

Paternalism and Pseudo-Rationality: An Illustration Based on Retirement Savings*

Itzik Fadlon, University of California, San Diego and NBER

David Laibson, Harvard University and NBER

October 5, 2022

Abstract

Resource allocations are jointly determined by the actions of social planners and households. We study economies in which households have private information about their tastes and have a distribution of behavioral propensities: optimal, myopic, or passive. In such economies, we show that utilitarian planners enact policies such as Social Security and default savings that cause equilibrium consumption smoothing (on average in the cross-section of households). Our framework resolves tensions in the household savings literature by simultaneously explaining evidence of consumption smoothing and optimal savings on average across households while also predicting behavioral anomalies at the household level. In general, common tests of consumption smoothing are not sufficient to show that households are optimizers, but natural experiments can be used to identify the fractions of optimizing, myopic, and passive households.

Keywords: paternalism, consumption, savings, social planner, optimization, myopia, passivity.

*We acknowledge extremely helpful feedback from our editor, Johannes Spinnewijn, and three referees. We also thank Youssef Benzarti, Raj Chetty, David Cutler, Emir Kamenica, Scott Kominers, Benjamin Lockwood, and participants at the ASSA and National Tax Association annual meetings. 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. Corresponding author: David Laibson, dlaibson@harvard.edu, Department of Economics, Harvard University, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138.

1 Introduction

Resource allocations are jointly determined by many actors, including households, firms, and governments. Consequently, allocative efficiency is influenced by the complex interplay among these intertwined economic actors. For example, if households save optimally for retirement, it may not be clear what causal role other institutions played, including paternalistic savings policies like defaults, defined benefit pensions, and Social Security. When we observe an economy with many paternalistic policies, it is hard to know what counterfactual allocation would have occurred if those paternalistic policies were removed. This attribution problem is closely related to seemingly contradictory findings in the household savings literature, including average levels of total saving (mandatory and voluntary) that are approximately optimal and household-level data that exhibits many behavioral anomalies like excess sensitivity and default effects (see Skinner 2007, Poterba 2014, and Beshears et al. 2018 for reviews). Could consumption smoothing arise, in part, because passive and myopic households are partially hemmed in by paternalistic policies?

To answer this question, we study the interaction between households and a benevolent government, which we will refer to as the social planner.¹ We analyze the case in which (i) the social planner is a rational utilitarian, (ii) households are heterogeneous in their degree of rationality and in their taste shocks, both of which are private information, and (iii) the social planner has some scope to influence the decisions of households – e.g., default savings and mandatory savings, which mirror institutions like 401(k) auto-enrollment and Social Security.² Without any information asymmetries, planner optimization would be a perfect substitute for household optimization. In essence, if the planner were rational, all-knowing, and all-powerful, the planner could simply force households to optimize, thereby producing efficient consumption dynamics.

We show that consumption smoothing arises, even when households know more about

¹Naturally, many governments are not trying to maximize the (equal-weighted) well-being of their citizens. Also, confusion on the part of government actors may affect their policy-making decisions. We briefly discuss deviations from the rational-utilitarian-government benchmark in the conclusion.

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

their preferences than the social planner and even when the government has a restrictive set of policy tools. In our main result, we show that consumption smoothing arises as long as the government has utilitarian preferences and rational expectations.

Accordingly, consumption smoothing does not identify household rationality. This is because consumption smoothing generated by household optimization is *also* implied by planner optimization. Hence, economies display pseudo-rationality: consumption smoothing that derives from planner rationality, regardless of the scope of household rationality. Pseudo-rationality arises for all concave utility functions with general taste shocks (under the maintained assumption that the planner is a rational utilitarian). We believe that this result is intuitive (and explain why below), but it has been overlooked in previous research despite its important implications for our interpretation of household resource allocation.

Consumption smoothing is the focus of the modern consumption literature (e.g., Bernheim, Skinner, Weinberg 2001; Aguiar and Hurst 2005; Shapiro 2005; Hastings and Washington 2010; Aguila et al 2011; Olafsson and Pagel 2018; Stephens and Toohey 2018; Ganong and Noel 2019; Gerard and Naritomi 2021, Kolsrud et al. 2021).³ To motivate consumption smoothing, consider the household-level Euler equation in a fully *rational* economy, where each household, i , has an idiosyncratic multiplicative taste shifter, θ_i , known to the household at date t , with mean $\bar{\theta}$, (and no other source of uncertainty),

$$u'(c_{it}) = \delta R \theta_i u'(c_{it+1}) \quad \forall i, t. \tag{1}$$

This equation makes predictions about consumption smoothing on average in the population. For example, with \ln utility equation (1) implies

$$\frac{c_{it+1}}{c_{it}} = \theta_i \delta R.$$

³In keeping with a methodological emphasis on household-level microfoundations, we do *not* base our analysis on a representative agent framework.

Following the consumption literature, we construct an average across I households.

$$\frac{1}{I} \sum_{i=1}^I \frac{c_{it+1}}{c_{it}} \rightarrow_D \bar{\theta} \delta R. \quad (2)$$

Hence, average household-level consumption growth converges (in distribution) to $\bar{\theta} \delta R$.⁴

The key contribution of the current paper is to show that this consumption smoothing property for the *rational* economy also applies to a *behavioral* economy with a paternalistic planner. To illustrate the effect of a paternalistic government, we study a stripped-down 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 (i.e., a defined contribution retirement account with auto-enrollment).⁵ We consider an economy with three types of households: optimizing households, who behave optimally throughout their life-cycle; myopic households, who opt out of the default and consume all of their *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 stack the deck 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 sub-optimal behavior of myopic and passive households.

Our behavioral economy is populated by an arbitrary mass of myopic and passive households with privately known taste shifters. The government only knows the aggregate distribution of types. In the paternalistic equilibrium, the Euler equation does *not* hold for each individual household in the behavioral economy, but it *does* hold when household-level marginal utilities are averaged over the cross-section of households. When the taste shifters

⁴The practice of studying *average* consumption growth (averaging growth rates across households) is methodologically motivated in part by idiosyncratic noise (including measurement error).

⁵In the appendix, we show that this restriction to only two policy levers is without loss of generality.

are multiplicative we obtain,

$$\frac{1}{I} \sum_{i=1}^I [u'(c_{it}) - \delta R \theta_i u'(c_{it+1})] \rightarrow_D 0 \quad \forall i, t.$$

This, in turn, implies that the *same* average consumption growth rate (equation 2) emerges for the rational and the behavioral economies. In other words, consumption smoothing (on average) is the same regardless of the fraction of rational, myopic, and passive agents in the economy. We show how this result changes only slightly when we generalize the utility function from \ln utility to constant relative risk aversion.

Although we show that consumption smoothing is not generally diagnostic of household optimization, we demonstrate that other types of economic information – both cross-sectional and time series evidence – identify the extent of household optimization. For example, in our particular model, bunching 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 provide ways to overcome the ambiguous attribution problems that we highlight. We show how to use existing quasi-experimental methods to identify the mass of optimizing, myopic, and passive households in our stylized model.

Our choice of the particular framework for illustrating our arguments – namely, savings over the life-cycle – is motivated by the partially contradictory findings of the recent research on household savings. As we referenced above, some papers find evidence that is consistent with optimal savings while others highlight savings anomalies.⁶ Our model predicts that household-level sub-optimization, arising from myopia and passivity, will be partially offset by paternalistic policies, like defaults and Social Security. Hence, the economy that our

⁶For example, Scholz et al. (2006) estimate that less than twenty percent of households in the HRS under-save for retirement. Aguiar and Hurst (2005) and Aguila et al. (2011) report evidence of consumption smoothing between working life and retirement. However, Banks, Blundell, and Tanner (1998), Bernheim, Skinner, and Weinberg (2001), and Stephens and Toohey (2018) report that consumption declines in the transition to retirement. Another body of research studies employer retirement savings plans, and shows that automatic enrollment and other institutional nudges have an effect on savings (e.g., Madrian and Shea 2001, Choi et al. 2004, Beshears et al. 2009, Chetty et al. 2014, Blumenstock et al. 2018, Choukmane 2019). A related body of research finds that consumers have a high marginal propensity to consume out of both anticipated and unanticipated windfalls (e.g., Soulelos 1999, Parker 1999, 2017, Stephens 2003, Stephens and Unayama 2011, Parker et al. 2013).

model describes simultaneously features both household-level ‘mistakes’ *and* consumption smoothing on *average* across all 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 fully 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 a paper that introduced the concept of 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 freedom-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), Loginova and Persson (2012), Bubb and Pildes (2014), Alcott and Taubinsky (2015), Bubb, Corrigan, and Warren (2015), Chetty (2015), Spinnewijn (2015), Moser and de Souza e Silva (2015), Bubb and Warren (2016), Lockwood (2020), and Farhi and Gabaix (2020). Another closely related literature studies *self-binding* policies, in other words, self-directed paternalism: e.g., Ashraf, Karlan, and Yin (2006), Giné, 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 (2019).

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 privately observed taste shifters that these households experience. We also describe the policy levers available to the government. Finally, we characterize the equilibrium behavior of the households. Section 3 derives the equilibrium behavior of

the social planner. In this section we show that the classical Euler equation holds (averaging marginal utilities across households in the cross-section), regardless of the proportions of optimizers, myopes, and passives, as long as the social planner is a rational utilitarian. Then we focus on our primary object of interest: consumption smoothing. In Section 4, we study the special case of multiplicative taste shifters and constant relative risk aversion. Using this special case 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 (averaging the growth rate of consumption across households in the cross-section) is a robust feature of our model. In Section 5, we discuss the issue of identification – how can the distribution of optimizing, myopic, and passive households be identified in our setting? Section 6 concludes and discusses political economy extensions.

2 Baseline Model

Setup. We consider a two-period model, where period $t = 1$ is working life and period $t = 2$ 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 $R \equiv 1 + r$. The life-time budget constraint household i faces is

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

Here we assume that income is exogenous.⁷ We also simplify our framework by imposing a within-household budget constraint, which rules out redistribution.

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

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

⁷Endogenous income would not change the basic result of the paper: a utilitarian planner adopts policies that moves the economy toward allocative efficiency. We leave model extensions that include endogenous decisions about how long and how much to work to future research. See Lockwood (2020) for this type of analysis.

where δ is a discount factor.⁸ The taste parameter $\theta \in \Theta$ varies across households and is independently drawn from a cumulative distribution function $F(\theta)$ over a (non-degenerate) compact space Θ . We assume that the realization of θ is known to each household at time 0. We assume that the population's cumulative distribution function, $F(\theta)$, is known by the government, but the government does not know the realization of θ for each household. 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 savings account. The forced savings account enables the planner to control the behavior of myopes. We assume that the planner sets a default level of savings in the voluntary savings account, s_D , which enables the planner to control the behavior of passives. Households are able to opt-out of this default at zero cost. In addition, the planner sets forced savings (a minimum level of savings from labor income) of s_F , which is deposited into the forced savings account during working life.⁹ Because s_F is *defined* as the minimum level of savings, passively following the default must engender a level of savings at least as great as s_F . In notation, $s_F + s_D \geq s_F$. Hence, an operational default requires

$$s_D \geq 0. \tag{3}$$

In Appendix II, 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: *Optimizers*, *Myopes*, and *Passives*.

⁸The 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 $\widehat{\delta} = \delta\theta$.

⁹This is essentially what has been adopted by Australia, Israel, and Singapore, and has similarities to Social Security in the US.

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

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

subject to two constraints:

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

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

The first inequality is the budget constraint. The second inequality 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. For unconstrained optimizers, 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.¹⁰ 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

$$\begin{aligned} W \equiv & \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta)] dF(\theta) \\ & + \mu_M \int_{\Theta} [u_1(c_1^M; \theta) + \delta u_2(c_2^M; \theta)] dF(\theta) \\ & + \mu_P \int_{\Theta} [u_1(c_1^P; \theta) + \delta u_2(c_2^P; \theta)] dF(\theta). \end{aligned} \quad (4)$$

Recall that this expression is maximized subject to the constraint $s_D \geq 0$ (equation (1)).¹¹ This framework implies no inter-household transfers, so the government’s only role is to influence intertemporal allocations and not to create redistribution across households. For an analysis of government policies in a setting with redistribution, see Moser and Olea de Souza e Silva (2019) and Beshears et al (2020).

In the introduction, we motivated and summarized our results by reporting averages of finite populations – e.g., $\frac{1}{I} \sum x_i$ – and describing convergence in distribution for large I . This

¹⁰A savings subsidy is not part of the optimal mechanism in our setting. A savings subsidy would not affect the behavior of myopes or passives. It would (adversely) distort the choices of optimizers. Accordingly, it is socially suboptimal. However, in an economy with present-biased agents (with homogeneous present bias) a savings subsidy is optimal (see Beshears et al. 2020; this other setting allows inter-household fiscal transfers). Nonetheless, if agents with present bias have highly heterogeneous degrees of present bias, savings subsidies once again cease to be an important part of the optimal mechanism.

¹¹We assume zero (administrative) cost for forced savings (e.g., Social Security) and defaults. If these programs were costly to implement on the margin, the utilitarian planner would use them less and our results would accordingly be weakened. However, if the cost of such programs is a fixed cost, and if those fixed costs are small in relation to the aggregate economy, as many have argued (e.g., National Academy of Social Insurance 2003; Benartzi et al. 2017), our results would not change.

is natural notation when one works with a discrete set of observations. It is more notationally compact to assume we are studying a continuum of agents. We therefore shift from describing an average of a finite set of households to an integral of a continuum of households. Accordingly, we use the expectation operator, $E[\cdot]$, which is the *integral* 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 shifters (θ), which are random variables from the perspective of an outside observer:

$$E[x(\theta)] \equiv \mu_O \int_{\Theta} x^O(\theta) dF(\theta) + \mu_M \int_{\Theta} x^M(\theta) dF(\theta) + \mu_P \int_{\Theta} x^P(\theta) dF(\theta). \quad (5)$$

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 (and do not face binding constraints), 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 attributed to household optimization also appear in our partially myopic and partially passive economy as a result of the planner's intervention.

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

$$E [u'_1(c_1; \theta)] = \delta RE [u'_2(c_2; \theta)]. \quad (6)$$

Recall that the household-level subscript i has been suppressed to simplify notation. As we explained above, the expectation operator is taken across households (with different values of θ).

Proposition 1 establishes that the Euler equation (6) holds for any mass vector μ characterizing the fraction of optimizing, myopic, and passive households. The results in the rest of the paper also have the property that a classical optimality condition holds, averaging across households in the population, regardless of the fraction of optimizing households. Proposition 1 uses the property that policies that change consumption at the margin are optimally implemented (in this setting) by using a default savings rate and a required minimum savings rule. If by contrast, we studied an economy where it were optimal to change consumption at the margin with tax incentives, then a price wedge would be introduced and the Euler equation would no longer exactly hold in the aggregate.

The proof of Proposition 1 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. Mandatory savings, s_F , has a negative impact on optimizers, some of whom are constrained by s_F , which becomes a binding lower bound for optimizers with sufficiently low values of θ . This binding

lower bound lowers the welfare of optimizers (relatively to the case of laissez faire). Note that s_D only affects the choices of passives, so it has no negative spillover effects on either optimizers or passives. This asymmetry – s_F has negative spillovers on optimizers and s_D does not – leads the planner to use s_D to supplement s_F . The planner sets the default savings rate, s_D , strictly greater than zero, which implies that, in equilibrium, myopes save less than passives: $s_F < s_F + s_D$. This first lemma ($s_D > 0$) is proved in Appendix I.

Our second lemma describes a change of variables for the planner’s optimization problem.

Lemma 2 (Change of Variables) *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 (4) as choosing c_1^D and \bar{c}_1 to maximize

$$\begin{aligned}
 W \equiv & \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(R(y - c_1(\theta)); \theta)] dF(\theta) \\
 & + \mu_M \int_{\Theta} [u_1(\bar{c}_1; \theta) + \delta u_2(R(y - \bar{c}_1); \theta)] dF(\theta) \\
 & + \mu_P \int_{\Theta} [u_1(c_1^D; \theta) + \delta u_2(R(y - c_1^D); \theta)] dF(\theta),
 \end{aligned} \tag{7}$$

subject to the constraint

$$c_1^D \leq \bar{c}_1, \tag{8}$$

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

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

Proof: Lemma 1 establishes that at an optimum $\bar{c}_1 > c_1^D$. In other words, myopes consume in period 1 *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 (8)) 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 $c_2^P = R \times (y - c_1^D)$. 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

$$\int_{\Theta} u'_1(c_1^P; \theta) dF(\theta) = \delta R \int_{\Theta} u'_2(c_2^P; \theta) dF(\theta) \quad (9)$$

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

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

$$\begin{aligned} \int_{\Theta} \left[\mu_O u'_1(c_1(\theta); \theta) \frac{dc_1(\theta)}{d\bar{c}_1} + \mu_M u'_1(c_1^M; \theta) \frac{dc_1^M(\theta)}{d\bar{c}_1} \right] dF(\theta) \\ + \delta \int_{\Theta} \left[\mu_O u'_2(c_2(\theta); \theta) \frac{dc_2(\theta)}{d\bar{c}_1} + \mu_M u'_2(c_2^M; \theta) \frac{dc_2^M}{d\bar{c}_1} \right] dF(\theta) = 0. \quad (10) \end{aligned}$$

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 $\theta \in \Gamma(\bar{c}_1)$.¹² This implies that

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

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 (10) therefore reduces to:

$$\begin{aligned} \int_{\theta \in \Gamma(\bar{c}_1)} \mu_O u'_1(c_1(\theta); \theta) dF(\theta) + \int_{\theta \in \Theta} \mu_M u'_1(c_1^M; \theta) dF(\theta) \\ = \delta R \int_{\theta \in \Gamma(\bar{c}_1)} \mu_O u'_2(c_2(\theta); \theta) dF(\theta) + \delta R \int_{\theta \in \Theta} \mu_M u'_2(c_2^M; \theta) dF(\theta). \end{aligned} \quad (11)$$

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

$$u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta). \quad (12)$$

Combining equations (9), (11), and (12), we have

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

■

¹²This 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 $\frac{dc_1(\theta)}{d\bar{c}_1} = 0$.

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

In summary, Proposition 1 implies that a classical Euler equation characterizes equilibrium allocations in the economy, regardless of the proportions of optimizers, myopes, and passives. Note that the Euler equation holds in expectation across all households, but it does not hold at the level of each individual household. Some households consume too little in period 1 (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 in period 1 (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 $c_1 < c_1^D$). The Euler equation is only satisfied on *average* in the population (averaging marginal utilities across households).

3.1 Could an Omniscient 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, at the household level, the Euler equation will not be satisfied (though it will be satisfied for each unconstrained optimizing household). An omniscient econometrician with full information – i.e., an econometrician who knows 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 these equations are *not* universally satisfied. However, such a direct test of the Euler equation is not implementable in practice, because it relies on knowledge of the unobservable household-level taste shifter, θ .

We now turn to the sorts of consumption smoothing tests that are implementable in practice and are frequently undertaken.

4 Consumption Smoothing

To study consumption smoothing we need to place some structure on our utility function. Accordingly, we 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 the case of multiplicative taste shocks (c.f., Atkeson and Lucas 1992; Amador, Werning, and Angeletos 2006; and Beshears et al. 2020). Without loss of generality and to simplify notation, we assume that $\delta R = 1$. We maintain these assumptions throughout this section.

As a first step in the analysis of consumption smoothing, we 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 in the fully optimizing economy yields

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

When $E[\theta] = 1$ (a leading 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 lemma we prove that these relationships hold in our behavioral economy, which contains a mix of optimizers, myopes and passives.

Lemma 4 *Assume a rational utilitarian 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 (9) 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 (11) 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

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

■

The degree of consumption smoothing between working life and retirement is often used as proxy for 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 utility functions with constant relative risk aversion.

Specifically, we study the case of $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, where $\gamma > 0$.¹³ We analyze consumption growth, c_2/c_1 , under the assumption that $E[\theta] = 1$. We continue to assume that unconstrained optimizers set c_1 and c_2 as optimal functions of θ ; constrained optimizers set $c_1 = \bar{c}_1$ and $c_2 = R \times s_F$; myopes behave the same as constrained optimizers; and passives accept the default allocation $c_1 = c_1^D$ and $c_2 = R \times (s_F + s_D)$. When we write an expectation operator, $E[\cdot]$, we are integrating over types and (as appropriate) values of θ . Recall equation (5).

For \ln utility the results are very simple and do not depend on the fraction of optimizers, myopes, and passives in the population.

Corollary 1 *Assume \ln utility and any distribution of optimizing, myopic, and passive households. Assume taste shifters have an average value of unity ($E[\theta] = 1$). Then*

$$E \left[\frac{c_2 - c_1}{c_1} \right] = 0.$$

¹³As $\gamma \rightarrow 1$, this function converges to $\ln(c)$.

Proof: Lemma 4 implies that

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

For the case of \ln utility and $E[\theta] = 1$, this equality implies

$$E \left[\frac{c_2 - c_1}{c_1} \right] = 0.$$

■

There is an analogous result for the case of constant relative risk aversion (*CRRA*), which is the case typically studied in the modern consumption literature.

Corollary 2 *Assume CRRA utility and any distribution of optimizing, myopic, and passive households. With no taste shocks on average at retirement ($E[\theta] = 1$), then*

$$E \left[\left(\frac{c_2}{c_1} \right)^\gamma - 1 \right] = 0.$$

This corollary is also proven by manipulating Lemma 4. Corollary 2 has the advantage that it can be directly evaluated using consumption data and the (homogeneous) coefficient of relative risk. It has the disadvantage that it is the expectation of a (power) transformation of consumption growth and not consumption growth itself. If we simply studied consumption growth without the γ power, then we would find that (outside the \ln utility case) average consumption growth is dependent on the distribution of optimizing, myopic, and passive households.

Consumption growth averaged across households can be calculated for the *CRRA* economy when all households are *optimizers*. Then, the classical Euler equation applies for each household, so

$$u'(c_1) = \theta u'(c_2),$$

which implies

$$E \left[\frac{c_2 - c_1}{c_1} \right] = E \left[\theta^{1/\gamma} - 1 \right].$$

This sets the benchmark for a classical economy. Note that the average consumption growth rate is not generically 0 (with an exception being the joint case $E[\theta] = 1$ and $\gamma = 1$).

4.1 Numerical illustrations

Once we add myopes and passives, we calculate the average growth rate of consumption (by numerically solving for the minimum savings level that the planner will choose – see equation (11)). We now present two examples of this mixed-type case, to illustrate the degree of consumption smoothing that is achieved in equilibrium when the population contains a large mass of myopes and passives.

Henceforth, we assume θ is distributed uniformly on $[1/2, 3/2]$.¹⁴ Assuming that all agents are optimizers and that $\gamma = 2$ implies

$$E\left[\frac{c_2 - c_1}{c_1}\right] = E\left[\theta^{1/2} - 1\right] = -1.1\%.$$

By contrast, replace the homogeneous population of optimizers with a mixed population of optimizers, myopes, and passives, so that $\mu_O = \mu_M = \mu_P = 1/3$. In this case,

$$E\left[\frac{c_2 - c_1}{c_1}\right] = -0.2\%.$$

Now, we consider a value of relative risk aversion below unity (and continue to assume the same distribution of θ). Assuming that $\gamma = 1/2$ and that all households are optimizers implies that

$$E\left[\frac{c_2 - c_1}{c_1}\right] = E\left[\theta^2 - 1\right] = 8.3\%.$$

Replace again the homogeneous population of optimizers with a mixed population of optimizers, myopes, and passives, so that $\mu_O = \mu_M = \mu_P = 1/3$. This implies

¹⁴With this assumption, we find using equation (11) that the cutoff value $\bar{\theta} \in [1/2, 3/2]$ such that all optimizers with lower values of θ are constrained is $\bar{\theta} = 0.9142$. The consumption ratio, $\frac{c_2}{c_1}$, is $\frac{c_2(\bar{\theta})}{c_1(\bar{\theta})} = \bar{\theta}^{1/\gamma}$ for unconstrained optimizers, $\frac{c_2^M}{c_1^M} = \bar{\theta}^{1/\gamma}$ for myopes and constrained optimizers, and $\frac{c_2^P}{c_1^P} = 1$ for passives.

$$E \left[\frac{c_2 - c_1}{c_1} \right] = 1.7\%.$$

4.2 Special case of no optimizers

In the absence of any optimizing agents (i.e., $\mu_O = 0$), average consumption growth is exactly equal to 0 for every household (regardless of the coefficient of relative risk aversion). For this case, the planner constrains all households using the same minimum savings threshold (and utilizes no additional savings for passives). Intuitively, households cannot be relied upon to reveal their private information about θ because the households are all either myopes or passives. If all households are being forced to have the same values of c_1 and c_2 , then the planner's problem reduces to

$$\max_{c_1} E [u(c_1) + \delta \theta u(R(y - c_1))].$$

Now, the planner's first order condition is

$$u'(c_1) = E[\theta] \delta R u'(R(y - c_1)) = u'(c_2).$$

Hence, consumption growth is zero for every household.

This special case (no optimizers) also yields a specific prediction for wealth at the end of working life (i.e., retirement wealth):

$$W_2 = c_2 = \frac{y}{2}.$$

In other words, the household is forced to save half of their lifetime resources for retirement. This result generalizes for the case in which retirement represents fraction α of the adult lifespan. Then fraction α of adult lifetime resources are saved for retirement.

4.3 From theory to practice: consumption smoothing

As discussed in the introduction, many papers use consumption smoothing as a key test of household optimization. The tests in this literature map onto the tests that we have described in this section: i.e., studying the average value of

$$\frac{c_{i,t+1} - c_{i,t}}{c_{i,t}},$$

or the closely related average value of $\Delta \ln c_{i,t+1}$ (for all $i \in 1, \dots, I$).¹⁵ Our results show that utilitarian planners will elicit consumption smoothing (on average) even when some economic agents are myopic or passive. Hence, in our framework, an observed failure of consumption smoothing (e.g., Bernheim, Skinner, Weinberg 2001; Shapiro 2005; Hastings and Washington 2010; Olafsson and Pagel 2018; Stephens and Toohey 2018; Ganong and Noel 2019; Gerard and Naritomi 2021, Kolsrud et al. 2021) is sufficient but not necessary for imperfect household optimization. Likewise, an observed pattern of consumption smoothing (Aguiar and Hurst 2005; Aguila et al 2011) is not sufficient for household optimization.

The same logic applies to other moments from the household balance sheet. Consumption smoothing is closely related to savings. For example, we showed in the previous subsection that if the economy is comprised solely of myopes and passives, households will be forced to save fraction α of their adult lifetime resources for retirement (where retirement is fraction α of the adult lifespan). In our framework, finding that households accumulate enough wealth to smooth consumption (e.g., Scholz, Seshadri, and Khitatrakun 2006) is not sufficient to argue that households are optimizers. Wealth formation to support consumption smoothing (on average) is a property of economies with behavioral agents and utilitarian planners. However, certain balance sheet moments will distinguish behavioral agents from optimizers:

¹⁵For example, consider the benchmark estimation equation in Bernheim, Skinner and Weinberg (2001; equation (7)), which studies $\Delta \ln c_t$ and is motivated as an Euler Equation for CRRA utility. Likewise, Ganong and Noel (2019) also study consumption growth. Closely related, Aguiar and Hurst (2005) and Stephens and Toohey (2018) study dynamics of $\ln c_t$ which is isomorphic to studying $\Delta \ln c_t$. Because these papers all use disaggregated (household) data, the estimating equations are implicitly studying average consumption growth where the average is taken in the cross-section across households, and compared across time periods (e.g., varying by age or UI status).

e.g., households that routinely save in partially illiquid assets with much lower returns than the interest rates on credit cards on which they hold large balances (see Angeletos et al. 2001; Laibson, Maxted, and Moll 2021; Laibson et al 2022).

5 Identification of Optimizers, Myopes, and Passives

While consumption smoothing on average does not identify the distribution of household types, other types of already widely used empirical strategies do reveal the share of optimizing and non-optimizing households. In this section, we use our positive model to illustrate empirical strategies that identify the share of different types of agents.

First, we describe a strategy for identifying all of the variables of the model using only existing observational data. Then we discuss how one could complement this analysis by exploiting experimental variation.

Assume that the econometrician only has existing observational data and wishes to estimate all of the underlying parameters, including the mass vector μ . First, the mass of passive agents, μ_P , is given by the mass of agents who are saving exactly $s_F + s_D$ (e.g., Madrian and Shea 2001). Second, identify the set of agents who are saving in the interior of the action space, $s > s_F$ and *not* at the passive default. These are the optimizers who are unconstrained. Use their consumption choices to identify a truncated set of taste shifters. Specifically, for each household in the set of unconstrained optimizers invert the Euler equation, $u'_1(c_1(\theta); \theta) = \delta R u'_2(c_2(\theta); \theta)$, to calculate a household-specific value θ_i . (Here we study the special case where the Euler equation is invertable.) These values of θ are drawn from a truncated density. This density could be structurally estimated and projected into the full space (i.e., the shape and mass of the unobserved tail could be inferred by the observed part of the distribution). This projection enables the researcher to infer the “missing mass” of (constrained) optimizers, completing inference of μ_O . Finally, it follows that $\mu_M = 1 - \mu_O - \mu_P$.

We now discuss two types of experimental inferences¹⁶ that could be made independently

¹⁶For instance, policy changes might arise as the government learns about the efficacy of new policies (e.g.,

of the analysis that we just discussed. Specifically, we discuss quasi-experiments that study behavioral responses to policy variations, 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 $s_F + s'_D$ (e.g., Chetty, Friedman, Leth-Petersen, Nielsen, and Olsen 2014).

Second, changes to the level of forced savings can shed light on the proportion of myopic households and of optimizing households. Recall that in our setting, the mass of households that choose a savings rate s equal to the mandatory (minimum) savings rate, s_F , is given by

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

Consider a change in the level of mandatory savings and evaluate the change in the mass of all households that continue to choose to save at the mandatory minimum 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 in equilibrium because there are optimizing households on the boundary for the case of a mixed-type economy. Using the utility function and the distribution of θ , $F(\theta)$, the derivative $\frac{d\Pr(\theta \in \Gamma(\bar{c}_1))}{ds_F}$ can be calculated. Hence, in such a case, observing the empirical magnitude $\frac{d\Pr(s=s_F)}{ds_F}$ (obtained, for example, from a natural experiment as in Lindeboom and Montizaan 2020, for the isomorphic case of a decrease in public pension wealth) enables calculation of the mass of optimizing agents, μ_O . Once μ_O is known, the mass of myopes is given by:

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

The analysis discussed in this subsection is derived from our specific positive model. However, the broader point is that an analysis of cross-sectional distributions of economic outcomes can provide tests for household optimization that are not confounded by the pres-

the research that led to passage of the Pension Protection Act of 2006).

ence of a rational utilitarian planner.

6 Conclusion

We study a simple setting that illustrates the interactions between a utilitarian social planner and heterogeneous households, some of whom optimize, some of whom are myopic, and some of whom are passive. In this setting, planner optimization is a partial substitute for household optimization. This substitution arises because the utilitarian 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). Specifically, the Euler equation holds in the sense that average marginal utility (across households) before retirement is equal to average marginal utility after retirement. These Euler equation properties arise generally, whether or *not* all (or even some) households are optimizers.

In turn, these results imply that consumption-smoothing tests 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. 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 optimization than consumption growth on average in the cross section.

Our conclusions depend upon the assumption that the government is a fully optimizing utilitarian. In practice, governments often fall short of this benevolent benchmark, despite (or, sometimes, because of) the pressures that they face to obtain and hold power. This leads to a follow-up question: how would our results change if the government is not utilitarian, but is instead a self-interested (and/or ideological) politician or political party? In a democracy, the answer depends on three key considerations: what is the voting frequency of different types of households, to what degree do households see their behavioral propensities as biases that they want the government to address, and to what degree do altruistic motives (i.e., helping other households overcome biases) influence voting?

For example, a government with a populist policy agenda might try to appeal to the myopic types in our model by advancing policies that maximize the scope for current consumption. Such a government could be elected if myopic types represent an effective voting block (and/or form coalitions with other types), and if myopic types do not view myopia as a problem/bias in their own behavior or in the behavior of others.

On the other hand, a government might cater to the preferences of socio-economic elites, and therefore over-weight their interests (relative to a utilitarian's uniform weights). If socio-economic elites tend to have more education and to be more like optimizers (and therefore less myopic and passive), then the resulting equilibrium would be characterized by subsidies that are disproportionately used by elites (cf. Kolsrud et al 2021 in the context of pension incentives for delayed retirement). Relatedly, a government with an ideological libertarian or laissez faire orientation would intervene less than our (utilitarian) model implies.

Alternatively, a paternalistic government that uses greater welfare weights for myopes and passives relative to optimizers, would introduce more forced savings mechanisms, thereby engendering a jump up in average consumption between working life and retirement (driven by optimizers with high values of θ who choose to save more than they are forced to save).

We have also made assumptions about the rationality of the government. A utilitarian or non-utilitarian government might be confused about the distribution of behavioral types (e.g., optimizers, myopes, and passives), and/or by the distribution of θ . If the government is a utilitarian with imperfect rationality, then the resulting policy implications would not be materially different from the ones that we have derived, as long as the government's deviations from rationality are small. In our setting, the equilibrium government policies are continuous in the deep parameters as perceived by the government, so the policy predictions are robust to a small degree of government irrationality. We speculate that governments would be right on average, which would tend to make the analysis that we undertake true (to first order) when averaging across utilitarian governments, though not true for each utilitarian government.

A complete positive political economy model would need to (i) predict the high-level goals (e.g., utilitarian or non-utilitarian) of governments that come to power, (ii) predict their de-

viations from rationality, and also (iii) predict the policies they accordingly adopt, sometimes through complicated intra-governmental bargaining processes. We have completed the contingent analysis for one important boundary case: an optimizing utilitarian government with rational expectations.

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] Banks, James, Richard Blundell, and Sarah Tanner (1998): "Is There a Retirement-Savings Puzzle?," *American Economic Review*, 88(4): pp. 769-788.
- [9] Benartzi, Shlomo, John Beshears, Katherine L. Milkman, Cass R. Sunstein, Richard H. Thaler, Maya Shankar, Will Tucker-Ray, William J. Congdon, and Steven Galing (2017): "Should Governments Invest More in Nudging?" *Psychological Science*, 28(8): pp. 1041-1055.

- [10] Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg (2001): "What Accounts for the Variation in Retirement Wealth among US Households?," *American Economic Review*, 91(4): pp. 832-857.
- [11] Beshears, John, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian (2020): "Optimal Illiquidity," NBER Working Paper 27459.
- [12] Beshears, John, James J. Choi, Christopher Harris, David Laibson, Brigitte C. Madrian, and Jung Sakong (2020): "Which Early Withdrawal Penalty Attracts the Most Deposits to a Commitment Savings Account?," *Journal of Public Economics*, 183, Article 104144.
- [13] Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2018) "Behavioral Household Finance" in eds. B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, *Handbook of Behavioral Economics - Foundations and Applications 1*, Volume 1, Elsevier.
- [14] 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.
- [15] Blumenstock, Joshua, Michael Callen, and Tarek Ghani (2018): "Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan," *American Economic Review*, 108(10): pp. 2868-2901.
- [16] Bryan, Gharad, Dean Karlan, and Scott Nelson (2010): "Commitment Devices," *Annual Review of Economics*, 2(1): pp. 671-698.
- [17] 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.
- [18] Bubb, Ryan, and Richard Pildes (2014): "How Behavioral Economics Trims Its Sails and Why," *Harvard Law Review*, Vol 127.

- [19] Bubb, Ryan, and Patrick L. Warren (2016): "An Equilibrium Theory of Retirement Plan Design," NYU Law and Economics Research Paper No. 16-28.
- [20] 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.
- [21] Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick (2009): "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 124(4): pp. 1639-1674.
- [22] Chetty, Raj (2015): "Behavioral Economics and Public Policy: A Pragmatic Perspective," *American Economic Review*, 105(5): pp. 1-33
- [23] 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.
- [24] Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick (2003): "Optimal Defaults," *American Economic Review*, 93(2): pp. 180-185
- [25] 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.
- [26] Choukhmane, Taha (2020): "Default Options and Retirement Saving Dynamics," Working Paper.
- [27] DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F. Schmieider (2017): "Reference-Dependent Job Search: Evidence from Hungary," *Quarterly Journal of Economics*, 132(4): pp. 1969-2018.

- [28] Diamond, Peter A. (1977): "A Framework for Social Security Analysis," *Journal of Public Economics*, 8(3): pp. 275–298.
- [29] Fang, Hanming, and Dan Silverman (2009): "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY," *International Economic Review*, 50(4): pp. 1043-1077.
- [30] Farhi, Emmanuel, and Xavier Gabaix (2020): "Optimal Taxation with Behavioral Agents," *American Economic Review*, 110 (1): 298-336.
- [31] Feldstein, M. (1985): "The Optimal Level of Social Security Benefits," *Quarterly Journal of Economics*, 100(2): pp. 303-320.
- [32] Ganong, Peter, and Pascal Noel (2019): "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, 109(7): pp.2383-2424.
- [33] Gerard, François, and Joana Naritomi (2021): "Job Displacement Insurance and (the Lack of) Consumption-Smoothing," *American Economic Review*, 111(3): pp.899-942.
- [34] 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.
- [35] Gruber, Jonathan, and Botond Köszegi (2001): "Is Addiction ‘Rational’? Theory and Evidence," *Quarterly Journal of Economics*, 116(4): pp. 1261-1303.
- [36] Hastings, Justine, and Ebonya Washington (2010): "The First of the Month Effect: Consumer Behavior and Store Responses," *American Economic Journal: Economic Policy*, 2(2): pp.142-162.
- [37] Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan (2015): "Self-Control at Work," *Journal of Political Economy*, 123(6): pp. 1227-1277.

- [38] Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn (2021): "Retirement Consumption and Pension Design," Working Paper.
- [39] Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers (1982): "The Adequacy of Savings," *American Economic Review*, 2(5): pp. 1056–1069.
- [40] Laibson, David, Sean Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman (2022). "Estimating Discount Functions with Consumption Choices Over the Lifecycle." Harvard mimeo.
- [41] Laibson, David, Peter Maxted, and Benjamin Moll. (2021) "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy." No. w29094. National Bureau of Economic Research, 2021.
- [42] Lindeboom, Maarten and Raymond Montizaan (2020): "Disentangling retirement and savings responses," *Journal of Public Economics*, Vol. 192.
- [43] Lockwood, Benjamin B. (2020) "Optimal income taxation with present bias." *American Economic Journal: Economic Policy* 12.4: 298-327.
- [44] Loginova, Uliana, and Petra Persson (2012): "Paternalism, Libertarianism, and the Nature of Disagreement," Working Paper.
- [45] 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.
- [46] Moser, Christian, and Pedro Olea de Souza e Silva (2019): "Paternalism vs Redistribution: Designing Retirement Savings Policies with Behavioral Agents," Working Paper, Columbia University.
- [47] National Academy of Social Insurance (1998): "Evaluating Issues in Privatizing Social Security: Report of the Panel on Privatization of Social Security."

- [48] 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.
- [49] O'Donoghue, Ted, and Matthew Rabin (2006): "Optimal Sin Taxes," *Journal of Public Economics*, 90(10): pp. 1825-1849.
- [50] Olafsson, Arna, and Michaela Pagel (2018): "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *Review of Financial Studies*, 31(11): pp. 4398-4446.
- [51] Parker, Jonathan A. (1999): "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 89(4): pp. 959-973.
- [52] Parker, Jonathan A. (2017): "Why Don't Households Smooth Consumption? Evidence from a \$25 Million Experiment," *American Economic Journal: Macroeconomics*, 9(4): pp.153-183.
- [53] Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland (2013): "Consumer Spending and The Economic Stimulus Payments of 2008," *American Economic Review*, 103(6): pp. 2530-2553.
- [54] Poterba, James M. (2014): "Retirement Security in an Aging Population," *American Economic Review*, 104(5): pp. 1-30.
- [55] Schilbach, Frank (2019): "Alcohol and Self-Control: A Field Experiment in India," *American Economic Review*, 109 (4): 1290-1322.
- [56] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006): "Are Americans Saving 'Optimally' for Retirement?," *Journal of Political Economy*, 114(4): pp. 607-643.
- [57] Shapiro, Jesse M. (2005): "Is There A Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 89(2-3): pp. 303-325.
- [58] Skinner, Jonathan (2007): "Are You Sure You're Saving Enough for Retirement?," *Journal of Economic Perspectives*, 21(3): pp. 59-80.

- [59] Souleles, Nicholas S. (1999): "The Response of House hold Consumption to Income Tax Refunds," *American Economic Review*, 89(4), pp. 947-58.
- [60] Spinnewijn, Johannes (2015): "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 13(1): pp. 130-167.
- [61] Stephens, Melvin (2003): "'3rd of tha Month': Do Social Security Recipients Smooth Consumption Between Checks?," *American Economic Review*, 93(1): pp. 406-422.
- [62] Stephens, Melvin, and Desmond Toohey (2018): "Changes in Nutrient Intake at Retirement," NBER Working Paper 24621.
- [63] Stephens, Melvin, and Takashi Unayama (2011): "The Consumption Response to Seasonal Income: Evidence from Japanese Public Pension Benefits," *American Economic Journal: Applied Economics*, 3(4): 86-118.
- [64] Sunstein, Cass R., and Richard H. Thaler (2003): "Libertarian Paternalism is not an Oxymoron," *University of Chicago Law Review*: pp. 1159-1202.
- [65] Thaler, Richard H., and Cass R. Sunstein (2003): "Libertarian Paternalism," *American Economic Review*, 93(2): pp.175-179.
- [66] Thaler, Richard H., and Cass R. Sunstein (2008): "Nudge: Improving Decisions about Health, Wealth, and Happiness," Yale University Press.

Appendix I. 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 $s_F + s_D$ (holding s_F fixed from the first sub-problem) 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,

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

which can also be expressed as

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

This constraint follows from the constraint $s_D \geq 0$ (equation (3)). At the end of this proof we confirm that an even stronger condition applies: $s_D > 0$, which is the claim in the lemma.

To summarize, we characterize optimal policy for the myopes and optimizers – the minimum savings level s_F – without taking account of optimal policy for the passives. Then we characterize optimal policy for the passives – the *sum* $s_F + s_D$ – holding fixed s_F . At the end of the analysis we verify the constraint $s_D \geq 0$.

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

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

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

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

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, further simplifying the equation above, 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 $\theta \in \Gamma(\bar{c}_1)$. Therefore, for the general case in which the mass of optimizers is non-zero and the mass of myopes is non-zero, then

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

and accordingly

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

At $s_D = 0$,

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

This shows that $\frac{dW}{ds_D} > 0$ at $s_D = 0$. Because the optimization with respect to s_D is globally concave,¹⁷ 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.

■

¹⁷The 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.

Appendix II. Generalization as a Mechanism Design Problem

In this appendix, we show that the equilibrium of our model in Section 3 (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

$$\begin{aligned}
 W \left(c_1(\theta), c_2(\theta), c_1^M, c_2^M, c_1^P, c_2^P \right) &\equiv \mu_O \int_{\Theta} [u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta)] dF(\theta) \\
 &\quad + \mu_M \int_{\Theta} [u_1(c_1^M; \theta) + \delta u_2(c_2^M; \theta)] dF(\theta) \\
 &\quad + \mu_P \int_{\Theta} [u_1(c_1^P; \theta) + \delta u_2(c_2^P; \theta)] dF(\theta), \tag{13}
 \end{aligned}$$

subject to the household behavioral models (respectively, optimizer, myopic, and passive) summarized in Section 2.

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^P , and c_2^P to maximize equation (13) subject to the within-household budget constraints,

$$c_1(\theta) + \frac{c_2(\theta)}{R} \leq y \text{ for all } \theta, \tag{14}$$

$$c_1^M + \frac{c_2^M}{R} \leq y, \tag{15}$$

$$c_1^P + \frac{c_2^P}{R} \leq y, \tag{16}$$

incentive compatibility constraints for optimizers,

$$u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta) \geq u_1(c_1(\theta'); \theta) + \delta u_2(c_2(\theta'); \theta), \quad \forall \theta, \theta' \tag{17}$$

$$u_1(c_1(\theta); \theta) + \delta u_2(c_2(\theta); \theta) \geq u_1(c_1^P; \theta) + \delta u_2(c_2^P; \theta), \quad \forall \theta \tag{18}$$

and a maximally impatient reporting rule for myopes,

$$c_1^M = \max \left\{ \sup_{\theta \in \Theta} c_1(\theta), c_1^P \right\}. \quad (19)$$

We now summarize the mechanism design problem: the planner chooses $\{c_1(\theta)\}_{\theta \in \Theta}$, $\{c_2(\theta)\}_{\theta \in \Theta}$, c_1^P , and c_2^P to maximize the objective in equation (13), subject to the budget constraints and incentive compatibility constraints in equations (14)-(19). 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^P and c_2^P . With this setup, we can now present the following proposition.

Proposition 2 (Characterization as a Mechanism Design Problem) *For a rational utilitarian planner, the mechanism design problem and the constrained problem (i.e., maximizing equation (13) by choosing arguments s_F and s_D) generate the same equilibrium allocation.*

This implies an immediate corollary.

Corollary 3 *For a rational utilitarian planner, the mechanism design problem and the constrained problem (maximizing equation (13) by choosing arguments s_F and s_D) generate the same Euler equation averaging marginal utility across households:*

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

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

Proof (Characterization as a Mechanism Design Problem): We first show that the (second-best) optimal allocation is characterized by a *maximum consumption rule* for optimizers, which we will define after describing an unconstrained allocation. For every type θ , there exists a (full-information) unconstrained allocation, $(c_1^*(\theta), c_2^*(\theta))$, which satisfies budget

balance

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

and the first order condition

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

Next, consider a *different* allocation that assigns consumption $(c_1(\theta), c_2(\theta))$, for all $\theta \in \Theta$.

We now define a *maximum consumption rule*.

Definition: Consider an unconstrained allocation $(c_1^*(\theta), c_2^*(\theta))$, for all $\theta \in \Theta$. Now consider an alternative allocation $(c_1(\theta), c_2(\theta))$, for all $\theta \in \Theta$. This alternative allocation is a ‘*maximum consumption rule*’ if and only if two conditions are both satisfied:

- (i) every type with $c_1^*(\theta) \leq \sup_{\theta \in \Theta} c_1(\theta)$ obtains $c_1(\theta) = c_1^*(\theta)$ and $c_2(\theta) = c_2^*(\theta)$ in the allocation; and
- (ii) every type with $c_1^*(\theta) > \sup_{\theta \in \Theta} c_1(\theta)$ obtains $c_1(\theta) = \sup_{\theta \in \Theta} c_1(\theta)$ and $c_2(\theta) = R(y - \sup_{\theta \in \Theta} c_1(\theta))$ in the allocation.

To prove that an optimal mechanism generates an allocation that is a maximum consumption rule for optimizers, consider a candidate allocation given by $\{\hat{c}_1(\theta)\}_{\theta \in \Theta}$, $\{\hat{c}_2(\theta)\}_{\theta \in \Theta}$, \hat{c}_1^M , \hat{c}_2^M , \hat{c}_1^P , and \hat{c}_2^P that is incentive compatible (i.e., satisfying equations (14)-(19)), and is *not* a maximum consumption rule for the optimizers. Because the candidate allocation is incentive compatible, it follows that $\sup_{\theta \in \Theta} \hat{c}_1(\theta) \leq \hat{c}_1^M$ and $\hat{c}_1^P \leq \hat{c}_1^M$.

Now perturb $\{\hat{c}_1(\theta)\}_{\theta \in \Theta}$, $\{\hat{c}_2(\theta)\}_{\theta \in \Theta}$, \hat{c}_1^M , \hat{c}_2^M , \hat{c}_1^P , and \hat{c}_2^P in the following way. Let $\bar{c}_1 = \sup_{\theta \in \Theta} \hat{c}_1(\theta)$ and construct a new allocation such that (i) every optimizer 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; (ii) every optimizer 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; and (iii) \hat{c}_1^M , \hat{c}_2^M , \hat{c}_1^P , and \hat{c}_2^P stay the same.

Note that this new allocation is a maximum consumption rule for the optimizers. This new allocation is also incentive compatible for optimizers: 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 for optimizers, the new allocation generates a strict improvement in welfare for at least one type of optimizer.¹⁸

Hence, only allocations that are maximum consumption rules for optimizers can be solutions to the mechanism design problem. It follows that the planner's problem can be reduced to the choice of a maximum consumption rule for optimizers – i.e., the choice of \bar{c}_1 – and the choice of c_1^P (note that $c_1^P = c_1^D$). By Lemma 2, the mechanism design problem (choosing $\{c_1(\theta)\}_{\theta \in \Theta}$, $\{c_2(\theta)\}_{\theta \in \Theta}$, c_1^P , and c_2^P to maximize the objective in equation (13) subject to the budget constraints and incentive compatibility constraints in equations (14)-(19)) and the constrained problem (maximizing equation (13) by choosing s_F and s_D) are isomorphic optimization problems. Accordingly, they have the same equilibrium allocation.

■

¹⁸There 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.