste shifters θ) would be needed

⁹For instance, policy changes might arise as the government learns about the efficacy of new policies (e.g., the research that led to passage of the Pension Protection Act of 2006).

to impute the population mass of optimizers.

The tests that we highlight here are derived from our positive model. However, the broader point is that an analysis of cross-sectional distributions of economic outcomes and how they change in response to variation in the economic environment can provide tests for household optimization that are not confounded by the presence of a rational social planner.

6 Generalization as a Mechanism Design Problem

In this section, we show that the equilibrium of our model above (which has a restricted policy space) exactly matches the equilibrium that arises when the government's policy tools are maximally generalized and the problem is treated as a mechanism design problem. To recap, the problem posed in Sections 2 and 3 is to choose the two policy variables s_F (mandatory savings) and s_D (additional default savings) to maximize the social planner's objective

$$\begin{aligned}
W(c_1(\theta), c_2(\theta), c_1^D, c_2^D) &= \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta)] dF(\theta) \\
&\quad + \mu_M \int_{\Theta} [u_1(c_1^M(\theta); \theta) + \delta u_2(c_2^M(\theta); \theta)] dF(\theta) \\
&\quad + \mu_P \int_{\Theta} [u_1(c_1^D(\theta); \theta) + \delta u_2(c_2^D(\theta); \theta)] dF(\theta). \tag{10}
\end{aligned}$$

We can also study the generalized version of this problem using a mechanism design framework. Now the planner chooses $\{c_1(\theta)\}_{\theta \in \Theta}$, $\{c_2(\theta)\}_{\theta \in \Theta}$, c_1^D , and c_2^D to maximize equation (10) subject to the within-household budget constraints,

$$c_1(\theta) + \frac{c_2(\theta)}{R} \leq y \text{ for all } \theta, \tag{11}$$

$$c_1^M + \frac{c_2^M}{R} \leq y, \tag{12}$$

$$c_1^D + \frac{c_2^D}{R} \leq y, \tag{13}$$

incentive compatibility constraints for optimizers,

$$u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta) \geq u_1(c_1(\theta'); \theta) + \delta u_2(c_2(\theta'); \theta), \quad \forall \theta, \theta' \quad (14)$$

$$u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta) \geq u_1(c_1^D; \theta) + \delta u_2(c_2^D; \theta), \quad \forall \theta \quad (15)$$

and a maximally impatient reporting rule for myopes,

$$c_1^M = \max \left\{ \sup_{\theta \in \Theta} c_1(\theta), c_1^D \right\}. \quad (16)$$

We now summarize the mechanism design problem: the planner chooses $\{c_1(\theta)\}_{\theta \in \Theta}$, $\{c_2(\theta)\}_{\theta \in \Theta}$, c_1^D , and c_2^D to maximize the objective in equation (10), subject to the budget constraints and incentive compatibility constraints in equations (11)-(16). In this mechanism design problem, optimizers report their type (truthfully in equilibrium), myopes always report the type that has the highest immediate consumption in the mechanism, and passives follow the defaults c_1^D and c_2^D .

Proposition 3 (Characterization of the Mechanism Design Problem) *The mechanism design problem and the constrained problem (i.e., maximizing equation (10) by choosing arguments s_F and s_D) generate the same equilibrium allocation.*

This implies an immediate corollary.

Corollary 1 *The mechanism design problem and the constrained problem (maximizing equation (10) by choosing s_F and s_D) generate the same aggregate Euler equation:*

$$E [u'_1(c_1; \theta)] = \delta RE [u'_2(c_2; \theta)].$$

Accordingly, the institutional assumptions that are made in Section 2 are made without loss of generality.

Proof (Characterization of the Mechanism Design Problem). We first show that the (second-best) optimal allocation is characterized by a “maximum consumption rule,”

which we now define. For every optimizer, θ , there exists a (full-information) unconstrained optimum, $(c_1^*(\theta), c_2^*(\theta))$, which satisfies budget balance

$$c_2^*(\theta) = R(y - c_1^*(\theta)),$$

and the first order condition

$$u'(c_1^*(\theta); \theta) = \delta R u'(c_2^*(\theta); \theta).$$

We now describe an allocation that is characterized by a maximum consumption rule. Consider an allocation that assigns consumption $(c_1(\theta), c_2(\theta))$, for all $\theta \in \Theta$, and consumption for the passives: c_1^D and c_2^D . Let $\bar{c}_1 = \max \{ \sup_{\theta \in \Theta} c_1(\theta), c_1^D \}$. The allocation satisfies a maximum consumption rule if and only if two conditions are satisfied: (i) every type with $c_1^*(\theta) \leq \bar{c}_1$ obtains $c_1(\theta) = c_1^*(\theta)$ and $c_2(\theta) = c_2^*(\theta)$ in the allocation; and (ii) every type with $c_1^*(\theta) > \bar{c}_1$ obtains $c_1(\theta) = \bar{c}_1$ and $c_2(\theta) = R(y - \bar{c}_1(\theta))$ in the allocation.

To prove that an optimal mechanism generates an allocation that is a maximum consumption rule, consider a candidate allocation given by $\{\hat{c}_1(\theta)\}_{\theta \in \Theta}$, $\{\hat{c}_2(\theta)\}_{\theta \in \Theta}$, \hat{c}_1^D , and \hat{c}_2^D that is incentive compatible (i.e., satisfying equations (11)-(16)), and does *not* satisfy the maximum consumption rule. Then perturb $\{\hat{c}_1(\theta)\}_{\theta \in \Theta}$, $\{\hat{c}_2(\theta)\}_{\theta \in \Theta}$, \hat{c}_1^D , and \hat{c}_2^D in the following way. Let $\bar{c}_1 = \max \{ \sup_{\theta \in \Theta} \hat{c}_1(\theta), \hat{c}_1^D \}$ and construct a new allocation that is a maximum consumption rule such that (i) every type with $c_1^*(\theta) \leq \bar{c}_1$ obtains $c_1(\theta) = c_1^*(\theta)$ and $c_2(\theta) = c_2^*(\theta)$ in the mechanism; and (ii) every type with $c_1^*(\theta) > \bar{c}_1$ achieves $c_1(\theta) = \bar{c}_1$ and $c_2(\theta) = R(y - \bar{c}_1(\theta))$ in the mechanism. This new allocation is incentive compatible: households either achieve their first-best allocation or they obtain $(\bar{c}_1, R(y - \bar{c}_1))$, which is their most preferred consumption pair in the set of all offered consumption pairs (because of the concavity of u). The new allocation improves welfare weakly for every agent compared to their welfare in the candidate allocation: agents that achieve their first-best allocation in the new allocation obtain a weak improvement in welfare (because they are now at their unconstrained optimum), and agents that obtain the maximum level of consumption obtain a weak improvement in welfare

(because they previously had $\widehat{c}_1(\theta) \leq \bar{c}_1$ and they now also have an allocation that uses their entire endowment). Because the original allocation was not a maximum consumption rule, the new allocation generates a strict improvement in welfare for at least one agent.¹⁰ Hence, only maximum consumption rules are solutions to the mechanism design problem. It follows that the planner’s problem can be reduced to the choice of a maximum consumption rule – i.e., the choice of \bar{c}_1 – and the choice of c_1^D . By Lemma 2, the mechanism design problem (choosing $\{c_1(\theta)\}_{\theta \in \Theta}$, $\{c_2(\theta)\}_{\theta \in \Theta}$, c_1^D , and c_2^D to maximize the objective in equation (10) subject to the budget constraints and incentive compatibility constraints in equations (11)–(16)) and the constrained problem (maximizing equation (10) by choosing s_F and s_D) are isomorphic optimization problems. Accordingly, they have the same equilibrium allocation.

■

7 Related Empirical Literature

Our choice of the particular framework for illustrating our ideas – namely, consumption dynamics over the life-cycle – is motivated by the tension between research that finds relatively smooth average lifecycle consumption dynamics (consistent with rational consumer behavior) and research that finds that individual households are highly sensitive to savings defaults (which is inconsistent with rational consumer behavior).

Analyzing average lifecycle consumption dynamics, many papers find evidence that is consistent with optimal retirement savings. For example, using a structural life-cycle analysis and accounting for government transfers and Social Security benefits, Scholz et al. (2006) estimate that less than twenty percent of households in the Health and Retirement Study under-save for retirement. Moreover, Scholz et al. (2006) find that the wealth deficits are generally a small fraction of life-time wealth. Bernheim, Skinner, and Weinberg (2001) report that expenditure on food drops at retirement, but Aguiar and Hurst (2005) report that the

¹⁰There are two ways for an allocation to fail to be a maximum consumption rule. Either (i) there is a type with $c_1^*(\theta) \leq \bar{c}_1$, and that type fails to obtain $c_1(\theta) = c_1^*(\theta)$ and $c_2(\theta) = c_2^*(\theta)$ in the mechanism; or (ii) there is a type with $c_1^*(\theta) > \bar{c}_1$, and that type fails to obtain $c_1(\theta) = \bar{c}_1$ and $c_2(\theta) = R(y - \bar{c}_1(\theta))$ in the mechanism. If case (i) applies, the type is made strictly better off by construction. If case (ii) applies, the type is made strictly better off because u is strictly concave.

drop in food expenditure at retirement is illusory, in the sense that caloric consumption and other measures of meal quality do not change when households enter retirement.¹¹ Studying overall nondurable expenditure, Aguila et al. (2011) do not find a statistically significant drop in consumption at retirement.

Another body of research studies employer retirement savings plans, and shows that automatic enrollment and other institutional nudges have a large effect on household savings. Many households do not deviate from the default contribution rate as well as the default asset allocation, and are largely unresponsive to government- or employer-subsidies for retirement savings contributions (Madrian and Shea 2001, Choi et al. 2004, Beshears et al. 2009, and Chetty et al. 2014).

Our model is consistent with these seemingly contradictory sets of findings. Our framework explains why a population that is comprised of a mix of optimizing, passive, and myopic households will have relatively smooth average consumption dynamics around retirement and will satisfy an aggregate Euler equation. We show that a benevolent government will set policy that elicits this aggregate (second-best) efficiency property, although significant deviations from optimality will arise at the level of individual households, including sensitivity to default savings.

8 Conclusion

We study a simple setting that illustrates the interactions between optimizing social planners and heterogeneous households, some of whom are optimizers, some of whom are passive, and some of whom are myopic. In this setting, planner optimization is a partial substitute for household optimization. This substitution arises because the social planner has the ability to design institutions – e.g., default savings and Social Security – that influence the consumption profiles of households. In equilibrium, classical Euler equations hold on average in the cross-section of households (but not for each household). These Euler equation properties arise

¹¹However, in a recent working paper, Stephens and Toohey (2016) revisit the analysis in Aguiar and Hurst with a broader set of data sources and methodologies and report that caloric consumption does decline on average around the time of retirement.

generally, whether or not households are optimizers.

These results imply that Euler equation tests and related consumption-smoothing tests – e.g., the lack of an average drop in consumption at retirement – do not differentiate between an optimizing social planner and optimizing households. However, even in the economy that we have studied, planner optimization is distinguishable, in principle, from household optimization. Under the assumption of household optimization, the Euler equation will hold for *each* household, but, with planner optimization (and without universal household optimization), it will only hold on *average* in the cross section. Although household-by-household tests are theoretically determinative, such fine-grained analysis is practically problematic if data are measured with noise or if some variables are unobservable (e.g., household-level taste shocks). However, other tests do distinguish between planner optimization and household optimization. Exogenous changes in policy (e.g., a default change at the level of a firm, or a natural experiment in the Social Security system), reveal more about household rationality than averages of observational data in the cross section.

Our conclusions depend upon the assumption that the government is a fully rational utilitarian. It is likely that flesh-and-blood governments fall short of this benevolent benchmark, despite (or because of) the pressures that they face to get re-elected. This leads to a natural follow-up question: how would our results change if the government is not utilitarian, but is instead a self-interested political party? The answer depends on two key considerations: what is the voting frequency of different types of households and to what degree do altruistic motives influence voting? These are important extensions, which we leave for future research.

References

- [1] Aguiar, Mark, and Erik Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5): pp. 919-948.
- [2] Aguila, Emma, Orazio Attanasio, and Costas Meghir (2011): "Changes in Consumption at Retirement: Evidence from Panel Data," *Review of Economics and Statistics*, 93(3): pp. 1094-1099.
- [3] Allcott, Hunt, and Dmitry Taubinsky (2015): "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 105(8): pp. 2501-2538.
- [4] Amador, Manuel, Iván Werning, and George-Marios Angeletos (2006): "Commitment vs. Flexibility," *Econometrica*, 74(2): pp. 365-396.
- [5] Ashraf, Nava, Dean Karlan, and Wesley Yin (2006): "Tying Odysseus to The Mast: Evidence from a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, 121(2): pp. 635-672.
- [6] Atkeson, Andrew, and Robert E. Lucas (1992): "On Efficient Distribution with Private Information," *Review of Economic Studies*, 59(3): pp. 427-453.
- [7] Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015): "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 130(3): pp. 1067-1115.
- [8] Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg (2001): "What Accounts for the Variation in Retirement Wealth among US Households?" *American Economic Review*, pp. 832-857.
- [9] Beshears, John, James J. Choi, Christopher Harris, David Laibson, Brigitte C. Madrian, and Jung Sakong (2015): "Self Control and Commitment: Can Decreasing the Liquidity

of a Savings Account Increase Deposits?," Working Paper 21474, National Bureau of Economic Research.

- [10] Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2009): "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States," in Jeffrey Brown, Jeffrey Liebman, and David A. Wise, eds. *Social Security Policy in a Changing Environment*, pp. 167-195. University of Chicago Press.
- [11] Bryan, Gharad, Dean Karlan, and Scott Nelson (2010): "Commitment Devices," *Annual Review of Economics*, 2(1): pp. 671-698.
- [12] Bubb, Ryan, Patrick Corrigan, and Patrick L. Warren (2015): "A Behavioral Contract Theory Perspective on Retirement Savings." *Connecticut Law Review*, Commentary Issue of Volume 47.
- [13] Bubb, Ryan, and Richard Pildes (2014): "How Behavioral Economics Trims Its Sails and Why," *Harvard Law Review*, Vol 127.
- [14] Bubb, Ryan, and Patrick L. Warren (2016): "An Equilibrium Theory of Retirement Plan Design." NYU Law and Economics Research Paper No. 16-28.
- [15] Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin (2003): "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'," *University of Pennsylvania Law Review* 151(3): pp. 1211-1254.
- [16] Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte Madrian, and Andrew Metrick (2009): "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics* 124(4): pp. 1639-1674.
- [17] Chetty, Raj (2015): "Behavioral Economics and Public Policy: A Pragmatic Perspective," *American Economic Review* 105(5): pp. 1-33

- [18] Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen (2014): "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 129(3): pp. 1141-1219.
- [19] Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick (2003): "Optimal Defaults." *American Economic Review* 93(2): pp. 180-185
- [20] Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick (2004): "For Better or for Worse: Default Effects and 401 (k) Savings Behavior," in David A. Wise, ed. *Perspectives on the Economics of Aging*, pp. 81-126. University of Chicago Press.
- [21] Diamond, Peter A. (1977): "A Framework for Social Security Analysis," *Journal of Public Economics*, 8(3): pp. 275-298.
- [22] Fang, Hanming, and Dan Silverman (2009): "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY," *International Economic Review* 50(4): pp. 1043-1077.
- [23] Farhi, Emmanuel, and Xavier Gabaix (2015): "Optimal Taxation with Behavioral Agents," Working Paper No. w21524, National Bureau of Economic Research.
- [24] Feldstein, M. (1985): "The Optimal Level of Social Security Benefits," *Quarterly Journal of Economics*, 100(2): pp. 303-320.
- [25] Giné, Xavier, Dean Karlan, and Jonathan Zinman (2010): "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation," *American Economic Journal: Applied Economics* 2(4): pp. 213-235.
- [26] Gruber, Jonathan, and Botond Köszegi (2001): "Is Addiction 'Rational'? Theory and Evidence," *Quarterly Journal of Economics*, 116(4): pp. 1261-1303.
- [27] Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan (2015): "Self-Control at Work," *Journal of Political Economy* 123(6): pp. 1227-1277.

- [28] Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers (1982): “The Adequacy of Savings,” *American Economic Review*, 2(5): pp. 1056–1069.
- [29] Lockwood, Benjamin (2016): “Optimal Income Taxation with Present Bias,” Working paper.
- [30] Madrian, Brigitte C., and Dennis F. Shea (2001): “The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior,” *Quarterly Journal of Economics*, 116(4): pp. 1149-1187.
- [31] Moser, Christian, and Pedro Olea de Souza e Silva (2015): “Paternalism vs Redistribution: Designing Retirement Savings Policies with Behavioral Agents,” Working paper Columbia University.
- [32] O’Donoghue, Ted, and Matthew Rabin (2003): "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes," *American Economic Review*, 93(2): pp. 186-191.
- [33] O’Donoghue, Ted, and Matthew Rabin (2006): “Optimal Sin Taxes,” *Journal of Public Economics* 90(10): pp. 1825-1849.
- [34] Poterba, James M. (2014): "Retirement Security in an Aging Population," *American Economic Review* 104(5): pp. 1-30.
- [35] Schilbach, Frank (2015): “Alcohol and Self-Control: A Field Experiment in India,” Working paper.
- [36] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006): “Are Americans Saving ‘Optimally’ for Retirement?,” *Journal of Political Economy*, 114(4): pp. 607-643.
- [37] Skinner, Jonathan (2007): “Are You Sure You’re Saving Enough for Retirement?” *Journal of Economic Perspectives* 21(3): pp. 59-80.
- [38] Sunstein, Cass R., and Richard H. Thaler (2003): “Libertarian Paternalism is not an Oxymoron,” *University of Chicago Law Review*: pp. 1159-1202.

- [39] Thaler, Richard H., and Cass R. Sunstein (2003): "Libertarian Paternalism," *American Economic Review*, 93(2): pp.175-179.
- [40] Thaler, Richard H., and Cass R. Sunstein (2008): "Nudge: Improving Decisions about Health, Wealth, and Happiness" Yale University Press.

9 Appendix

9.1 Proof of Lemma 1 ($c_1^D < \bar{c}_1$; i.e., $s_D > 0$).

First, we break the planner's problem down into two separable sub-problems. In the first sub-problem, we solve for $s_F = y - \bar{c}_1$ and *ignore* the passives. In the second sub-problem, we solve for s_D and *ignore* both the myopes and the optimizers. This separation of the two problems is only admissible if the resulting optima, s_F and s_D , satisfy an incentive compatibility (IC) constraint,

$$c_1^D \leq \bar{c}_1,$$

which can also be expressed as

$$y - s_F - s_D \leq y - s_F.$$

This IC constraint implies that myopes do not have an incentive to pretend to be passives. Note that this IC constraint is the *only* condition that could link the passives to the optimizers and the myopes. At the end of this proof we confirm that this IC constraint is satisfied by the solutions of the *separated* problems. Hence, the separation is without loss of generality. In other words, we can characterize optimal policy for the passives – the *sum* $s_F + s_D$ – without taking account of the policy parameter that influences the myopes and optimizers, s_F . Likewise, we can characterize optimal policy for the myopes and optimizers, without taking account of optimal policy for the passives.

In the separated problem, the optimal level of \bar{c}_1 is given by the following Euler equation:

$$\int_{\Theta} \left[\mu_O u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} + \mu_M u'_1(c_1^M; \theta) \frac{dc_1^M(\theta)}{d\bar{c}_1} \right] dF(\theta) \\ + \delta \int_{\Theta} \left[\mu_O u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} + \mu_M u'_2(c_2^M; \theta) \frac{dc_2^M}{d\bar{c}_1} \right] dF(\theta) = 0.$$

We can rearrange this by grouping together the optimizer terms and the myope terms:

$$\begin{aligned} \mu_O \int_{\Theta} \left[u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} + \delta u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} \right] dF(\theta) \\ + \mu_M \int_{\Theta} \left[u'_1(c_1^M; \theta) \frac{dc_1^M(\theta)}{d\bar{c}_1} + \delta u'_2(c_2^M; \theta) \frac{dc_2^M}{d\bar{c}_1} \right] dF(\theta) = 0. \end{aligned}$$

Further simplifying, note that $u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} + \delta u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} = 0$ for all $\theta \notin \Gamma(\bar{c}_1)$.

Consequently, we have

$$\begin{aligned} \mu_O \int_{\theta \in \Gamma(\bar{c}_1)} [u'_1(c_1(\theta); \theta) - \delta R u'_2(c_2(\theta); \theta)] dF(\theta) \\ + \mu_M \int_{\Theta} [u'_1(c_1^M; \theta) - \delta R u'_2(c_2^M; \theta)] dF(\theta) = 0. \end{aligned}$$

Note that $u'_1(c_1(\theta); \theta) - \delta R u'_2(c_2(\theta); \theta) > 0$ for all $\theta \in \Gamma(\bar{c}_1)$, where $\Gamma(\bar{c}_1) \subset \Theta$ denotes the set of θ values that would induce an optimizer to be strictly constrained if period-one consumption were bounded above by \bar{c}_1 . Therefore,

$$\int_{\theta \in \Gamma(\bar{c}_1)} [u'_1(c_1(\theta); \theta) - \delta R u'_2(c_2(\theta); \theta)] dF(\theta) > 0,$$

and accordingly

$$\int_{\Theta} [u'_1(c_1^M; \theta) - \delta R u'_2(c_2^M; \theta)] dF(\theta) < 0.$$

At $s_D = 0$,

$$\begin{aligned}
\frac{dW}{ds_D} &= - \int_{\Theta} [u'_1(c_1^P; \theta) - \delta R u'_2(c_2^P; \theta)] dF(\theta) \\
&= - \int_{\Theta} [u'_1(y - s_F; \theta) - \delta R u'_2(R s_F; \theta)] dF(\theta) \\
&= - \int_{\Theta} [u'_1(c_1^M; \theta) - \delta R u'_2(c_2^M; \theta)] dF(\theta) \\
&> 0.
\end{aligned}$$

This shows that $\frac{dW}{ds_D} > 0$ at $s_D = 0$. Because the optimization with respect to s_D is globally concave,¹² it follows that $s_D > 0$. This proves the lemma *and* confirms that the IC constraint is satisfied by the solutions of the separated problems. ■

¹²The second derivative of the objective is

$$\int_{\Theta} [u''_1(y - s_F - s_D; \theta) + \delta R^2 u''_2(R(s_F + s_D); \theta)] dF(\theta) < 0,$$

implying that the objective is concave.