Subsidising Self-Insurance

Charles Brendon,*
Christian Hellwig†
Gerard Maideu Morera‡

PRELIMINARY DRAFT

February 14, 2025

Abstract

We study income and savings taxation in a broad class of economies with persistent earnings shocks and incomplete insurance. We show there is a general efficiency case for savings to be subsidised whenever labour income is taxed. Persistence is central to this result. Workers who experience lower productivity save more for precautionary reasons, conditional on income. This violates the separability requirement of the Atkinson-Stiglitz (1976) theorem, implying the tax system should incentivise savings relative to consumption. Surprisingly, our result applies even in settings where the well-known Inverse Euler Equation (IEE) holds. The IEE is often interpreted as motivating the opposite intervention – a savings tax. The apparent difference between the two results rests on different definitions of the savings wedge. Our definition is based on the budgetary effects of saving: there is a subsidy if the government incurs a cost when individuals save. We show this is the relevant measure of the wedge for recovering traditional optimal tax rules based on observable statistics – compensated elasticities, and the realised distributions of production and savings.

*University of Cambridge and Toulouse School of Economics
†Toulouse School of Economics and CEPR
‡Toulouse School of Economics

1 Introduction

This paper considers the appropriate joint, nonlinear tax treatment of income and savings when insurance markets are incomplete. Its main message is that uninsurable, persistent labour productivity shocks imply a general case for governments to subsidise individuals' savings decisions, in the event that labour income is being taxed.

The basic intuition works as follows. *Ceteris paribus*, a marginal savings subsidy – like any price reduction – raises the real wage of affected earners, incentivising their labour supply. Since labour income is taxed, this generates fiscal revenue, which more than offsets the subsidy cost initially. Whilst the same incentive effect could equivalently have been achieved by subsidising both savings and consumption, a savings subsidy is less beneficial to higher earners. This is because higher earners – who bear less future risk, due to persistence – have a lower tendency to save from any given income. This tips the balance in the efficiency-equity trade-off. A savings subsidy lowers the distortionary effects of the direct labour income tax, with minimal cost to redistributive objectives.

This main result arises most starkly in a setting where individuals have additively separable preferences between consumption and labour supply. In this case, we show that it is the only qualitative channel that matters for the relative tax treatment of savings. This is notable, because it is precisely under separability that an inverse Euler equation characterises optimally-designed dynamic insurance schemes. The inverse Euler equation (IEE) is typically interpreted as implying a savings tax. Our result – that it is optimal to subsidise savings – *applies to environments where the IEE holds.* This seems to present a direct contradiction.

The difference between the results derives from an ambiguity in what is meant by the tax wedge on saving, in the context of a distorted economy. Savings are a derived demand. A decision to save more implies a decision to purchase an incremental vector of consumption and leisure quantities, spread across multiple future time periods and states. By conventional envelope logic, optimising individuals will be indifferent at the margin whether they use their savings to take extra leisure or extra consumption, and whether to do so immediately, or many periods into the future. In distorted economies, the policymaker does not share this indifference. If future labour supply is being taxed relative to future consuption, then a pattern of demand that is skewed towards leisure is more costly than one that is skewed towards consumption.

We measure the savings subsidy as the fiscal cost that the government incurs at the margin when an individual saves more, given the pattern of demand that extra saving will in-

duce. Our result is that this cost should, optimally, be positive, if there is a positive labour income tax. Analysis based on the inverse Euler equation typically focuses on a different measure of the savings wedge, based on departures from a classic one-period consumption Euler equation. This shows that *if* individuals *were* to use their savings exclusively for consumption in the next time period, *then* the fiscal authority would raise revenue as savings increase. The fact that individuals do not do this, instead using incremental wealth to raise leisure, is what allows the inverse Euler equation to hold even though savings are costly to the government.

Precisely to abstract from these complications, we initially exhibit our main result in a two-period setting where labour earnings in the second period are untaxed, and so second-period choices are undistorted. In this setting, the marginal tax rate on savings is defined unambiguously, independently of the way savings are used.

Throughout the paper, the way that we set up the optimal tax problem departs from standard practice, and so requires some prior comment. We consider an economy that exists over multiple discrete time periods, but we limit attention to the improvements that can be obtained by redesigning the tax treatment of consumption, production and saving within a single period.[1] For subsequent time periods, we assume that there exists a given set of policy functions that link consumption and production to the value of period-0 savings, as well as the subsequent evolution of productivity. Thus consumption in period $t > 0$ is described by a policy function $c_t^* \left(m_0, \theta^t\right)$, where m_0 is savings in period 0, θ^t is the history of productivity draws from period 1 to t. The function c_t^* is treated as given in period 0. By varying the tax treatment of savings in period 0, policy can influence future allocations – but only within the restricted set of options that m_0 indexes.

Any incentive-compatible allocation can, by the revelation principle, be expressed through policy functions that map directly from individuals' type histories to their allocations, of the form $c_t \left(\theta_0, \theta^t\right)$ for consumption in period t, say. We exploit the fact that the dependence on initial type, θ_0, can always be replaced by dependence on an alternative proxy variable m_0 that is itself dependent on θ_0. This m_0 will affect future allocations in a manner that is equivalent to the effect of savings on future choice in market economies. So long as the stochastic evolution of productivity is Markovian, it is equivalent for allocations to depend directly on θ_0, or to depend indirectly on it, through a decentralised mechanism that induces individuals of different types to make different savings choices.

The advantage of formulating the problem in this way is that it allows efficiency require-

[1] This period is framed as the initial one, and denoted period 0, but this is without loss.

ments for tax design in period 0 to studied in isolation from future allocations. Our results show that clear principles should guide the appropriate relative tax treatment of labour income and savings at any moment in time, and these principles are independent of the way the savings variable may go on to influence future budget constraints. Future allocations may embed insurance optimally, imperfectly, or not at all: the results that we derive will continue to be valid.

Our main result requires very little structure to be placed on the social welfare criterion, beyond the Pareto principle. It follows from an arbitrage requirement across tax bases, holding constant the profile of utilities across individuals. The welfare criterion will certainly matter for whether labour income is taxed in the first place; but whether or not this is true, we show that marginal labour and savings tax wedges must take opposite signs under any tax system that is constrained Pareto-efficient.

The initial analysis uses an information-theoretic approach to optimal taxation, in the tradition of Mirrlees (1971). Given this, it is natural to wonder how our results relate to elasticity-based characterisations of optimal tax design, as in the extensive literature following Diamond and Saez (1998) and Saez (2001).

In this vein, we show that our main arbitrage condition linking earnings and savings taxes can be re-expressed by reference to two familiar classes of measurable statistics: compensated elasticities of earnings and savings with respect to their relative prices, and the localised Pareto parameters from the earnings and savings distributions.[2] The result states that it is Pareto efficient to equalise across tax bases the overall distortion induced by marginally higher taxes – as measured by the interaction between compensated elasticity responses and the Pareto statistic. Setting utility to be separable and types persistent is, we show, equivalent to restricting the elasticities and Pareto parameters to guarantee that the savings and earnings tax wedges take opposite signs.

This elasticity-based analysis is of significant independent interest, because it provides a direct and novel generalisation to a nonlinear tax environment of well-known Ramsey commodity tax results – notably the elasticity formula first derived by Samuelson (1986). It is not straightforward to reconcile the conclusions for commodity taxation of Ramsey analysis, which prioritises measurable elasticities, with information-theoretic analyses – which link critical results to properties of the utility function and type process. Our formulae provide this link for the setting we consider.[3] Crucially, the generalisation becomes possible only

[2]Both of these objects are familiar from the sufficient statistics tax literature. See, for instance, Piketty and Saez (2013).

[3]Atkinson and Stiglitz (1976) showed that uniform commodity taxation was optimal when labour supply

when the savings tax is measured by reference to the fiscal measure of the savings wedge that we introduce. This is intuitive: elasticity-based tax rules link optimal tax design to the behavioural effects of tax changes on tax revenue. By summarising the full revenue implications of a change in individuals' saving levels, our measure of the savings wedge allows this to be done.

The paper is organised as follows. In Section 2 we set out the general notation and framework. In Section 3 we derive our main result in a two-period setting where definitional complications relating to the savings wedge do not arise, because labour earnings in the second period are untaxed. In Section 4 we then extend our results to general settings where future margins are distorted, including situations where a dynamically optimal allocation is implemented. Section 5 explains in more detail how our results can be reconciled with a more familiar characterisation of the savings wedge that is based on the IEE. Section 6 then provides the re-statement of our results by reference to measurable statistics, and shows how assumptions on the utility function and type process imply consequential cross-restrictions on these statistics. Section 7 concludes.

2 General setup and problem statement

2.1 Primitives and definitions

We consider a discrete-time environment with time indexed by $t \in \{0, ..., \tau\}$. $\tau = \infty$ is possible. There is a unit measure of individuals, who differ cross-sectionally according to an idiosyncratic unobservable, random productivity parameter $\theta_t \in \left[\underline{\theta}, \bar{\theta}\right] \equiv \Theta \subset \mathbb{R}_+$ that follows a stochastic Markov process from one period to the next, with continuous conditional distribution $\Pi\left(\theta_t|\theta_{t-1}\right)$ and density $\pi\left(\theta_t|\theta_{t-1}\right)$ for $t > 0$, and $\Pi_0\left(\theta_0\right)$ and $\pi_0\left(\theta_0\right)$ for the initial draw. We assume $\underline{\theta} > 0$ and $\bar{\theta} < \infty$. $\pi\left(\theta_t|\theta_{t-1}\right)$ is assumed to be strictly positive and continuous in both arguments, and continuously differentiable in lagged type θ_{t-1}. It is helpful in what follows to distinguish between the shock draw in period 0 and the subsequent history of draws. Accordingly, $\theta^t \in \Theta^t$ will denote a history of productivities from θ_1 to θ_t inclusive, excluding the θ_0 draw.

The persistence of (log) productivity is summarised by the variable $\rho\left(\theta_t|\theta_{t-1}\right)$, defined as

was weakly separable from consumption goods in a consumer's utility function, and non-linear taxation was available. Kaplow (2024) provides a recent discussion of the perceived tensions between this result and the classic, elasticity-based Ramsey tax analysis that builds on Ramsey (1927) and Samuelson (1986).

follows:

$$\rho\left(\theta_t|\theta_{t-1}\right) := \frac{\theta_{t-1}\frac{d\left(1-\Pi\left(\theta_t|\theta_{t-1}\right)\right)}{d\theta_{t-1}}}{\theta_t \pi\left(\theta_t|\theta_{t-1}\right)} \tag{1}$$

It is easy to show:

$$\frac{d\mathbb{E}\left[\log\theta_t|\,\theta_{t-1}\right]}{d\log\theta_{t-1}} = \mathbb{E}\left[\rho\left(\theta_t|\theta_{t-1}\right)|\theta_{t-1}\right] \tag{2}$$

and so when ρ is constant, it is the autoregressive parameter for the effect of $\log\theta_{t-1}$ on $\log\theta_t$. We assume $\rho\left(\theta_t|\theta_{t-1}\right) \in [0,1)$ throughout. Types are iid when ρ is identically equal to zero, and persistent when it is strictly positive for all (θ_{t-1},θ_t) pairs. We restrict attention to these two cases.

We also define the pass-through rates from θ_t to θ_s as the compounded products of successive $\rho\left(\theta_t|\theta_{t-1}\right)$ values. For all $s > t \geq 0$, we set $D_{t,t}\left(\theta_0,\theta^t\right) = 1$ and define:

$$D_{t,s}\left(\theta_0,\theta^s\right) := D_{t,s-1}\left(\theta_0,\theta^{s-1}\right)\rho\left(\theta_s|\theta_{s-1}\right) = \prod_{r=t+1}^{s}\rho\left(\theta_r|\theta_{r-1}\right). \tag{3}$$

Individuals have preferences across contingent sequences of consumption and labour earnings, conditional on their productivity type. In period 0, these are represented by the lifetime utility function:

$$U_0 := \mathbb{E}\left[\sum_{t=0}^{\tau}\beta^t u\left(c_t,y_t;\theta_t\right)\middle|\,\theta_0\right] \tag{4}$$

where $\beta \in (0,1)$ is a discount factor, u is a period utility function and c_t and y_t represent period-t consumption and earnings respectively, both measurable with respect to the individual's history of type draws up to t. For notational convenience, we drop the dependence on c_t, y_t, and θ_t whenever this does not cause ambiguity. We restrict attention to contingent sequences where the utility function is well-defined for all types. There is no source of randomness aside from productivity. Preferences over continuation sequences can be defined from arbitrary period t, in the usual way.

We make standard differentiability and regularity assumptions. Using subscripts to denote partial derivatives, we assume $u_c > 0$, $u_y < 0$, and strict concavity of u in (c,y) given θ. In addition, $u_\theta > 0$ and $u_{y\theta} > 0$: being more productive is beneficial and reduces the marginal disutility of an extra unit of earnings. The derivative $u_{c\theta}$ also exists and may take any sign, but is restricted by single crossing in its Spence-Mirrlees form: an increase in θ must reduce the marginal compensation that individuals require to produce an extra unit of

output at all (c, y, θ). Here, this implies:

$$\frac{u_{c\theta}}{u_c} - \frac{u_{y\theta}}{u_y} > 0 \tag{5}$$

for all (c, y, θ).

An **allocation**, generically $X_0 \in \mathcal{X}_0$, is a set of mappings from histories of types $\left(\theta_0, \theta^t\right) \in \Theta^{t+1}$ to consumption and production levels, $c_t\left(\theta_0, \theta^t\right)$ and $y_t\left(\theta_0, \theta^t\right)$, for all $t \geq 0$. $\mathcal{X}_0$ is the space of all such allocations. Since productivity is private information, an incentive compatibility constraint will restrict the admissible set of allocations across types. Given a history, denote by $U_t\left(\theta_t'; \theta_t\right)$ the value from t onwards to type θ_t of receiving the allocation 'intended' for type θ_t'.[4] Incentive compatibility requires:

$$U_t\left(\theta_t; \theta_t\right) \geq U_t\left(\theta_t'; \theta_t\right) \tag{6}$$

for all $\left(\theta_t', \theta_t\right) \in \Theta^2$, at all history nodes and in all time periods.[5] An allocation that satisfies this requirement is called an *incentive-compatible allocation*.

Resources can be saved from one period to the next by the policymaker at a fixed real interest rate $R > 1$. Thus any allocation has a present-value **resource cost**, given by the following expression:

$$\mathbb{E}_{-1}\left[\sum_{t=0}^{\tau} R^{-t}\left(c_t - y_t\right)\right] \tag{7}$$

where expectations $\mathbb{E}_{-1}[\cdot]$ are taken over the population distribution of the entire sequence of shock histories $\left(\theta_0, \theta^t\right)$ from period 0 on.

Our analysis centres on a standard notion of constrained Pareto efficiency:

Definition. An allocation is **constrained Pareto efficient** if there is no other incentive-compatible allocation that implies the same U_0 for all period-0 types in Θ, and a strictly lower resource cost.

[4] That is, the allocation $\left(c_t\left(\theta_0, \theta^{t-1}, \theta_t'\right), y_t\left(\theta_0, \theta^{t-1}, \theta_t'\right)\right)$ in period t, followed by $\left(c_{t+1}\left(\theta_0, \theta^{t-1}, \theta_t', \theta_{t+1}\right), y_{t+1}\left(\theta_0, \theta^{t-1}, \theta_t', \theta_{t+1}\right)\right)$ in period $t+1$, where θ_{t+1} is that period's true type draw, and so on for subsequent periods.

[5] We will use $U_t\left(\theta_t\right)$ in place of $U_t\left(\theta_t; \theta_t\right)$ for cases where actual and reported type are identical.

2.2 Wealth as a variable

Our aim is to understand the joint properties of saving and income taxation when allocations are constrained Pareto efficient. The main complication for this is that neither saving nor its close counterpart, wealth, is a primitive variable. This makes it difficult to incorporate either into a mechanism design problem that is defined on the space of direct allocations – i.e., history-contingent consumption and production levels.

To resolve this we identify an individual's wealth at the end of period 0 with their continuation allocation. A **continuation allocation**, generically $X_1 \in \mathcal{X}_1$, is defined near-identically to an allocation, but restricted to time period 1 onwards. Thus it is a set of mappings from histories of types $\theta^t \in \Theta^t$ to consumption and production levels, $c_t\left(\theta^t\right)$ and $y_t\left(\theta^t\right)$, for all $t \geq 1$. $\mathcal{X}_1$ is the space of all continuation allocations. Again, we restrict attention to incentive-compatible continuation allocations.

An **indexed set of continuation allocations** is a combination of an indexing variable $m_0 \in I \subset \mathbb{R}$, with I an interval, and a mapping $X_1^* : I \to \mathcal{X}_1$, so that $X_1^*\left(m_0\right)$ defines a continuation allocation for each $m_0 \in I$. We denote policy functions from an indexed set with asterisks, so that $c_t^*\left(m_0', \theta^t\right)$ is the consumption level at node θ^t when the index variable takes value m_0', and so on.

Any incentive-compatible allocation X_0 embeds an indexed set of continuation allocations $X_1^*\left(m_0\right)$, namely the set of continuation allocations associated with each initial type, θ_0, and an assignment function $m_0 : \Theta \to I$ that maps each type to their continuation allocation, i.e. $X_1\left(\theta_0\right) = X_1^*\left(m_0\left(\theta_0\right)\right)$. Conversely, given any indexed set of incentive-compatible continuation allocations $X_1^*\left(\cdot\right)$, a joint assignment $c_0\left(\theta_0\right)$, $y_0\left(\theta_0\right)$ and $m_0\left(\theta_0\right)$ for all $\theta_0 \in \Theta$ generates an allocation $X \in \mathcal{X}_0$, with $c_t\left(\theta_0, \theta^t\right) = c_t^*\left(m_0\left(\theta_0\right), \theta^t\right)$ and $y_t\left(\theta_0, \theta^t\right) = y_t^*\left(m_0\left(\theta_0\right), \theta^t\right)$, for all $\left(\theta_0, \theta^t\right)$. The efficiency of an original allocation can, therefore, be explored by considering alternative $c_0\left(\theta_0\right)$, $y_0\left(\theta_0\right)$ and $m_0\left(\theta_0\right)$ assignments, given the mapping $X_1^*\left(m_0\right)$ implied by the original allocation.

The index variable m_0 here provides a single degree of freedom over future consumption and production allocations. Provided $X_1^*\left(m_0\right)$ respects certain regularity conditions – so that an increase in m_0 meaningfully corresponds to an improvement in an individual's future allocation – this allows m_0 to be interpreted as an individual's wealth level. For instance, this will make sense if $X_1^*\left(m_0\right)$ is such that $c_t^*\left(m_0, \theta^t\right)$ is increasing in m_0 and $y_t^*\left(m_0, \theta^t\right)$ decreasing in m_0 at all history nodes. This, in turn, will follow when $X_1^*\left(m_0\right)$ was constructed from an original allocation where $c_t\left(\theta_0, \theta^t\right)$ and $y_t\left(\theta_0, \theta^t\right)$ were correspondingly monotone in θ_0.

Thus we treat m_0 as a composite good, chosen in period 0, and explore requirements for its efficient tax treatment. Since continuation allocations are limited to a one-dimensional set by construction, this analysis will not be *sufficient* for efficiency. Insofar as we are nonetheless able to obtain non-trivial policy prescriptions, this limited scope is an advantage: focusing on constrained Pareto efficiency in period 0, the results we obtain apply regardless of how future policy is to be designed, with the unrestricted second-best policy of the dynamic Mirrleesian economy a special case of our analysis.

Though not essential, it helps the exposition to assume $X_1^*(m_0)$ is such that the policy functions $c_t^*(m_0, \theta^t)$ and $y_t^*(m_0, \theta^t)$ are differentiable in m_0 for each θ^t and t pair. Where this is assumed, the corresponding derivatives are denoted by a subscript: $c_{m,t}^*(m_0, \theta^t)$ and $y_{m,t}^*(m_0, \theta^t)$ respectively. For convenience we suppress the arguments (m_0, θ^t) where this does not cause ambiguity, writing c_t^* in place of $c_t^*(m_0, \theta^t)$, $c_{m,t}^*$ in place of $c_{m,t}^*(m_0, \theta^t)$, and so on.

2.3 Relaxed policy problem

In common with many papers in the dynamic Mirrlees literature, we simplify incentive compatibility through a first-order approach. Central to this is the notion of an **information rent**, which captures the marginal effect of an increase in an individual's period-t productivity on their utility, fixing a given state-contingent allocation from t onwards. For an individual whose true type is θ_t receiving the allocation intended for θ_t', this information rent is given by:

$$U_{\theta,t}(\theta_t'; \theta_t) := \frac{1}{\theta_t} \mathbb{E}\left[\sum_{s=t}^{\tau} \beta^{s-t} D_{t,s} \theta_s u_\theta(c_s, y_s; \theta_s) \middle| \theta_t \right] \tag{8}$$

where the values for c_s and y_s are those intended for the period-t type θ_t', assuming subsequent truthful type reporting. As for $U_t(\theta_t)$, we will write $U_{\theta,t}(\theta_t)$ where the allocation is the one intended for the given true type. Notice that the information rent is non-negative.

Technically, it will be convenient to focus on allocations that have bounded information rents at each history node, defined as follows.

Definition. An allocation has **bounded information rents** at each history node if, for all (θ_0, θ^{t-1}), there exists a function $b(\cdot; \theta_0, \theta^{t-1}) : \Theta \to \mathbb{R}$, integrable on Θ, such that $|U_{\theta,t}(\theta_t'; \theta_t)| \leq b(\theta_t; \theta_0, \theta^{t-1})$ for all $\theta_t' \in \Theta$ at almost all $\theta_t \in \Theta$.

Given this, the following envelope result is established in the appendix. At its core is an

application of Theorem 2 in Milgrom and Segal (2002).

Proposition 1. *An incentive-compatible allocation with bounded information rents must satisfy the following restriction for all $\theta_t' \in \Theta$, at all history nodes $\left(\theta_0, \theta^{t-1}\right) \in \Theta^t$:*

$$U_t\left(\theta_t'\right) = U_t\left(\underline{\theta}\right) + \int_{\underline{\theta}}^{\theta_t'} U_{\theta,t}\left(\theta_t\right) d\theta_t \tag{9}$$

Constraint (9) implies a single restriction for each $\theta_t' \in \Theta$, and so is a dimension simpler than (6) – which provides a restriction for every *pair* of types in Θ^2. It is for this reason that analysis typically uses (9) as a substitute for (6), the so-called 'first-order approach' to incentive compatibility. The complication is that, whilst (9) is a necessary requirement for (6) to hold, it is not sufficient. Non-local incentive constraints may be violated, and additional ex-post verification strategies are required to confirm that an allocation satisfying (9) also satisfies (6) for all pairs.

We consider perturbations around a benchmark incentive-compatible allocation whose Pareto efficiency is being assessed. These perturbations are constructed to hold lifetime utility constant for all types in period 0. By (9), if lifetime utility is constant then information rents must additionally be held constant constant for (almost) all types, so as to preserve incentive compatibility.

Formally, suppose that a triplet $\{c_0\left(\theta_0\right), y_0\left(\theta_0\right), m_0\left(\theta_0\right)\}$ is Pareto efficient. Then, it must be that for given lifetime utility value and information rents $U_0\left(\theta_0\right)$ and $U_{\theta,0}\left(\theta_0\right)$ defined by $\{c_0\left(\theta_0\right), y_0\left(\theta_0\right), m_0\left(\theta_0\right)\}$ according to (9), $\{c_0\left(\theta_0\right), y_0\left(\theta_0\right), m_0\left(\theta_0\right)\}$ solves the following cost minimization problem:

Problem 1.

$$\min_{\substack{\tilde{c}_0 : \Theta \to \mathbb{R}_+ \\ \tilde{y}_0 : \Theta \to \mathbb{R}_+ \\ \tilde{m}_0 : \Theta \to I}} \mathbb{E}_{-1}\left\{\tilde{c}_0\left(\theta_0\right) - \tilde{y}_0\left(\theta_0\right) + \sum_{t=1}^{\tau} R^{-t}\left[c_t^*\left(\tilde{m}_0\left(\theta_0\right), \theta^t\right) - y_t^*\left(\tilde{m}_0\left(\theta_0\right), \theta^t\right)\right]\right\} \tag{10}$$

subject to:

$$\begin{aligned} U_0\left(\theta_0\right) = & u\left(\tilde{c}_0\left(\theta_0\right), \tilde{y}_0\left(\theta_0\right); \theta_0\right) \\ & + \mathbb{E}\left[\sum_{t=1}^{\tau} \beta^t u\left(c_t^*\left(\tilde{m}_0\left(\theta_0\right), \theta^t\right), y_t^*\left(\tilde{m}_0\left(\theta_0\right), \theta^t\right); \theta_t\right) \middle| \theta_0\right] \end{aligned} \tag{11}$$

for all $\theta_0 \in \Theta$, and:

$$\begin{aligned}\theta_0 U_{\theta,0}(\theta_0) =& \theta_0 u_\theta\left(\tilde{c}_0(\theta_0), \tilde{y}_0(\theta_0); \theta_0\right) \\ &+ \mathbb{E}\left[\sum_{t=1}^{\tau} \beta^t D_{0,t} \theta_t u_\theta\left(c_t^*\left(\tilde{m}_0(\theta_0), \theta^t\right), y_t^*\left(\tilde{m}_0(\theta_0), \theta^t\right); \theta_t\right) \middle| \theta_0\right]\end{aligned} \tag{12}$$

for all $\theta_0 \in \Theta$, given $U_0(\theta_0)$ and $U_{\theta,0}(\theta_0)$ satisfying (9) for all θ_0, and given policy functions c_t^* and y_t^* for all θ^t and t.

Using a perturbation approach, we derive necessary conditions for an allocation to be Pareto efficient, and then translate these into optimal distortions, or tax wedges, on income and savings. The precise definition of these distortions is critical to our analysis, particularly in the case of the savings wedge. We turn to this next.

2.4 Defining wedges

We focus on the joint, qualitative implications of Pareto efficiency for the earnings and savings wedges in period 0. Both of these wedges are defined by reference to a period-0 allocation $\{c_0(\theta_0), y_0(\theta_0), m_0(\theta_0)\}$, given an indexed set of continuation allocations $X^*(m_0)$.

The **earnings wedge** on type θ_0 is denoted $T'(y_0(\theta_0))$, or simply $T'(\theta_0)$ where convenient. It is defined as the difference between the real product of work for θ_0 – normalised to 1 – and the marginal compensation that θ_0 requires to generate this extra unit of output. That is:[6]

$$T'(\theta_0) = 1 + \frac{u_{y,0}}{u_{c,0}} \tag{13}$$

where the marginal rate of substitution on the right-hand side is defined for type θ_0, at θ_0's allocation. Note that $T'(y_0(\theta_0))$ corresponds to the marginal social gain, when type θ_0 produces an extra unit of output, and is just compensated for doing so.

The savings wedge for type θ_0 is a critical object in our analysis. Its definition is more complicated than the earnings wedge, because different assumptions about how savings are to be spent will generate different expressions for the wedge. In distorted economies, there is no reason for these different expressions to be equivalent.

The optimal taxation literature typically focuses on a one-period **private savings wedg**e, defined as the deviation of an agent's expected inter-temporal marginal rate of substitution

[6]Recall $u_y < 0$.

in consumption from the technological return on savings, R:

$$s'(\theta_0) = \mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] - 1 \tag{14}$$

This private savings wedge measures the average distortion that an individual faces on a decision between allocating consumption to period 0 and allocating consumption to period 1, relative to a risk-free instrument that pays the uniform return R. It is of particular relevance when the inverse Euler equation applies, since this directly implies that $s'(\theta_0)$ must be positive.

However, $s'(\theta_0)$ will not summarise the full *public finance* implications of an individual's decision to save more or less at the margin. This is because incremental savings will not, in equilibrium, be used exclusively for consumption in period 1. Additional leisure and consumption at arbitrary horizons will also be taken. This matters if there are relative tax distortions between the period-1 consumption good and these other alternatives. The decision to save will have direct implications for taxes raised on these margins, and efficient policy design will need to take these effects into account.

This motivates what we define as the **public savings wedge**: the markup of the implicit price to an individual of an extra unit of m_0, over the real marginal cost of providing it. This corresponds to the textbook '$\frac{\mathcal{P}-\mathcal{C}}{\mathcal{C}}$' definition of a wedge, where $\mathcal{P}$ is price and $\mathcal{C}$ marginal cost. But crucially, $\mathcal{P}$ and $\mathcal{C}$ are defined by reference to the changed pattern of consumption and production, across all future periods, that m_0 induces.

For simplicity, we focus the discussion on the case where the functions c_t^* and y_t^* are differentiable in m_0 along all (or Π-almost all) history paths. Given this, the consumer's marginal rate of substitution between c_0 and m_0 can be written as: $\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}c_{m,t}^* + \frac{u_{y,t}}{u_{c,0}}y_{m,t}^*\right]\right]$. In any market equilibrium, this must correspond to the relative price of m_0 in c_0 numeraire, and so $\mathcal{P} = \mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}c_{m,t}^* + \frac{u_{y,t}}{u_{c,0}}y_{m,t}^*\right]\right]$.

The marginal resource *cost* of raising m_0 is the discounted value of the change in consumption, less the change in production that this increase implies. This can be written as $\mathcal{C} = \mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}\left[c_{m,t}^* - y_{m,t}^*\right]\right]$. Consequently, we define the public savings wedge $S'(\theta_0)$

as:

$$S'(\theta_0) = \frac{\overbrace{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}c^*_{m,t} + \frac{u_{y,t}}{u_{c,0}}y^*_{m,t}\right]\right]}^{\mathcal{P}} - \overbrace{\mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}\left[c^*_{m,t} - y^*_{m,t}\right]\right]}^{\mathcal{C}}}{\underbrace{\mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}\left[c^*_{m,t} - y^*_{m,t}\right]\right]}_{\mathcal{C}}} \tag{15}$$

When $S'(\theta_0) > 0$, the marginal value to consumers of an extra unit of m_0 exceeds the marginal cost *of the composite goods these savings are used to purchase*. The consumer will pay more to increase m_0 than the real cost of providing this increase. The difference is a *de facto* tax wedge. When $S'(\theta_0) < 0$, every increase in private savings instead costs the planner more in period-0 numeraire than individual consumers are willing to sacrifice. Savings are costly to the fiscal authority at the margin – i.e., they are subsidised.

The point is not that we believe S' to be any more or less 'correct' than s' as a measure of the savings wedge. Both can be instructive about the nature of distortions. What matters is what is useful. As we show, when savings taxes are analysed by reference to the public savings wedge, the analysis becomes much clearer. When it is used, the optimal tax analysis of savings in period 0 parallels that of any other consumer good. This makes it important to understand the link between the different measures of the wedge. We do this through a series of brief examples.

Example 1. Suppose two time periods ($\tau = 1$), no period-1 distortions ($-\frac{u_{y,1}}{u_{c,1}} = 1$), uniform returns to saving ($c^*_{m,1} - y^*_{m,1} \equiv \Delta > 0$, independent of θ_1). In this case:

$$S'(\theta_0) = \frac{\mathbb{E}_0\left[\frac{\beta u_{c,1}}{u_{c,0}}\Delta\right] - R^{-1}\Delta}{R^{-1}\Delta} = \mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] - 1 = s'(\theta_0) \tag{16}$$

Thus our wedge coincides with the departure from the standard Euler equation on risk-free one-period bonds. This result extends easily to an arbitrary, finite number of time periods, so long as the consumption-production margin remains undistorted in all $t > 0$, and savings between t and $t+1$ earn a uniform, undistorted gross return equal to R for all $t \in \{1, ..., \tau - 1\}$. This example is developed extensively in Section 3.

Example 2. Two time periods ($\tau = 1$), uniform period-1 labour distortion ($-\frac{u_{y,1}}{u_{c,1}} \equiv 1 - \mu$), uniform returns to saving, denominated in consumption good ($c^*_{m,1} - y^*_{m,1}(1-\mu) \equiv \Delta > 0$.

In this case:

$$S'(\theta_0) = \frac{1}{1-\mu\frac{\mathbb{E}_0\left[y^*_{m,1}\right]}{\Delta}}\mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] - 1 \tag{17}$$

$$= \frac{1}{1-\mu\frac{\mathbb{E}_0\left[c^*_{m,1}\right]}{\Delta}}\mathbb{E}_0\left[-\frac{\beta R u_{y,1}}{u_{c,0}}\right] - 1 \tag{18}$$

$\frac{\mathbb{E}_0\left[y^*_{m,1}\right]}{\Delta}$ and $\frac{\mathbb{E}_0\left[c^*_{m,1}\right]}{\Delta}$ are the average income effects on production and consumption, respectively, per unit of additional resources the consumer holds in period 1. So long as leisure is a normal good, $y^*_{m,1} < 0$, and so – by equation (17) – the public savings wedge differs from the private to the extent that higher savings reduce production, and lower production is more costly than higher consumption (μ).

There is no analytical reason to privilege the departure from the *consumption* Euler condition over alternatives. In an undistorted economy, marginal rates of substitution are equated with rates of transformation across all margins. Equation (18) illustrates how S' differs from the wedge associated with a decision to use savings for *leisure* alone – i.e., the gap between the value to the consumer of a uniform reduction in production in period 1, and its real cost. If $\mu \neq 0$ and leisure and consumption are both strictly normal goods,[7] this wedge and the wedge on the consumption Euler equation lie either side of our measure. Assuming $\mu > 0$:[8]

$$\mathbb{E}_0\left[-\frac{\beta R u_{y,1}}{u_{c,0}}\right] - 1 < S'(\theta_0) < \mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] - 1 = s'(\theta_0) \tag{19}$$

Finally on this example, notice that S' would equal zero if the consumer in period 0 could save at a risk-free rate of return of $\frac{R}{1-\mu\frac{\mathbb{E}_0\left[y^*_{m,1}\right]}{\Delta}}$, denominated in consumption good. At this rate of return, the option to reduce consumption in period 0 by a unit, and raise consumption and leisure in period 1 in the amounts given by the period-1 policy functions, would be a choice that is available to consumers at an undistorted price.

Example 3. Three time periods ($\tau = 2$), no labour distortions in periods 1 or 2 ($-\frac{u_{y,1}}{u_{c,1}} = -\frac{u_{y,2}}{u_{c,2}} = 1$), uniform returns to saving equal to $\frac{1}{1+\mu}R$ between periods 1 and 2,[9] uniform

[7] That is, $y^*_{m,1} < 0$ and $c^*_{m,1} > 0$.

[8] For $\mu < 0$ the inequalities are reversed.

[9] That is, $\mathbb{E}_1\left[\frac{\beta u_{c,2}}{u_{c,1}}\right] = (1+\mu)R^{-1}$.

returns to saving equal to Δ between periods 0 and 1.[10] In this case:

$$S'(\theta_0) = \frac{1}{1 - \frac{\mu}{1+\mu}\frac{R^{-1}(1+\mu)\mathbb{E}_0\left[c^*_{m,2} - y^*_{m,2}\right]}{\Delta}} \mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] - 1 \tag{20}$$

Again, S' differs from the wedge on the classic Euler condition to the extent that savings will be used to purchase goods that are further distorted. $\frac{R^{-1}(1+\mu)\mathbb{E}_0\left[c^*_{m,2} - y^*_{m,2}\right]}{\Delta}$ measures the fraction of the returns to saving in period 1 that are deferred to period 2, and $\frac{\mu}{1+\mu}$ captures the wedge on this subsequent margin. If $\mu > 0$, so that additional savings between periods 1 and 2 are taxed, then $S'(\theta_0)$ will generally exceed $s'(\theta_0)$. This is because additional savings in period 0 imply additional savings in period 1, which are also to be taxed.

Example 4. Two time periods ($\tau = 1$), no period-1 distortions ($-\frac{u_{y,1}}{u_{c,1}} = 1$), returns to saving measurable with respect to θ_1 ($c^*_{m,1} - y^*_{m,1} := \Delta(\theta_1)$). In this case:

$$S'(\theta_0) = \mathbb{E}_0\left[\frac{\beta R u_{c,1}}{u_{c,0}}\right] + \text{Cov}\left[\frac{\beta R u_{c,1}}{u_{c,0}}, \frac{\Delta(\theta_1)}{\mathbb{E}_0\left[\Delta(\theta_1)\right]}\right] - 1 \tag{21}$$

With nonlinear tax systems it is possible for the effective return on savings to vary in period-1 type. This will be true, for instance, if the period-1 income tax schedule is wealth-contingent, so that the realised post-tax returns to saving vary in period-1 labour income. In these settings, S' allows for the fact that the consumer will value savings more highly when returns correlate positively with their marginal utility, and an increase in $\mathcal{P}$ that merely reflects this should not be considered part of a wedge.

Given the correlation structure of the savings instrument, the measure S' indicates that savings are not distorted when the right-hand side of (21) equals zero. This is appropriate, as the right-hand side would also equal zero at an undistorted complete-markets allocation. However, it is immediate from (21) that the Euler condition associated with risk-free savings instruments *will* exhibit a 'wedge' in this case, equal in absolute magnitude to the covariance term. The wedge on the classic consumption Euler condition would, therefore, give a misleading impression of distortions.

To summarise, so long as future allocations are distorted, individual indifference does not imply social indifference. The marginal cost of raising m_0 will generally differ across alternatives that hold the marginal utility from it constant. Any measure of the savings wedge must,

[10] That is, $\left[\left(c^*_{m,1} - y^*_{m,1}\right) + (1+\mu)R^{-1}\left(c^*_{m,2} - y^*_{m,2}\right)\right] \equiv \Delta$.

therefore, make some assumptions about the end use of savings. Our definition of $S'(\theta_0)$ assumes savings are to be used in the way they would actually be used in equilibrium. As we show, this significantly clarifies the factors that matter for the optimal tax treatment of savings relative to consumption.

3 Two-period example

We first develop the intuition for our main result in a simplified two-period setting. The environment is deliberately constructed to be without distortions in period 1, so that the savings wedge has a unique, unambiguous definition.

3.1 Assumptions and main result

There are two time periods: period 0 and period 1. The within-period utility function is assumed to take the isoelastic separable form:

$$u(c, y; \theta) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \frac{1}{1+\nu}\left(\frac{y}{\theta}\right)^{1+\nu} \tag{22}$$

for $\sigma > 0$ and $\nu > 0$. This implies:

$$\theta u_\theta(c, y; \theta) = \left(\frac{y}{\theta}\right)^{1+\nu} \tag{23}$$

In period 1, we assume that there is no insurance of any kind, with consumption and production chosen optimally subject to a linear budget constraint of the form:

$$c_1 = y_1 + Rm_0 \tag{24}$$

where Rm_0 corresponds to the individual's period-1 wealth. Note that this implies the present-value cost of providing the period-1 allocation, viewed in period 0, is m_0, which is independent of ex-post choice, or the distribution over period-1 types.

The period-1 choice problem can be described by policy functions $c_1^*(m_0, \theta_1)$ and $y_1^*(m_0, \theta_1)$,

which are the unique solutions to the pair of equations:

$$c_1^*(m_0, \theta_1) = y_1^*(m_0, \theta_1) + Rm_0 \tag{25}$$

$$c_1^*(m_0, \theta_1)^{-\sigma} = \frac{1}{\theta_1}\left(\frac{y_1^*(m_0, \theta_1)}{\theta_1}\right)^{\nu} \tag{26}$$

It is easy to show that c_1^* and y_1^* are both strictly positive and differentiable in m_0, given θ_1. Their partial derivatives (denoted by subscripts) satisfy $c_{m,1}^* > 0$ and $y_{m,1}^* < 0$ respectively, with $y_{m,1}^* \to 0$ as $\sigma \to 0$.

Since θ_1 is stochastic, the absence of insurance in period 1 will not generally be *ex-ante* Pareto efficient. We leave this future inefficiency aside, and simply ask whether Pareto improvements are possible in period 0 by varying period-0 consumption, production, and retained wealth.

Period-0 allocations are described by functions $c_0(\theta_0)$, $y_0(\theta_0)$ and $m_0(\theta_0)$, with $m_0(\theta_0)$ inducing an allocation in period 1 via the policy functions $c_1^*(m_0(\theta_0), \theta_1)$ and $y_1^*(m_0(\theta_0), \theta_1)$. The nonlinear tax functions $T(y_0)$ and $S(m_0)$ give a quantity of the period-0 consumption good to be paid for each choice of y_0 and m_0 respectively. The consumer's budget constraint in period 0 is thus:

$$y_0 \geq c_0 + m_0 + T(y_0) + S(m_0) \tag{27}$$

And, at an interior optimum, first-order conditions require:

$$1 - T'(y_0(\theta_0)) = \frac{\frac{1}{\theta_0}\left(\frac{y_0(\theta_0)}{\theta_0}\right)^{\nu}}{c_0(\theta_0)^{-\sigma}} \tag{28}$$

$$1 + S'(m_0(\theta_0)) = \frac{\beta R \mathbb{E}\left[c_1^*(m_0(\theta_0), \theta_1)^{-\sigma} | \theta_0\right]}{c_0(\theta_0)^{-\sigma}} \tag{29}$$

Notice that condition (29) here is a special case of (15), making use of (25) and (26).

We derive the following result for this setting.

Proposition 2. *(a) Suppose productivity is iid. Then at any constrained Pareto-efficient allocation, $S'(\theta_0) = 0$ everywhere.*

(b) Suppose productivity is persistent. Then at any constrained Pareto-efficient allocation, $S'(\theta_0)$ takes the opposite sign to $T'(\theta_0)$ for all $\theta_0 \in \Theta$.

Proof outline

The Proposition follows directly from the following arbitrage condition, established in the appendix, which is a requirement for Pareto efficiency:

$$\frac{T'(\theta_0)}{\frac{1}{\theta_0}\left(\frac{y_0}{\theta_0}\right)^{\nu}} = \frac{S'(\theta_0)}{\beta\mathbb{E}\left[\rho(\theta_1|\theta_0)\cdot\frac{y^*_{m,1}}{\theta_1}\left(\frac{y^*_1}{\theta_1}\right)^{\nu}\middle|\,\theta_0\right]} \tag{30}$$

The marginal tax rates in the numerators here represent the resource surpluses that are generated when individuals produce an extra unit of output (T') or raise their period-1 wealth by a unit (S'), whilst being left indifferent. The denominators in the two main fractions represent the effect on period-0 information rents of the corresponding changes in production (left-hand side) and saving (right-hand side).

The economics behind the condition is as follows. An increase in $y_0(\theta_0)$ by $\varepsilon > 0$ generates additional marginal tax revenue of $T'(\theta_0)\varepsilon$ – and so is desirable per se, to the extent that this value is positive. But it also implies a change to information rents, by an amount $\frac{1}{\theta_0}\left(\frac{y_0}{\theta_0}\right)^{\nu}\varepsilon$ – implying higher utility for more productive agents and reducing the scope for redistribution from types $\theta' > \theta_0$ to types $\theta' < \theta_0$.

Likewise, an increase in $m_0(\theta_0)$ by $\eta > 0$ increases tax revenue by $S'(\theta_0)\eta$ but changes information rents by $\beta\mathbb{E}\left[\rho(\theta_1|\theta_0)\cdot\frac{y^*_{m,1}}{\theta_1}\left(\frac{y^*_1}{\theta_1}\right)^{\nu}\middle|\,\theta_0\right]\eta$. Hence by setting

$$\eta = -\frac{\frac{1}{\theta_0}\left(\frac{y_0}{\theta_0}\right)^{\nu}}{\beta\mathbb{E}\left[\rho(\theta_1|\theta_0)\cdot\frac{y^*_{m,1}}{\theta_1}\left(\frac{y^*_1}{\theta_1}\right)^{\nu}\middle|\,\theta_0\right]}\varepsilon,$$

the combined increase in earnings y_0 and reduction in savings m_0 leaves information rents at θ_0 unchanged but increases tax revenues if:

$$\left\{T'(\theta_0) - \frac{\frac{1}{\theta_0}\left(\frac{y_0}{\theta_0}\right)^{\nu}}{\beta\mathbb{E}\left[\rho(\theta_1|\theta_0)\cdot\frac{y^*_{m,1}}{\theta_1}\left(\frac{y^*_1}{\theta_1}\right)^{\nu}\middle|\,\theta_0\right]}S'(\theta_0)\right\}\varepsilon > 0 \tag{31}$$

A symmetric movement in the opposite direction – reducing y_0 and increasing m_0 – will generate a surplus if the bracket on the right-hand side of equation (31) is, instead, strictly negative. Thus, at an optimum, (30) must hold.

Proposition 2 then follows by inspecting the denominator on the right-hand side of (30). The iid case is simplest. It arises when $\rho(\theta_1|\theta_0)$ is zero everywhere, so that the equation can

only hold with $S'(\theta_0) = 0$. When types are persistent, $\rho(\theta_1|\theta_0) > 0$. In this case all objects in the denominator on the right-hand side are strictly positive except $y^*_{m,1}$ – the marginal effect of an increase in m_0 on y^*_1. This is strictly negative, and the results follow.

3.2 Discussion

Part (a) of Proposition 2 is essentially an application of the Atkinson-Stiglitz theorem. When types are iid, an individual's relative preferences between savings and consumption in period 0 are not influenced by their current type. As is well known, this implies that there are no gains from distorting the consumption-savings margin, relative to what can be achieved by an income tax alone.

Part (b) of Proposition 2 is the more significant result, in light of a literature that generally suggests savings should be *taxed* when insurance is incomplete. So long as the marginal income tax is positive, we have obtained precisely the opposite conclusion. To understand the result, we start with the role of normality in leisure. The following helps account for its presence:

Lemma 1. *The following are equivalent:*

1. $y^*_{m,1}(m_0, \theta_1) < 0$ *for all* $\theta_1 \in \Theta$.
2. $c^*_1(m_0, \theta_1)^{-\sigma}$ *is strictly decreasing in* θ_1*, given* m_0.

In words, normality of leisure is equivalent to diminishing marginal utility of income in period-1 type. This is significant in the context of type persistence, because persistence means that lower types in period 0 will rationally expect lower type draws in period 1 – and, thus, a higher marginal utility from their savings. Given separability, the utility from c_0 is type-independent in period 0. This implies that lower types will have a greater relative incentive to save out of any *given* level of post-tax income in period 0.

The type-independence of preferences that gives rise to the Atkinson-Stiglitz result is, therefore, broken: savings are a behavioural complement to work, since those who work longer hours for the same disposable income are, by definition, of lower type. Subsidising savings raises the incentives to work for those who place a high marginal value on savings, and this is desirable so long as a positive marginal income tax is keeping hours inefficiently low. But because higher types do not value savings by so much, this subsidy has less of an effect on higher types' welfare than a direct cut in the labour income tax. This is consistent with the general prescription, dating at least to Corlett and Hague (1953), that complements to work should be subsidised in distorted economies.

4 General setting

This section shows that the economic content of Proposition 2 generalises well beyond the stylised example of Section 3. Three of the simplifying assumptions made in Section 3 are relaxed: first, we consider settings with an arbitrary number of time periods; second, we allow for insurance to exist after period 0; third, we let the within-period utility function take an arbitrary non-separable form, subject to satisfying the Spence-Mirrlees condition. Though these complicate, and – in the case of non-separability – qualify the analysis, we show that the analytical case for subsidising self-insurance remains as a relevant consideration throughout.

In Section 3 we derived the policy functions $c_1^*(m_0, \theta_1)$ and $y_1^*(m_0, \theta_1)$ from an assumption that allocations were generated by optimal within-period choice under autarky, given m_0. For the remainder of the paper we dispense with explicit modelling of later time periods, and – as previewed in Section 2 – simply assume that there exists a fixed set of possible continuation allocations $X^*(m_0)$, parameterised by wealth variable $m_0 \in I \subseteq \mathbb{R}$. These are associated with policy functions $c_t^*(m_0, \theta^t)$ and $y_t^*(m_0, \theta^t)$ for every θ^t and $t > 0$. The main content of this assumption is that the institutional framework in which individual choice takes place after period 1 is determined independently of period-0 policy design, and so varying insurance in the initial period does not affect its future character.

Whilst not essential, it is convenient to place regularity assumptions on c_t^* and y_t^*, so as to take derivatives in the usual way. The assumptions we make are summarised as follows:

Assumption 1. *There exist given policy functions $c_t^*(m_0, \theta^t)$ and $y_t^*(m_0, \theta^t)$ describing consumption and production choices at each node $\theta^t \in \Theta$, for each wealth level m_0. These satisfy the following pair of regularity assumptions:*

(a) $c_t^*(m_0, \theta^t)$ is weakly increasing in m_0 and $y_t^*(m_0, \theta^t)$ is weakly decreasing in m_0 for all $\theta^t \in \Theta^t$ and all $t > 0$.

(b) For each $m_0 \in I$, $c_t^*(m_0, \theta^t)$ and $y_t^*(m_0, \theta^t)$ are differentiable in m_0 for Π-almost all history nodes θ^t.

Tax wedges in period 0 are as defined in Section 2. We derive the following arbitrage result for this general setting.

Proposition 3. *Suppose an allocation is constrained Pareto efficient. Then for all $\theta_0 \in \Theta$:*

$$\frac{\frac{T'(\theta_0)}{(1-T'(\theta_0))}}{\frac{\theta_0 u_{c\theta,0}}{u_{c,0}} - \frac{\theta_0 u_{y\theta,0}}{u_{y,0}}} = \frac{\frac{S'(\theta_0)}{(1+S'(\theta_0))}}{\frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau} \beta^t D_{0,t} \theta_t \left[u_{c\theta,t} c_{m,t}^* + u_{y\theta,t} y_{m,t}^*\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau} \beta^t \left[u_{c,t} c_{m,t}^* + u_{y,t} y_{m,t}^*\right]\right]} - \frac{\theta_0 u_{c\theta,0}}{u_{c,0}}} \tag{32}$$

This implies $T'(\theta_0)$ takes the opposite sign to $S'(\theta_0)$ if and only if:

$$\frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}c^*_{m,t}+u_{y\theta,t}y^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}c^*_{m,t}+u_{y,t}y^*_{m,t}\right]\right]}-\frac{\theta_0 u_{c\theta,0}}{u_{c,0}}<0 \tag{33}$$

As before, this provides a condition linking the ratios of marginal tax revenue and information rent changes as taxed behaviours are varied. The denominator on the left-hand side of (32) is positive under the Spence-Mirrlees condition, and so the sign of the left-hand side is given by the sign of the marginal income tax rate. Implications for the marginal savings tax rate depend on the sign of the denominator on the right-hand side, as (33) – making use of (15) – affirms.

The object in (33) is the derivative with respect to θ_0 of the log of the marginal rate of substitution between c_0 and m_0: does an increase in type raise or lower an individual's relative preference for saving? The effects of type on incentives for self insurance tend to lower it. To see this, define:

$$U_{m,1}(m_0,\theta_1):=\mathbb{E}\left[\sum_{t=1}^{T}\beta^{t-1}\left[u_{c,t}c^*_{m,t}+u_{y,t}y^*_{m,t}\right]\middle|\, m_0,\theta_1\right]$$

Thus $U_{m,1}(m_0,\theta_1)$ gives the marginal utility value of a differential increase in m_0, from the perspective of type θ_1. From the envelope condition in period 1, we can show:

Lemma 2. *The following are equivalent:*

1. *$U_{m,1}(m_0,\theta_1)$ is strictly decreasing in θ_1, given m_0.*
2. *$\mathbb{E}\left[\sum_{t=1}^{T}\beta^{t-1}D_{1,t}\theta_t\left[u_{c\theta,t}c^*_{m,t}+u_{y\theta,t}y^*_{m,t}\right]\middle|\, m_0,\theta_1\right]<0$ for all $\theta_1\in\Theta$.*

This is an immediate generalisation of Lemma 1, and the proof is omitted.

When $U_{m,1}(m_0,\theta_1)$ is decreasing, savings provide greater marginal utility when period-1 productivity is low – i.e., they provide self insurance. Since $\rho(\theta_1|\theta_0)D_{1,t}=D_{0,t}$, the first term in (33) will be negative so long as additional savings provide self insurance, and productivity is persistent.

The second term in (33) accounts for direct labour-consumption complementarities in the utility function. If consumption is a utility complement to leisure, then higher types – who work less at given income – will have a higher marginal utility of consumption. In this case $u_{c\theta,0}>0$. Conversely, $u_{c\theta}<0$ when consumption is a utility complement to work. Given this, the following corollary to Proposition 3 is immediate:

Corollary 1. *The following conditions are together sufficient for the marginal savings tax to take the opposite sign to the marginal income tax:*

1. *Savings provide self-insurance;*
2. *Productivity is persistent; and*
3. $u_{c\theta,0} \geq 0$.

Two brief points are worth highlighting here. First, note that the separable case, where $u_{c\theta,0} = 0$, is particularly relevant because it gives rise to an Inverse Euler equation under a fully optimal dynamic scheme – a result that has received substantial attention in the literature. We discuss the comparison between our main result and the Inverse Euler Equation in the next section.

Second, note that the Corollary's conditions are sufficient but *not* necessary for the two tax rates to take the opposite sign. A negative value for $u_{c\theta,0}$ implies that there is a trade-off between direct consumption-labour complementarities and the self-insurance channel. Taxing savings *and* income can be justified in these circumstances only if the direct channel outweighs the self-insurance channel. Whether the self-insurance channel outweighs direct preference complementarities or vice versa directly depends on the importance of the latter (summarized by the relative magnitude of $\theta u_{c\theta}/u_c$ vs. $\theta u_{y\theta}/u_y$) relative to the former (which depends on the persistence of shocks as summarised by the pass-through coefficients $D_{0,t}$ and the wealth effects on labour supply $y^*_{m,t}/y^*_t$). Since the impact of uninsurable income risk on savings decisions vanishes at the top of the wealth distribution, our results are likely to be of more relevance to tax treatment at the lower ends of the wealth and earnings distributions – where self-insurance matters.

5 Relationship to the inverse Euler equation

5.1 Discussion

We have obtained particularly strong results on the character of Pareto-efficient tax systems where utility is separable between consumption and leisure. To recap: the public savings wedge should be zero in this case if productivity is iid, and should take the opposite sign to the earnings wedge when productivity is persistent, so long as savings provide self-insurance. A sufficient condition for the latter is that leisure after period 0 should be a normal good – i.e., increasing in m_0.

These results are surprising, because they seem to conflict with the inverse Euler equation (IEE) – a cornerstone of the dynamic Mirrleesian literature. Recall that the IEE is a necessary condition for the design of an optimal dynamic insurance scheme across multiple time periods, in situations where consumption and leisure are separable in the utility function.[11] In our setting, it takes the following form between periods 0 and 1:

$$\frac{1}{u_{c,0}\left(\theta_0\right)} = \mathbb{E}_0\left[\frac{1}{\beta R u_{c,1}\left(\theta_1\right)}\right] \tag{34}$$

Applying Jensen's inequality to the right-hand side and rearranging yields:

$$u_{c,0}\left(\theta_0\right) \leq \mathbb{E}_0\left[\beta R u_{c,1}\left(\theta_1\right)\right] \tag{35}$$

The inequality here is strict so long as there is some variation in the marginal utility of consumption in period 1 – which, given concavity in the consumption utility function, is equivalent to the absence of complete insurance. This implies that for each θ_0 there is a positive private savings wedge:

$$s'\left(\theta_0\right) := \mathbb{E}_0\left[\frac{\beta R u_{c,1}\left(\theta_1\right)}{u_{c,0}\left(\theta_0\right)}\right] - 1 > 0 \tag{36}$$

Based on this wedge, Golosov, Kocherlakota and Tsyvinski (2003) write that "it is optimal for society to deter saving by taxing it" – a general conclusion, in a setting where productivities are allowed to follow arbitrary stochastic processes. This appears inconsistent with our conclusion that the public savings wedge should be zero when the productivities are iid over time, and is not easily reconciled with a conclusion that savings and labour income taxes take opposite signs with persistence.[12]

A first point to emphasise is that the arbitrage condition (32) and the IEE are both necessary conditions for an optimal insurance scheme. A Mirrleesian optimum links the realisation of θ_0 to an optimally-designed set of period-1 policy functions, $c_1\left(\theta_0, \theta_1\right)$ and $y_1\left(\theta_0, \theta_1\right)$ (and similar for $t > 1$). These functions are constructed to satisfy (34), together with other period-1 optimality conditions and incentive constraints. Their direct dependence on θ_0 can be replaced by dependence on a wealth index m_0, as outlined in Section 2, giving policy functions of the form $c_1^*\left(m_0\left(\theta_0\right), \theta_1\right)$ and $y_1^*\left(m_0\left(\theta_0\right), \theta_1\right)$.

[11] For derivations and discussion, see, for instance, Golosov, Kocherlakota and Tsyvinski (2003) and Farhi and Werning (2012)

[12] In their detailed simulations of the labour tax wedge under dynamically optimal Mirrleesian tax schemes, Farhi and Werning (2013) find that it is optimal to set the labour tax wedge to a positive value, in a setting where the IEE applies.

Our analysis takes these c_1^* and y_1^* functions as given, and explores whether an improvement is possible by varying the assignment of c_0, y_0 and m_0 as functions of θ_0. This admits only a very coarse search across continuation allocations, but it is clear that any alternative allocations that *are* considered in our exercise would also be available when designing the optimal dynamic Mirrlees scheme. Clearly Pareto improvements cannot be possible relative to this (constrained-optimal) scheme. Thus a dynamic Mirrleesian optimum must satisfy efficiency requirement (32), alongside the IEE.

The difference between the results instead derives from the different definitions of the savings wedge. Recall that we have defined $S'(\theta_0)$ as the gap between the value of savings and the *realised* marginal cost of providing them:

$$S'(\theta_0) = \frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}\cdot c_{m,t}^* + \frac{u_{y,t}}{u_{c,0}}\cdot y_{m,t}^*\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}\left[c_{m,t}^* - y_{m,t}^*\right]\right]} - 1 \tag{37}$$

The private savings wedge is defined as the the gap between the value of savings and the cost of providing a *uniform* increment to *consumption alone* in the *next period (alone)*. Thus, mechanically, $S'(\theta_0)$ could depart from $s'(\theta_0)$ because: (1) savings induce changes in consumption patterns across multiple periods, not just one; (2) savings yield returns that vary by future type history; and (3) savings induce changes to leisure, not just consumption.

Which of these channels is responsible for the difference between $s'(\theta_0)$ and $S'(\theta_0)$? The first possibility, which can be ruled out, is that the divergence comes from $S'(\theta_0)$ assessing costs over multiple time periods, rather than just the first. To see this, consider an alternative wedge measure to $s'(\theta_0)$, evaluated on a savings instrument that pays a profile of consumption returns across multiple future time periods – but retains a 'uniformity' restriction,[13] and assumes returns are spent on the consumption good alone. Using $d_{m,t}^*\left(\theta^t\right)$ to denote the

[13]The uniformity restriction is that the period-τ payout d_τ^* is measurable with respect to $\theta^{\tau-1}$, $d_{\tau-1}^* + \beta\left[\mathbb{E}_{\tau-1}\frac{u_{c,\tau}}{u_{c,\tau-1}}\right]d_\tau^*$ is measurable with respect to $\theta^{\tau-2}$, the object:

$$d_{\tau-2}^* + \beta\left[\mathbb{E}_{\tau-2}\frac{u_{c,\tau-1}}{u_{c,\tau-2}}\right]\left[d_{\tau-1}^* + \beta\left[\mathbb{E}_{\tau-1}\frac{u_{c,\tau}}{u_{c,\tau-1}}\right]d_\tau^*\right]$$

is measurable with respect to $\theta^{\tau-3}$, and so on.

marginal returns that are enjoyed in period t, any such wedge, $\bar{S}'_c(\theta_0)$, will satisfy:[14]

$$\bar{S}'_c(\theta_0) := \frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}\cdot d^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}d^*_{m,t}\right]} - 1 \geq s'(\theta_0) > 0 \tag{38}$$

with the first inequality strict whenever $d^*_{m,t} > 0$ for $t > 1$ along a positive-measure subset of histories, and incomplete insurance that endures beyond period 1.

The point is that the pure move from assessing a one-period wedge to a multi-period wedge will, if anything, imply a still greater positive gap between the value of savings and their cost. This is intuitive. If the IEE *also* applies between periods 1 and 2, 2 and 3, and so on, then a deferred consumption pattern can only compound the one-period consumption wedge already present in $s'(\theta_0)$.

Thus to explain how $S'(\theta_0)$ can be negative (or zero) whilst $s'(\theta_0)$ is positive, two possibilities remain. The first is that a dynamic Mirrleesian economy implicitly delivers returns to savings that are non-uniform across states ex-post. The second possibility is that future labour wedges raise costs.

To assess whether non-uniform returns are behind the divergence, consider a multi-period savings wedge that allows for non-uniform returns across type draws, but is still assessed as if all returns are consumed. The returns paid at history node θ^t are again denoted $d^*_t(\theta^t)$, now set to equal the consumption value of the returns individuals actually receive at the realised allocation, i.e.:

$$d^*_t(\theta^t) := c^*_{m,t}(\theta^t) + \frac{u_{y,t}(\theta^t)}{u_{c,t}(\theta^t)} y^*_{m,t}(\theta^t) \tag{39}$$

The wedge associated with this savings instrument is given by $S'_c(\theta_0)$, defined equivalently to $\bar{S}'_c$:

$$S'_c(\theta_0) := \frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[\frac{u_{c,t}}{u_{c,0}}\cdot d^*_t\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}R^{-t}d^*_t\right]} - 1 \tag{40}$$

The difference here relative to $\bar{S}'_c$ is that no measurability restrictions are imposed on $d^*_t(\theta^t)$. This can have implications for the wedge, becaue negative correlation between $u_{c,t}$ and $d^*_{m,t}$ (i.e., higher returns to saving at the upper end of the earnings distribution) can raise the real cost of asset provision, without changing its value.

[14] The overbar is used to denote uniform returns, and the c subscript the fact that costs are assessed based on the consumption good alone.

Yet if non-uniform returns are to account for the difference between S' and s', there are important implications for the effect of savings on leisure. The following Proposition makes this clear.

Proposition 4. *Suppose utility is separable between consumption and labour supply and $\tau < \infty$. If a constrained Pareto-efficient allocation satisfies the inverse Euler equation at all history nodes in all time periods, then either:*

1. $S'_c(\theta_0) \geq 0$, *or*
2. $y^*_{m,t}(\theta^t) > 0$ *for a positive-measure subset of θ^t histories and time periods.*

We know $S'(\theta_0) < 0$ in the separable case with persistence. Thus non-uniform returns can account for the difference in wedge values at an optimal allocation only if leisure ceases to be a normal good. We show directly, in a computational example with $\tau = 1$, that this is not generically true.

The remaining possibility is that the difference between $s'(\theta_0)$ and $S'(\theta_0)$ follows from including leisure in the cost of saving. Note:

$$1 + S'(\theta_0) = \frac{\mathbb{E}_0\left[\sum_{t=1}^{\tau} R^{-t}\left[c^*_{m,t} - (1 - T'_t)\, y^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau} R^{-t}\left[c^*_{m,t} - y^*_{m,t}\right]\right]}\left(1 + S'_c(\theta_0)\right) \tag{41}$$

where $\left(1 + \frac{u_{y,t}(\theta^t)}{u_{c,t}(\theta^t)}\right) := T'_t(\theta^t)$. If $y^*_{m,t} < 0$ and $T'_t(\theta_t) > 0$ along almost all histories, this implies $S'(\theta_0) < S'_c(\theta_0)$. This seems to be the relevant case. Since $\left(1 + \frac{u_{y,t}}{u_{c,t}}\right) := T'_t(\theta_t)$ gives the period-t earnings wedge, $S'(\theta_0)$ will fall below $S'_c(\theta_0)$ to the extent that savings induce a reduction in labour supply in the future, and future labour income is taxed. This tax externality effect is of direct relevance to policy when assessing the full fiscal cost of the savings decision, and its exclusion from $S'_c(\theta_0)$ limits the relevance of that measure for elasticity-based policy design, as we argue in Section 6.

5.2 Two-period simulation

We provide simulation results to augment the previous discussion, computing constrained-optimal allocations for a two-period dynamic Mirrleesian economy with type persistence, assuming a utilitarian welfare objective. The purpose of this is to confirm numerically that the optimal allocation obtained in a dynamic Mirrleesian economy can be consistent with

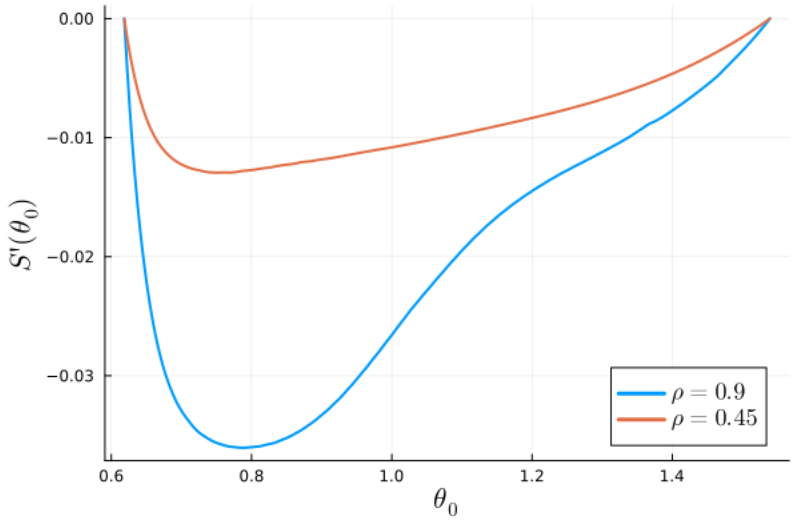


Figure 1: Marginal savings taxes as a function of period-0 type

a fiscal subsidy to saving, and that this occurs through the income effect of savings on on labour supply.

The simulations assume that utility takes the isoelastic, separable form specified in Section 3, with $\sigma = 1$ (i.e., log consumption utility) and $\nu = 2$. We set $\beta = R^{-1} = 0.975^{20}$, so that a period can be interpreted as 20 years at an annual interest rate around 2.5 per cent. Log productivity is modelled as following a truncated AR(1) process, with persistence parameter ρ.

Figure 1 charts computed values for the marginal savings tax as a function of θ_0, for the cases $\rho = 0.45$ and $\rho = 0.9$. The zero values of marginal taxes at endpoints are by-products of the bounded type distribution that we have assumed. The negative values for the marginal tax rate in at all other values for θ_0 are consistent with the theoretical results, given positive values for the marginal income tax.

Figure 2 shows three period-1 budget constraints, plotted for three distinct values for the savings level m_0, under the assumption $\rho = 0.45$.[15] The dots on the schedules identify the equilibrium pair $(c_1^*(m_0, \theta_1), y_1^*(m_0, \theta_1))$ at six specific values for θ_1, showing how these vary as m_0 increases. Two features of the plots are notable. The first is that leisure is a normal

[15]The lines plot *realised* equilibrium allocations for period 1. It is a trivial extension to prescribe consumption levels at values of y_1 that are not chosen by any type in equilibrium, so long as the additional allocations remain dominated by the intended equilibrium for all θ_1. The plots are optimal non-linear budget constraints for period 1. Their near-linearity is an interesting feature, which seems common to many calibrations.

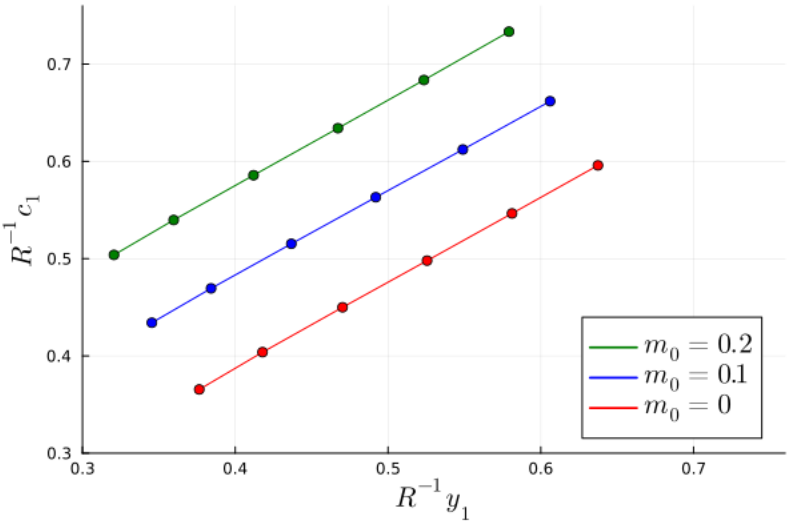


Figure 2: Period-1 budget constraints as m_0 varies

good in m_0, despite nonlinearities in the budget constraints. Under Lemma 2, and given the separability of the utility function (so $u_{c\theta} = 0$), this implies that m_0 provides self-insurance.

The second point to notice is that all three budget constraints have a slope that is less than unity, and so the horizontal distance between any pair of constraints is everywhere greater than the vertical distance.[16] Thus, savings allow individuals to reduce their production by a greater amount than they allow them to increase their consumption. This is the geometric counterpart to the fact that in distorted economies, leisure and consumption are not equally costly for the policymaker to provide in period 1. Positive marginal income taxes – which are what ensure the period-1 budget constraint has a slope less than unity – mean that leisure is kept inefficiently high by the distortionary income tax system. This means the use of savings to provide additional leisure is more costly to the policymaker than their use to provide additional consumption.

The full marginal resource cost of extra savings depends on the exact manner in which the extra resources are allocated between leisure and consumption, averaged across all period-1 productivity draws. Given the savings subsidies illustrated in Figure 1, the realised vector of choices is sufficiently skewed towards leisure to imply a subsidy. If savings *were* to be used exclusively for consumption, their resource cost would be lower – and, in the present case, individuals would receive less back than they paid in. There would be a savings tax. But this

[16] The units on both axes correspond to units of the period-0 consumption good.

is not the way in which the extra resources are actually used.

6 Relationship to elasticities

The analysis so far has adopted an information-theoretic approach to the problem – emphasising the trade off that arises in determining equilibrium information rents, and the associated arbitrage requirement. In particular, equation (30) in the two-period setting, and – for the more general case – equation (32), linked the marginal tax rates on income and savings with the effect of changes in production and saving on 'information rents'. Though extremely useful theoretical devices, information rents are difficult objects to operationalise practically. They describe the derivative of one unobservable theoretical construct with respect to another – the utility function with respect to a latent type variable. This contrasts with the strong trend in the applied Public Economics literature towards summarising key lessons in the form of measurable, observable 'sufficient statistics' – an approach that aids both empirical applicability and intuitive understanding.

In this section we therefore explain how our results can be translated into statements about Pareto-efficient tax rates that are written in terms of behavioural elasticities, together with the distributions of realised choice variables. The rules that we obtain are based on the idea that the fiscal losses due to behavioural effects of tax distortions should be equalised across bases, for each individual type. Importantly, the relevant measure of the savings tax wedge for this exercise proves to be the public savings wedge that we have defined.

6.1 Compensated elasticities

We make particular use of the *compensated* elasticities of y_0 and m_0 with respect to their respective prices. Defining these objects in an environment with nonlinear pricing is non-trivial, because responses may be sensitive to the precise way the price schedule is being perturbed – a uniform proportional change in prices may induce different responses from a uniform level change, for instance. Our approach is constructed to preserve the key symmetry and concavity properties of the matrix of compensated demand changes.

Specifically, compensated demands are defined by the solution to the following expendi-

ture minimisation problem:

$$\min_{c_0, y_0, m_0} \{c_0 + pm_0 + S(m_0) - ay_0 + T(y_0)\} \tag{42}$$
$$\text{s.t.:} u(c_0, y_0; \theta_0) + \mathbb{E}_0 \left[\sum_{t=1}^{\tau} \beta^t u\left(c_t^*\left(m_0, \theta^t\right), y_t^*\left(m_0, \theta^t\right); \theta_t\right) \right] \geq \bar{U}$$

where $\bar{U}$ is a benchmark level of lifetime utility, and p and a are, respectively, the pre-tax prices of the savings good and production good in period 0, expressed in units of the period-0 consumption good. Both of these prices equal 1 in the main analysis, and their inclusion here is solely for the purposes of defining elasticities. Since the tax schedules and future policy functions are to be held constant throughout, the corresponding expenditure function can be denoted $e(p, a; \bar{U}, \theta_0)$, and the optimal choices similarly: $m_0(p, a; \bar{U}, \theta_0)$, etc.

It is simple to show that the expenditure function is concave in (p, a), and satisfies the following envelope results, replicating Shephard's Lemma:

$$\frac{de(p, a; \bar{U}, \theta_0)}{dp} = m_0(p, a; \bar{U}, \theta_0) \tag{43}$$
$$\frac{de(p, a; \bar{U}, \theta_0)}{da} = -y_0(p, a; \bar{U}, \theta_0) \tag{44}$$

The compensated effect of a change in p on m_0 is then $\frac{dm_0(p,a;\bar{U},\theta_0)}{dp}$, and so on. [17]

Finally, we define elasticities by reference to these compensated demand changes, as fol-

[17] These envelope results allow Slutsky symmetry to be retained; specifically:

$$\frac{dm_0(p, a; \bar{U}, \theta_0)}{da} = \frac{d}{da}\left[\frac{de(p, a; \bar{U}, \theta_0)}{dp}\right] \tag{45}$$
$$= \frac{d}{dp}\left[\frac{de(p, a; \bar{U}, \theta_0)}{da}\right] \tag{46}$$
$$= -\frac{dy_0(p, a; \bar{U}, \theta_0)}{dp} \tag{47}$$

lows:

$$\varepsilon_{yy} := \frac{a - T'(y_0)}{y_0} \frac{dy_0(p, a; \bar{U}, \theta_0)}{da} \tag{48}$$

$$\varepsilon_{ym} := \frac{p + S'(m_0)}{y_0} \frac{dy_0(p, a; \bar{U}, \theta_0)}{dp} \tag{49}$$

$$\varepsilon_{my} := \frac{a - T'(y_0)}{m_0} \frac{dm_0(p, a; \bar{U}, \theta_0)}{da} \tag{50}$$

$$\varepsilon_{mm} := \frac{p + S'(m_0)}{m_0} \frac{dm_0(p, a; \bar{U}, \theta_0)}{dp} \tag{51}$$

with ε_{my} thus giving the compensated change in m_0 as the price of y_0 is varied, and so on. In all cases of interest, these elasticities will be considered at values $a = p = 1$. It is important to stress that these elasticities are not 'primitives': they depend on the price system, utility level and type of the agent in question.

Finally, note that concavity of the expenditure function implies $\varepsilon_{mm} \leq 0$, $\varepsilon_{yy} \geq 0$ and $\frac{\varepsilon_{yy}}{\varepsilon_{my}} \geq \frac{\varepsilon_{ym}}{\varepsilon_{mm}}$.[18] These inequalities bind when the expenditure function is weakly concave, in which case original choices can remain unaffected by price changes – i.e., compensated demands are perfectly inelastic in response to certain price changes. Otherwise the inequalities are strict. When choice is perfectly inelastic, tax design can be nondistortionary. Policy design in this case is trivial, and so for the remainder we focus on the case in which the inequalities are strict.

6.2 Elasticities and types

The following Lemma shows that there is a direct and intuitive link between the responsiveness of production and saving to changes in type, and these compensated elasticities.

[18] The last of these follows from the fact that the determinant of a negative-definite matrix is positive.

Lemma 3. *For all $\theta_0 \in \Theta$ and at any allocation, the following two conditions hold:*

$$\frac{1}{y_0}\frac{dy_0}{d\theta_0} = \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{u_{y\theta,0}}{u_{y,0}}\right]\varepsilon_{yy} + \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t} + u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\middle|\,\theta_0\right]}\right]\varepsilon_{ym} \tag{52}$$

$$\frac{1}{m_0}\frac{dm_0}{d\theta_0} = \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{u_{y\theta,0}}{u_{y,0}}\right]\varepsilon_{my} + \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t} + u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}\right]\varepsilon_{mm} \tag{53}$$

This result holds irrespective of the efficiency, or otherwise, of the tax system. It is helpful because it links derivatives of marginal rates of substitution with respect to type (the objects in square brackets), to derivatives of choice with respect to prices. Intuitively, the Lemma says that changes to choice by type match the changes that *would* be induced by a particular vector of price changes – namely, price changes that correspond to the effect of θ_0 on marginal rates of substitution.

The link to observable objects can be made stronger by relating type derivatives to properties of the realised distribution of choice variables. Let $\Pi^y\left(y_0\right)$ denote the realised distribution (cdf) of output choices, with density $\pi^y\left(y_0\right)$, with $\Pi^m\left(m_0\right)$ and $\pi^m\left(m_0\right)$ the corresponding objects for m_0. Denote by $\alpha^y\left(y_0\right)$ the local Pareto statistic for y_0, defined as follows:

$$\alpha^y\left(y_0\right) := \frac{\pi^y\left(y_0\right)y_0}{1-\Pi^y\left(y_0\right)} \tag{54}$$

and symmetically for $\alpha^m\left(m_0\right)$. The following result is then a simple development of Lemma 3:

Proposition 5. *Suppose the expenditure function is strictly concave, and y_0 and m_0 both strictly increasing in type. Then:*

$$sign\left[\alpha^m\left(m_0\right)\cdot\varepsilon_{my} - \alpha^y\left(y_0\right)\cdot\varepsilon_{yy}\right] = sign\left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t} + u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\middle|\,\theta_0\right]}\right] \tag{55}$$

$$\alpha^m\left(m_0\right)\cdot\varepsilon_{mm} - \alpha^y\left(y_0\right)\cdot\varepsilon_{ym} < 0 \tag{56}$$

$\alpha^m(m_0) \cdot \varepsilon_{my}$ and $\alpha^y(y_0) \cdot \varepsilon_{yy}$ are summary measures of the aggregate responsiveness of m_0 and y_0, respectively, when the returns to work increase. They combine the elasticity with a measure of the density of demand for each good at the point of interest. The analysis in Section 4 showed that Pareto-efficient income and savings taxes take the opposite sign whenever the object on the right-hand side in (55) is positive. Proposition 5 shows that this occurs precisely when the aggregate responsiveness of savings is greater than that of income, as the returns to work are increased.

Condition (56) is derived identically to (55), and relates to the effect of an increased cost of acquiring period-1 wealth. The strictly positive sign comes from the Spence-Mirrlees condition, which guarantees the marginal rate of substitution between consumption and labour supply is monotone in type.

6.3 Sufficient statistics for Pareto efficiency

Combining Lemma 3 and Proposition 3, our Pareto efficiency condition can be re-expressed by reference to elasticities and the realised density of market choices:

Proposition 6. *At a constrained Pareto-efficient allocation, income and savings taxes satisfy the following two conditions for all $\theta_0 \in \Theta$:*

$$\Phi(\theta_0) = \alpha^y(y_0) \cdot \left\{ \frac{T'(y_0)}{(1-T'(y_0))} \varepsilon_{yy} - \frac{S'(m_0)}{(1+S'(m_0))} \varepsilon_{ym} \right\} \tag{57}$$

$$\Phi(\theta_0) = \alpha^m(m_0) \cdot \left\{ \frac{T'(y_0)}{(1-T'(y_0))} \varepsilon_{my} - \frac{S'(m_0)}{(1+S'(m_0))} \varepsilon_{mm} \right\} \tag{58}$$

for some constant $\Phi(\theta_0)$.

These expressions provide a novel link between the Ramsey and Mirrleesian approaches to optimal taxation. The bracketed terms on the right of (57) and (58) capture the extent to which the tax system is distorting, respectively, the work and savings of type θ_0. They are localised versions of the 'indices of distortion' first used by Samuelson (1986) to generalise Ramsey results for the linear, homogeneous-consumer commodity tax problem.[19]

These local distortion measures are interacted with the local Pareto parameter, which captures the intensity of consumption of the two goods at the point of interest, and is the only component of expressions (57) and (58) without direct analogue in Samuelson (1986).

[19] A textbook presentation can be found in Atkinson and Stiglitz (1980), Chapter 12. Close parallels exist in the regulatory literature on multi-good Ramsey pricing: see, for instance, Armstrong and Vickers (2018).

This generalisation of linear Ramsey formulae is, to the best of our knowledge, novel to our paper, but it is entirely intuitive: deterring the consumption of a taxed good is more costly the more people are consuming it, and the density, which features as a key component of the Pareto statistic, provides a localised measure of this total demand.[20] Thus equations (57) and (58) refine the idea, familiar from Ramsey problems, that an optimal tax system should deter the consumption of all taxed goods equally at the margin. The constant $\Phi(\theta_0)$ captures the extent of the distortion, common across goods. It will take different values depending on the degree of redistribution targeted.

It remains to link these elasticity-based characterisations to the earlier results. Combining (57) and (58), and rearranging:

$$\frac{T'(y_0)}{(1-T'(y_0))}\left[\varepsilon_{my}\alpha^m(m_0)-\varepsilon_{yy}\alpha^y(y_0)\right]=\frac{S'(m_0)}{(1+S'(m_0))}\left[\varepsilon_{mm}\alpha^m(m_0)-\varepsilon_{ym}\alpha^y(y_0)\right] \tag{59}$$

Intuitively, the left-hand side here summarises the relative distortionary implications of the headline labour income tax; the right-hand side summarises the effects of the tax on wealth holdings. Combining with proposition 5, we obtain

Proposition 7. *At a constrained Pareto-efficient allocation, for all $\theta_0 \in \Theta$, we have*

$$sign\left[T'(y_0)\cdot S'(m_0)\right]=-sign\left[\varepsilon_{my}\alpha^m(m_0)-\varepsilon_{yy}\alpha^y(y_0)\right] \tag{60}$$

Suppose for now that it is optimal to tax labor income ($T'(y_0)>0$). Proposition 7 states that savings should be subsidized ($S'(m_0)<0$) if the product of the local Pareto coefficient of savings times the compensated elasticity of savings to a change in the retention rate ($1-T'(y_0)$) exceeds the product of the local Pareto coefficient of earnings times the compensated labor supply elasticity to a change in $1-T'(y_0)$. This proposition is the analogue of the condition Hellwig and Werquin (2024) derive for optimal savings taxes on top income earners under generic period-0 preferences over consumption, earnings and savings.

The intuition behind the result can be understood as follows: following the intuition from Atkinson-Stiglitz, savings should be subsidized if higher types have a weaker intrinsic propensity to save, and taxed if their intrinsic propensity to save is stronger than for lower types. The ratio of local Pareto coefficients indicates how much savings co-vary locally with income. The ratio of compensated elasticities indicates how much a given change in net

[20] Saez (2001) showed the relevence of the empirical earnings density in weighting the importance of local distortionary effects in static Mirrlees models.

income is passed through to higher savings, for given preferences – or in other words, how much of the co-variation in savings can be explained by the behavioral response of savings to income changes. If the two ratios coincide, the behavioral response of savings fully accounts for the co-variation of savings with income, i.e. heterogeneous preferences do not play any role in accounting for the observed co-variation of savings with income, i.e. we are in the case of Atkinson-Stiglitz. If instead the two ratios differ, then heterogeneous preferences must account for the gap between the two, and there is scope to either tax or subsidize savings. While we previously tied the sign of a savings tax or subsidy to the existence of a precautionary savings motive (i.e. the value of σ and the persistence of types), proposition 7 links the strength of this precautionary savings motive to sufficient statistics and behavioral responses - Pareto coefficients and compensated elasticities.

Hellwig and Werquin (2024) focus on the top of the type distribution and argue, based on estimates of Pareto coefficients for consumption, income and savings, as well as estimates of consumption and savings responses to income tax changes, that $\varepsilon_{my}\alpha^m(m_0) < \varepsilon_{yy}\alpha^y(y_0)$ for top income earners, and hence that it is optimal to shift part of their tax burden from income to savings. This is intuitive if one takes into account that self-insurance motives are likely to be weak for top income earners, and therefore preference complementarity between consumption and earnings makes it optimal to subsidize their consumption - or equivalently tax their savings.

However, the same conditions and analysis can be applied also to low-income households by replacing the upper local Pareto coefficients $\alpha^m(m_0)$ and $\alpha^y(y_0)$ with their lower counter-parts: $\underline{\alpha}^m(m_0) = \frac{\pi^m(m_0)m_0}{\Pi^m(m_0)}$ and $\underline{\alpha}^y(y_0) = \frac{\pi^y(y_0)y_0}{\Pi^y(y_0)}$.[21] Hence we again obtain that the comparison of local Pareto coefficients and behavioral elasticities fully identify the strength of precautionary savings motives on low-income households, and thus provide an empirical test of the case we have emphasised, where Pareto efficient labour and savings taxes take opposite signs. In this case, for a given type a cut in labor income taxes induces a high demand for savings, as individuals use the tax cut to cover themselves against low future earnings shocks. On the other hand, a marginal increase in type has a much weaker effect on savings because the pass-through from higher current earnings is partly offset through improved future earnings prospects and thus weaker precautionary savings incentives. In other words, if the self-insurance motive is the dominant force shaping how the propensity to save co-moves with individual's propensity to work, then savings become a complement

[21] Since choices are assumed monotone in θ_0, $\Pi^m(m_0) = \Pi^y(y_0) = \Pi(\theta_0)$, and so the replacement of upper with lower Pareto tail just implies multiplying (59) through by the scalar $\frac{1-\Pi(\theta_0)}{\Pi(\theta_0)}$.

to earnings, and hence should be subsidized, rather than taxed.

7 Conclusion

A good is subsidised when increased purchases of it are costly to the government. Under this definition, our paper has shown that there is a general, Paretian case for savings to be subsidised, so long as labour income is taxed. This result relies on savings providing self-insurance, so that their ex-post value is greater when times are bad. When this is true, type persistence means that self-insurance will be more valuable to those whose current productivity is low. A savings subsidy can then provide incentives for low types to work – and so help to offset the distortionary effect of the labour income tax – at a comparatively low cost to distributive objectives.

The surprising feature of this argument is that it seems to contradict conventional understandings of the role of savings taxes in incomplete market economies. This is particularly the case when consumption and labour supply are assumed to be separable in the utility function. When this is true, we have shown that our channel is the *only* force that matters for the sign of the public savings wedge. This appears difficult to square with the conventional interpretation of the inverse Euler condition as implying a savings tax. As we have argued, the two are perfectly consistent once it is recognised that the relevant fiscal measure of the savings wedge differd from the 'private' wedge in the standard consumption Euler condition. The public wedge is a particularly useful construct, because – as we have just shown – it allows for a direct generalisation of elasticity-based rules for optimal tax design in a nonlinear setting.

The relevance of our message – 'subsidise self-insurance' – evidently depends on the importance of the insurance motive in accounting for observed savings behaviour. It is well known that this channel is weak at the upper ends of the wealth distribution, where just small complementarities between consumption and work will be sufficient to provide an offsetting case for a savings tax. Things are quite different at the lower end of the earnings distribution, where uncertainty and the need for a buffer stock can be a critical motive. Our result relies on very few assumptions about the appropriate welfare criterion, or the nature of insurance in the future. It provides a novel consideration when assessing the tax treatment of low earners.

References

[1] Armstrong, M., and J. Vickers (2018), "Multiproduct Pricing Made Simple", *Journal of Political Economy*, 126(4), pp. 1444–1471.

[2] Atkinson, A.B., and J.E. Stiglitz (1976), "The Design of Tax Structure: Direct Versus Indirect Taxation", *Journal of Public Economics* 6(1/2), pp. 55–75.

[3] Atkinson, A.B., and J.E. Stiglitz (1980), *Lectures on Public Economics*, Princeton University Press

[4] Corlett, W.J., and D.C. Hague (1953), "Complementarity and the excess burden of taxation", *The Review of Economic Studies* 21(1), pp. 21–30.

[5] Diamond, P.A., and J.A. Mirrlees (1971), "Optimal Taxation and Public Production II: Tax Rules", *The American Economic Review*, 61(3), pp. 261–278.

[6] Farhi, E., and I. Werning (2012), "Capital Taxation: Quantitative Explorations of the Inverse Euler Equation", *Journal of Political Economy* 120(3), pp. 398–445.

[7] Farhi, E., and I. Werning (2013), "Insurance and Taxation over the Life Cycle", *The Review of Economic Studies* 810(2), pp. 596–635.

[8] Golosov, M., N. Kocherlakota and A. Tsyvinski (2003), "Optimal Indirect and Capital Taxation", *The Review of Economic Studies*, 70(3), pp. 569–587.

[9] Heathcote, J., K. Storesletten, and G.L. Violante (2005), "Two Views of Inequality over the Life Cycle", *Journal of the European Economic Association*, 3 (2/3), pp. 765–775

[10] Hellwig, C., and N. Werquin (2024), "Using Consumption Data to Derive Optimal Income and Capital Taxes", Working Paper

[11] Kaplow, L. (2024), "Optimal Income Taxation", *Journal of Economic Literature*, 62(2), pp. 637–738.

[12] Kocherlakota, N.R. (2005), "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation", *Econometrica*, 73(5), pp. 1587–1621.

[13] Milgrom, P., and I. Segal (2002), "Envelope Theorems for Arbitrary Choice Sets", *Econometrica*, 70(2), pp. 583–601.

[14] Pavan, A., I. Segal and J. Toikka (2014), "Dynamic Mechanism Design: A Myersonian Approach", *Econometrica*, 82(2), pp. 601–653.

[15] Piketty, T., and E. Saez (2013), "Optimal Labor Income Taxation", Ch 7 in Auerbach et al. (eds), *Handbook of Public Economics*, Vol. 5, pp 391–474.

[16] Ramsey, F.P. (1927), "A Contribution to the Theory of Taxation", *Economic Journal*, 37, pp. 47–61.

[17] Saez, E. (2001), "Using Elasticities to Derive Optimal Income Tax Rates", *The Review of Economic Studies*, 68(1), pp. 205–229.

[18] Samuelson, P.A. (1986), "Theory of Optimal Taxation", *Journal of Public Economics*, 30(2), pp. 137–143

[19] Storesletten, Telmer and Yaron (2004), "Cyclical Dynamics in Idiosyncratic Labor Market Risk", *Journal of Political Economy*, 112(3), pp. 695–717.

A Proofs

A.1 Proof of Proposition1

To establish the envelope condition, it is convenient first to apply a change of variables to the stochastic process governing productivity.[22] Let θ_t be determined by a 'primitive' random variable drawn independently from a uniform distribution on $[0,1]$ each period. Let this variable be denoted α_t, with history $\alpha^t \in [0,1]^{t+1}$ up to period t, and $\alpha_t^s \in [0,1]^{s-t+1}$ denoting a partial history from periods t to s inclusive ($s \geq t$). For all $t \geq 0$ and any given α^t, $\theta_t\left(\alpha^t\right)$ is then defined implicitly by the recursions:

$$\Pi_0\left(\theta_0\left(\alpha^0\right)\right) = \alpha_0 \tag{61}$$

$$\Pi\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right) = \alpha_t \tag{62}$$

($t > 0$). We have assumed that $\Pi_0(\cdot)$ and $\Pi(\cdot|\theta_{t-1})$ are strictly monotone and continuously differentiable, and so the sequence of random variables α^t is measurable with respect to $\left(\theta_0, \theta^t\right)$, and vice-versa. $\theta^t\left(\alpha^t\right)$ then denotes the history $\left\{\theta_1\left(\alpha^1\right), \theta_2\left(\alpha^2\right), ..., \theta_t\left(\alpha^t\right)\right\}$, and $\theta_t^s\left(\alpha^s\right)$ the partial history $\left\{\theta_t\left(\alpha^t\right), \theta_{t+1}\left(\alpha^{t+1}\right), ..., \theta_s\left(\alpha^s\right)\right\}$.

Given policy functions $c_s\left(\theta_0, \theta^s\right)$ and $y_s\left(\theta_0, \theta^s\right)$, a corresponding set of objects can be defined on the space of α^t draws, in the obvious way. With some abuse, we write $c_s\left(\alpha^s\right)$ to denote $c_s\left(\theta_0\left(\alpha_0\right), \theta^s\left(\alpha^s\right)\right)$ and $y_s\left(\alpha^s\right)$ for $y_s\left(\theta_0\left(\alpha_0\right), \theta^s\left(\alpha^s\right)\right)$.

[22] A similar approach is taken by Pavan, Segal and Toikka (2014).

We start with some simple results on the link from α_t to subsequent productivity draws. Differentiating (61) and (62) gives:

$$\pi\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)\frac{d\theta_t\left(\alpha^t\right)}{d\alpha_t}=1 \tag{63}$$

$$\frac{d\Pi\left(\theta_{t+s}\left(\alpha^{t+s}\right)|\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\right)}{d\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}\frac{d\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}{d\alpha_t}$$
$$+\pi\left(\theta_{t+s}\left(\alpha^{t+s}\right)|\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\right)\frac{d\theta_{t+s}\left(\alpha^{t+s}\right)}{d\alpha_t}=0 \tag{64}$$

the latter for $s>0$. Thus:

$$\frac{d\theta_{t+s}\left(\alpha^{t+s}\right)}{d\alpha_t}=\frac{\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\frac{d\left(1-\Pi\left(\theta_{t+s}\left(\alpha^{t+s}\right)|\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\right)\right)}{d\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}}{\theta_{t+1}\left(\alpha^{t+1}\right)\pi\left(\theta_{t+s}\left(\alpha^{t+s}\right)|\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\right)} \tag{65}$$

$$\cdot\frac{\theta_{t+s}\left(\alpha^{t+s}\right)}{\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}\frac{d\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}{d\alpha_t} \tag{66}$$

$$=\rho\left(\theta_{t+s}\left(\alpha^{t+s}\right)|\theta_{t+s-1}\left(\alpha^{t+s-1}\right)\right)\frac{\theta_{t+s}\left(\alpha^{t+s}\right)}{\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}\frac{d\theta_{t+s-1}\left(\alpha^{t+s-1}\right)}{d\alpha_t} \tag{67}$$

$$=\left[\prod_{r=1}^{s}\rho\left(\theta_{t+r}\left(\alpha^{t+r}\right)|\theta_{t+r-1}\left(\alpha^{t+r-1}\right)\right)\right]\frac{\theta_{t+s}\left(\alpha^{t+s}\right)}{\theta_t\left(\alpha^t\right)}\frac{d\theta_t\left(\alpha^t\right)}{d\alpha_t} \tag{68}$$

Suppose a direct mechanism exists, in which individuals report a type $\tilde{\alpha}_t\in[0,1]$ period by period, and receive the corresponding allocations $c_t\left(\tilde{\alpha}^t\right)$ and $y_t\left(\tilde{\alpha}^t\right)$, as a function of their history of reports. Given a history of type draws α^{t-1} and reports $\tilde{\alpha}^{t-1}$, their period-t **choice** is a period-t type report $\tilde{\alpha}_t\in[0,1]$, coupled with a state-contingent plan for future reports, $\tilde{\alpha}_s\left(\alpha_{t+1}^s\right):[0,1]^{s-t}\to[0,1]$, for all $s>t$. Denote the space of such choices $\mathcal{X}$, with generic element x. Each choice in $\mathcal{X}$ induces a state-contingent sequence of consumption and production levels, via the allocation functions $c_s\left(\tilde{\alpha}^s\right)$ and $y_s\left(\tilde{\alpha}^s\right)$, which in turn implies a value for expected lifetime utility in t:

$$U_t\left(x|\,\alpha_t,\alpha^{t-1},\tilde{\alpha}^{t-1}\right):=\mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}\left[u\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\right]\middle|\,\alpha^{t-1},\tilde{\alpha}^{t-1},\alpha_t\right] \tag{69}$$

where $u\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)$ is shorthand for $u\left(c_s\left(\tilde{\alpha}^s\right),y_s\left(\tilde{\alpha}^s\right);\theta_s\left(\alpha^s\right)\right)$, and $\tilde{\alpha}^s$ is the report history:

$$\left\{\tilde{\alpha}^{t-1},\tilde{\alpha}_t,\tilde{\alpha}_{t+1}\left(\alpha_{t+1}\right),...,\tilde{\alpha}_s\left(\alpha_{t+1}^s\right)\right\}$$

induced by x. Dependence on the true and reported histories, α^{t-1} and $\tilde{\alpha}^{t-1}$, is left implicit in what follows.

Holding constant the choice $x\in\mathcal{X}$ (and α^{t-1} and $\tilde{\alpha}^{t-1}$), the function $U_t\left(x|\,\alpha_t\right)$ is differentiable with respect to α_t, its derivative given by:

$$
\begin{aligned}
U_{\alpha,t}\left(x|\,\alpha_t\right) =& \mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}\left[u_\theta\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\frac{d\theta_s\left(\alpha^s\right)}{d\alpha_t}\right]\middle|\,\alpha^t\right] \\
=& \frac{d\theta_t\left(\alpha^t\right)}{d\alpha_t}\mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}\frac{\left[\prod_{r=t}^{s}\rho\left(\theta_r\left(\alpha^r\right)|\theta_{r-1}\left(\alpha^{r-1}\right)\right)\right]}{\rho\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)}\left[u_\theta\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\frac{\theta_s\left(\alpha^s\right)}{\theta_t\left(\alpha^t\right)}\right]\middle|\,\alpha^t\right] \\
=& \frac{d\theta_t\left(\alpha^t\right)}{d\alpha_t}\frac{1}{\theta_t\left(\alpha^t\right)}\mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}D_{t,s}\left(\theta_0\left(\alpha_0\right),\theta^s\left(\alpha^s\right)\right)\left[\theta_s\left(\alpha^s\right)u_\theta\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\right]\middle|\,\alpha^t\right] \qquad (70)
\end{aligned}
$$

The assumption that information rents are bounded can be restated as:

Assumption. *For every* $\left(\alpha^{t-1},\tilde{\alpha}^{t-1}\right)$*, there exists a function* $b:\Theta\to\mathbb{R}$*, integrable on* Θ*, such that:*

$$
b\left(\theta_t\left(\alpha^{t-1},\alpha_t\right)\right)\geq\left|\frac{1}{\theta_t\left(\alpha^t\right)}\mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}D_{t,s}\left(\theta_0\left(\alpha_0\right),\theta^s\left(\alpha^s\right)\right)\left[\theta_s\left(\alpha^s\right)u_\theta\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\right]\middle|\,\alpha^{t-1},\alpha_t\right]\right| \qquad (71)
$$

for all $x\in\mathcal{X}$*.*

This implies:

$$
\begin{aligned}
&\left|\frac{1}{\theta_t\left(\alpha^t\right)\pi\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)}\mathbb{E}\left[\sum_{s=t}^{\tau}\beta^{s-t}D_{t,s}\left(\theta_0\left(\alpha_0\right),\theta^s\left(\alpha^s\right)\right)\left[\theta_s\left(\alpha^s\right)u_\theta\left(\tilde{\alpha}^s\left(x\right);\theta_s\left(\alpha^s\right)\right)\right]\middle|\,\alpha^t\right]\right| \\
\leq& b\left(\theta_t\left(\alpha^{t-1},\alpha_t\right)\right)\frac{1}{\pi\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)} \qquad (72)
\end{aligned}
$$

The composite function on the right-hand side here is integrable on $[0,1]$. To see this, note:

$$
\begin{aligned}
\int_{\underline{\theta}}^{\theta_t\left(\alpha^{t-1},\alpha_t\right)}b\left(\theta\right)d\theta &= \int_0^{\alpha_t}b\left(\theta_t\left(\alpha^{t-1},\alpha\right)\right)\frac{d\theta_t\left(\alpha^{t-1},\alpha\right)}{d\alpha}d\alpha \\
&= \int_0^{\alpha_t}b\left(\theta_t\left(\alpha^{t-1},\alpha\right)\right)\frac{1}{\pi\left(\theta_t\left(\alpha^{t-1},\alpha\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)}d\alpha
\end{aligned}
$$

$\theta_t\left(\alpha^{t-1},\alpha_t\right)$ is a nondecreasing, differentiable function of α_t, and so integrability of $b\left(\theta\right)$ on Θ implies integrability of $b\left(\theta_t\left(\alpha^{t-1},\alpha_t\right)\right)\frac{1}{\pi\left(\theta_t\left(\alpha^t\right)|\theta_{t-1}\left(\alpha^{t-1}\right)\right)}$ as a function of α_t on $[0,1]$. Given this, Theorem 2 in Milgrom and Segal (2002) can be applied. Assuming (without loss, by the revelation principle), that the mechanism induces truthful reporting, and using $U_t^\alpha\left(\alpha_t\right)$ to denote the maximised value that can be obtained by an individual with current type α_t at the given history, we have:

$$U_t^{\alpha}(\alpha_t) - U_t^{\alpha}(0) = \int_0^{\alpha_t} \frac{1}{\theta_t(\alpha)} \mathbb{E}\left[\sum_{s=t}^{\tau} \beta^{s-t} D_{t,s}\left(\theta_0(\alpha_0), \theta^s(\alpha^s)\right) \left[\theta_s(\alpha^s) u_\theta(\alpha^s; \theta_s(\alpha^s))\right] \middle| \alpha^{t-1}, \alpha \right] \frac{d\theta_t(\alpha^t)}{d\alpha_t} d\alpha \tag{73}$$

$$= \int_{\underline{\theta}}^{\theta_t(\alpha^{t-1}, \alpha_t)} \frac{1}{\theta} \mathbb{E}\left[\sum_{s=t}^{\tau} \beta^{s-t} D_{t,s}\left(\theta_0(\alpha_0), \theta^s(\alpha^s)\right) \left[\theta_s(\alpha^s) u_\theta(\alpha^s; \theta_s(\alpha^s))\right] \middle| \theta \right] d\theta \tag{74}$$

Reversing the change of variables, any direct mechanism that takes reports of α as inputs has an equivalent in which reports of θ are given each period, leaving the choice set unaffected.[23] For each report history $\tilde{\alpha}^{t-1}$ there is a corresponding sequence of θ reports $\left(\tilde{\theta}_0, \tilde{\theta}^{t-1}\right)$, and for each true history α^t, there is a single θ_t sufficient to summarise the effects of past type draws on current payoffs (types are Markovian). It follows that condition (74) will continue to apply under the strategically-equivalent mechanism in which θ reports are obtained period by period. Thus, for every report history:

$$U_t^{\alpha}(\alpha_t(\theta_t)) - U_t^{\alpha}(0) = U_t(\theta_t) - U_t(\underline{\theta}) \tag{75}$$

$$= \int_{\underline{\theta}}^{\theta_t} \frac{1}{\theta} \mathbb{E}\left[\sum_{s=t}^{\tau} \beta^{s-t} D_{t,s} \theta_s u_\theta(\theta^s; \theta_s) \middle| \theta \right] d\theta \tag{76}$$

$$= \int_{\underline{\theta}}^{\theta_t} U_{\theta,t}(\theta)\, d\theta \tag{77}$$

as claimed.

A.2 Proof of Proposition 2

Using the simplification that $c_1 = Rm_0 + y_1$, the Pareto problem is to minimise:

$$\int_{\Theta} \left\{ \tilde{c}_0(\theta_0) - \tilde{y}_0(\theta_0) + \tilde{m}_0(\theta_0) \right\} \pi(\theta_0)\, d\theta_0 \tag{78}$$

subject to given values for $U_0(\theta_0)$ and $U_{\theta,0}(\theta_0)$ at all $\theta_0 \in \Theta$.

[23] Recall α^t is measurable with respect to θ^t, and vice-versa.

This problem can be analysed piecewise, for each θ_0. The constraints are:

$$\begin{aligned} U_0(\theta_0) =& \frac{\tilde{c}_0(\theta_0)^{1-\sigma}-1}{1-\sigma} - \frac{1}{1+\nu}\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{1+\nu} \\ &+ \beta\int_\Theta \left[\frac{c_1^*(\tilde{m}_0(\theta_0),\theta_1)^{1-\sigma}-1}{1-\sigma} - \frac{1}{1+\nu}\left(\frac{y_1^*(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1}\right)^{1+\nu}\right] d\Pi(\theta_1|\theta_0) \end{aligned} \tag{79}$$

$$\theta_0 U_{\theta,0}(\theta_0) = \left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{1+\nu} + \beta\int_\Theta \rho(\theta_1|\theta_0)\left[\left(\frac{y_1^*(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1}\right)^{1+\nu}\right] d\Pi(\theta_1|\theta_0) \tag{80}$$

Letting $\mu(\theta_0)\pi(\theta_0)$ be the multiplier on (79) and $\mu_\theta(\theta_0)\pi(\theta_0)$ be the multiplier on (79), necessary conditions for cost minimisation are:

1. On $\tilde{c}_0(\theta_0)$:
$$-1+\mu(\theta_0)\tilde{c}_0(\theta_0)^{-\sigma} = 0 \tag{81}$$

2. On $\tilde{y}_0(\theta_0)$:
$$1-\mu(\theta_0)\frac{1}{\theta_0}\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{\nu} + \mu_\theta(\theta_0)\frac{1}{\theta_0}(1+\nu)\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{\nu} = 0 \tag{82}$$

3. On $\tilde{m}_0(\theta_0)$:
$$-1+\mu(\theta_0)\beta R\int_\Theta c_1^*(\tilde{m}_0(\theta_0),\theta_1)^{-\sigma} d\Pi(\theta_1|\theta_0) \tag{83}$$
$$+\mu_\theta(\theta_0)\beta(1+\nu)\int_\Theta\left[\rho(\theta_1|\theta_0)\frac{y_{m,1}^*(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1}\left(\frac{y_1^*(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1}\right)^{\nu}\right] d\Pi(\theta_1|\theta_0) = 0$$

where (83) has used the period-1 budget constraint (25), and the period-1 optimality condition (26). Equations (81) and (82) together give solutions for the multipliers:

$$\mu_\theta(\theta_0) = -\frac{1-\frac{\frac{1}{\theta_0}\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{\nu}}{\tilde{c}_0(\theta_0)^{-\sigma}}}{\frac{1}{\theta_0}(1+\nu)\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{\nu}} \tag{84}$$

$$\mu(\theta_0) = \tilde{c}_0(\theta_0)^{\sigma} \tag{85}$$

Using these in condition (83) we have:

$$\frac{\frac{\beta R \int_{\Theta} c_1^*(\tilde{m}_0(\theta_0),\theta_1)^{-\sigma} d\Pi(\theta_1|\theta_0)}{\tilde{c}_0(\theta_0)^{-\sigma}} - 1}{\beta \int_{\Theta} \left[\rho\left(\theta_1|\theta_0\right) \frac{y^*_{m,1}(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1} \left(\frac{y_1^*(\tilde{m}_0(\theta_0),\theta_1)}{\theta_1} \right)^{\nu} \right] d\Pi\left(\theta_1|\theta_0\right)} = \frac{1 - \frac{\frac{1}{\theta_0}\left(\frac{\bar{y}_0(\theta_0)}{\theta_0}\right)^{\nu}}{\tilde{c}_0(\theta_0)^{-\sigma}}}{\frac{1}{\theta_0}\left(\frac{\tilde{y}_0(\theta_0)}{\theta_0}\right)^{\nu}} \tag{86}$$

Using (28) and (29) in (86) gives the arbitrage formula (30). The main results then follow as set out in the main text.

A.3 Proof of Lemma 1

In period 1, optimal allocations for any given m_0 satisfy the envelope condition:

$$\begin{aligned}\frac{c_1^*\left(m_0,\theta_1'\right)^{1-\sigma}-1}{1-\sigma} - \frac{1}{1+\nu}\left(\frac{y_1^*\left(m_0,\theta_1'\right)}{\theta_1'}\right)^{1+\nu} =& \frac{c_1^*\left(m_0,\underline{\theta}\right)^{1-\sigma}-1}{1-\sigma} - \frac{1}{1+\nu}\left(\frac{y_1^*\left(m_0,\underline{\theta}\right)}{\underline{\theta}}\right)^{1+\nu} \\ &+ \int_{\underline{\theta}}^{\theta_1'} \frac{1}{\theta_1}\left(\frac{y_1^*\left(m_0,\theta_1\right)}{\theta_1}\right)^{1+\nu} d\theta_1\end{aligned} \tag{87}$$

Differentiating (87) with respect to m_0 and applying conditions (25) and (26):

$$c_1^*\left(m_0,\theta_1'\right)^{-\sigma} = c_1^*\left(m_0,\underline{\theta}\right)^{-\sigma} + (1+\nu)\int_{\underline{\theta}}^{\theta_1'} \frac{1}{\theta_1}\left\{\frac{y^*_{m,1}\left(m_0,\theta_1\right)}{\theta_1}\left(\frac{y_1^*\left(m_0,\theta_1\right)}{\theta_1}\right)^{\nu}\right\} d\theta_1 \tag{88}$$

Take any pair of types $\theta_1'' > \theta_1'$,. We have:

$$c_1^*\left(m_0,\theta_1''\right)^{-\sigma} - c_1^*\left(m_0,\theta_1'\right)^{-\sigma} = (1+\nu)\int_{\underline{\theta}}^{\theta_1'} \frac{1}{\theta_1}\left\{\frac{y^*_{m,1}\left(m_0,\theta_1\right)}{\theta_1}\left(\frac{y_1^*\left(m_0,\theta_1\right)}{\theta_1}\right)^{\nu}\right\} d\theta_1 \tag{89}$$

The right-hand side of (89) is negative when $y^*_{m,1}\left(m_0,\theta_1\right) < 0$ everywhere, which shows $c_1^*\left(m_0,\theta_1''\right)^{-\sigma}$ must be strictly decreasing when this is true.

Going the other way, if $c_1^*\left(m_0,\theta_1\right)^{-\sigma}$ is monotone in θ_1, then by Lebesgue's Theorem it is differentiable a.e., and its derivative is given by:

$$(1+\nu)\frac{1}{\theta_1}\left\{\frac{y^*_{m,1}\left(m_0,\theta_1\right)}{\theta_1}\left(\frac{y_1^*\left(m_0,\theta_1\right)}{\theta_1}\right)^{\nu}\right\}$$

This derivative must be negative, which can be true only if $y^*_{m,1}\left(m_0,\theta_1\right) < 0$, which must therefore hold a.e.. But this can only be true in the present setting if $\sigma > 0$, in which $y^*_{m,1}\left(m_0,\theta_1\right) < 0$ everywhere.

A.4 Proof of Proposition 3

The proof is a direct generalisation of the proof of Proposition 2. The Pareto problem is to minimise:

$$\int_{\Theta}\left\{\tilde{c}_0\left(\theta_0\right)-\tilde{y}_0\left(\theta_0\right)+\mathbb{E}\left[\sum_{t=1}^{\tau}R^{-t}\left[c_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right)-y_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right)\right]\middle|\theta_0\right]\right\}\pi\left(\theta_0\right)d\theta_0 \tag{90}$$

subject to given values for $U_0\left(\theta_0\right)$ and $U_{\theta,0}\left(\theta_0\right)$ at all $\theta_0\in\Theta$.

The constraints are:

$$U_0\left(\theta_0\right)=u\left(\tilde{c}_0\left(\theta_0\right),\tilde{y}_0\left(\theta_0\right);\theta_0\right)+\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t u\left(c_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right),y_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right);\theta_t\right)\middle|\theta_0\right] \tag{91}$$

$$\theta_0 U_{\theta,0}\left(\theta_0\right)=\theta_0 u_\theta\left(\tilde{c}_0\left(\theta_0\right),\tilde{y}_0\left(\theta_0\right);\theta_0\right)+\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t u_\theta\left(c_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right),y_t^*\left(\tilde{m}_0\left(\theta_0\right),\theta^t\right);\theta_t\right)\middle|\theta_0\right] \tag{92}$$

Letting $\mu\left(\theta_0\right)\pi\left(\theta_0\right)$ be the multiplier on (91) and $\mu_\theta\left(\theta_0\right)\pi\left(\theta_0\right)$ be the multiplier on (91), necessary conditions for cost minimisation at the benchmark allocation are:

1. On $\tilde{c}_0\left(\theta_0\right)$:
$$-1+\mu\left(\theta_0\right)u_{c,0}\left(\theta_0\right)+\mu_\theta\left(\theta_0\right)\theta_0 u_{c\theta,0}\left(\theta_0\right)=0 \tag{93}$$

2. On $\tilde{y}_0\left(\theta_0\right)$:
$$1+\mu\left(\theta_0\right)u_{y,0}\left(\theta_0\right)+\mu_\theta\left(\theta_0\right)\theta_0 u_{y\theta,0}\left(\theta_0\right)=0 \tag{94}$$

3. On $\tilde{m}_0\left(\theta_0\right)$:
$$\begin{aligned}&-\mathbb{E}\left[\sum_{t=1}^{\tau}R^{-t}\left[c_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)-y_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)\right]\middle|\theta_0\right] \\ &+\mu\left(\theta_0\right)\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)+u_{y,t}\left(\theta^t\right)y_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)\right]\middle|\theta_0\right] \\ &+\mu_\theta\left(\theta_0\right)\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)+u_{y\theta,t}\left(\theta^t\right)y_{m,t}^*\left(m_0\left(\theta_0\right),\theta^t\right)\right]\middle|\theta_0\right]=0\end{aligned} \tag{95}$$

The multipliers solve to give:

$$\mu_\theta(\theta_0)\,\theta_0 = -\frac{u_{c,0}(\theta_0) + u_{y,0}(\theta_0)}{u_{y\theta,0}(\theta_0)\,u_{c,0}(\theta_0) - u_{c\theta,0}(\theta_0)\,u_{y,0}(\theta_0)} \tag{96}$$

$$= -\frac{T'(\theta_0)}{u_{y\theta,0}(\theta_0) + (1 - T'(\theta_0))\,u_{c\theta,0}(\theta_0)} \tag{97}$$

$$\mu(\theta_0) = \frac{u_{c\theta,0}(\theta_0) + u_{y\theta,0}(\theta_0)}{u_{y\theta,0}(\theta_0)\,u_{c,0}(\theta_0) - u_{c\theta,0}(\theta_0)\,u_{y,0}(\theta_0)} \tag{98}$$

$$= \frac{1}{u_{c,0}(\theta_0)} \frac{u_{c\theta,0}(\theta_0) + u_{y\theta,0}(\theta_0)}{u_{y\theta,0}(\theta_0) + (1 - T'(\theta_0))\,u_{c\theta,0}(\theta_0)} \tag{99}$$

where the definition of the earnings tax (13) has been used in both cases. Using these in condition (95), and making use of (15), we have:

$$\left[\frac{u_{c\theta,0}(\theta_0) + u_{y\theta,0}(\theta_0)}{u_{y\theta,0}(\theta_0) + (1 - T'(\theta_0))\,u_{c\theta,0}(\theta_0)} - 1\right] + S'(\theta_0)\frac{u_{c\theta,0}(\theta_0) + u_{y\theta,0}(\theta_0)}{u_{y\theta,0}(\theta_0) + (1 - T'(\theta_0))\,u_{c\theta,0}(\theta_0)} \tag{100}$$

$$-\frac{T'(\theta_0)}{u_{y\theta,0}(\theta_0) + (1 - T'(\theta_0))\,u_{c\theta,0}(\theta_0)} \frac{\frac{1}{\theta_0}\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}(\theta^t)\,c^*_{m,t} + u_{y\theta,t}(\theta^t)\,y^*_{m,t}\right]\right]}{\mathbb{E}\left[\sum_{t=1}^{\tau} R^{-t}\left[c^*_{m,t}(m_0(\theta_0),\theta^t) - y^*_{m,t}(m_0(\theta_0),\theta^t)\right]\middle|\,\theta_0\right]} = 0$$

Rearranging gives expression (32), and (33) then follows by applying the Spence-Mirrlees condition and Assumption 1 to sign terms in the denominators on left- and right-hand sides of (32) respectively.

A.5 Proof of Proposition 4

To prove the result, it is sufficient to show that if $y^*_{m,t} \leq 0$ a.e., then $S'_c(\theta_0) > 0$ follows. The proof will work backwards from τ, given an arbitrary history $\theta^{\tau-1}$. Incentive compatibility in τ requires, for all $\theta'_\tau \in \Theta$:

$$\begin{aligned} u_{c,\tau}(\theta'_\tau)\,d^*_\tau(\theta'_\tau) =& u_{c,\tau}(\theta'_\tau)\,c^*_{m,\tau}(\theta'_\tau) + u_{y,\tau}(\theta'_\tau)\,y^*_{m,\tau}(\theta'_\tau) \\ & u_{c,\tau}(\underline{\theta})\,c^*_{m,\tau}(\underline{\theta}) + u_{y,\tau}(\underline{\theta})\,y^*_{m,\tau}(\underline{\theta}) \\ & + \int_{\underline{\theta}}^{\theta'_\tau}\left[u_{y\theta,\tau}(\theta_\tau)\,y^*_{m,\tau}(\theta_\tau)\right]d\theta_\tau \end{aligned} \tag{101}$$

Since $u_{y\theta,\tau}(\theta_\tau) > 0$, it follows that $u_{c,\tau}(\theta_\tau)\,d^*_\tau(\theta_\tau)$ is weakly decreasing in θ_τ. Now let the function $d'_\tau : \Theta \to \mathbb{R}_+$ be defined at any history node $\theta^{\tau-1}$ by:

$$d'_\tau(\theta_\tau) = \frac{1}{u_{c,\tau}(\theta_\tau)} \frac{\mathbb{E}_{\tau-1}\left[d^*_\tau(\theta_\tau)\right]}{\mathbb{E}_{\tau-1}\left[\frac{1}{u_{c,\tau}(\theta_\tau)}\right]} \tag{102}$$

This has the property that $u_{c,\tau}(\theta_\tau)\, d'_\tau(\theta_\tau)$ is a constant, and:

$$\mathbb{E}_{\tau-1}\left[d'_\tau(\theta_\tau)\right] = \mathbb{E}_{\tau-1}\left[d^*_\tau(\theta_\tau)\right] \tag{103}$$

Now consider the function:

$$\mathbb{E}_{\tau-1}\left[\beta\frac{u_{c,\tau}(\theta_\tau)}{u_{c,\tau-1}}\left(d^*_\tau(\theta_\tau) - d'_\tau(\theta_\tau)\right)\right] \tag{104}$$

Since $u_{c,\tau}(\theta_\tau)\, d^*_\tau(\theta_\tau)$ is weakly decreasing in θ_τ, $\frac{d^*_\tau(\theta_\tau)}{d'_\tau(\theta_\tau)}$ is likewise, and so the function $\left(d^*_\tau(\theta_\tau) - d'_\tau(\theta_\tau)\right)$ crosses zero once from above. Since it also has an expected value of zero, and $u_{c,\tau}(\theta_\tau)$ is decreasing in θ_τ, it follows that, at every θ^{t-1} node:

$$\mathbb{E}_{\tau-1}\left[\beta\frac{u_{c,\tau}(\theta_\tau)}{u_{c,\tau-1}}d^*_\tau(\theta_\tau)\right] \geq \mathbb{E}_{\tau-1}\left[\beta\frac{u_{c,\tau}(\theta_\tau)}{u_{c,\tau-1}}d'_\tau(\theta_\tau)\right], \tag{105}$$

$$= \beta\frac{1}{u_{c,\tau-1}}\frac{\mathbb{E}_{\tau-1}\left[d^*_\tau(\theta_\tau)\right]}{\mathbb{E}_{\tau-1}\left[\frac{1}{u_{c,\tau}(\theta_\tau)}\right]} \tag{106}$$

$$= R^{-1}\mathbb{E}_{\tau-1}\left[d^*_\tau(\theta_\tau)\right] \tag{107}$$

where the last line applies the IEE.

If $\tau = 1$ we are done. For $\tau > 1$ a recursive argument can be set up. For $1 < t \leq \tau$, suppose that at every node θ^{t-1} the following holds:

$$\mathbb{E}_{t-1}\left[\sum_{s=t}^{\tau}\beta^{s-t+1}\frac{u_{c,s}(\theta^s)}{u_{c,t-1}}d^*_s(\theta^s)\right] \geq \mathbb{E}_{t-1}\left[\sum_{s=t}^{\tau}R^{t-1-s}d^*_s(\theta^s)\right] \tag{108}$$

Using an identical argument to above, incentive compatibility after node θ^{t-2} implies:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}u_{c,s}(\theta^s)\, d^*_s(\theta^s)\,\middle|\,\theta^{t-2},\theta'_{t-1}\right] \\
=&\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\left[u_{c,s}(\theta^s)\, c^*_{m,s}(\theta^s) + u_{y,s}(\theta^s)\, y^*_{m,s}(\theta^s)\right]\,\middle|\,\theta^{t-2},\theta'_{t-1}\right] &(109)\\
=&\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\left[u_{c,s}(\theta^s)\, c^*_{m,s}(\theta^s) + u_{y,s}(\theta^s)\, y^*_{m,s}(\theta^s)\right]\,\middle|\,\theta^{t-2},\underline{\theta}\right] &(110)\\
&+\int_{\underline{\theta}}^{\theta'_t}\frac{1}{\theta_{t-1}}\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}D_{t-1,s}(\theta^s)\left[\theta_s u_{y\theta,s}(\theta^s)\, y^*_{m,s}(\theta^s)\right]\,\middle|\,\theta^{t-2},\theta_{t-1}\right]d\theta_t
\end{aligned}$$

Since $y^*_{m,s}(\theta^s)$ is weakly negative at all horizons, this implies the object:

$$\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}u_{c,s}(\theta^s)\,d^*_s(\theta^s)\,\middle|\,\theta^{t-2},\theta'_{t-1}\right]$$

is weakly decreasing in θ'_{t-1} at every θ^{t-2} history node. Now take a function $d'_{t-1}:\Theta\to\mathbb{R}_+$, defined by:

$$d'_{t-1}(\theta_{t-1}) := \frac{1}{u_{c,t-1}(\theta_{t-1})}\frac{\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\frac{u_{c,s}(\theta^s)}{u_{c,t-1}(\theta_{t-1})}d^*_s(\theta^s)\,\middle|\,\theta^{t-2}\right]}{\mathbb{E}\left[\frac{1}{u_{c,t-1}(\theta_{t-1})}\,\middle|\,\theta^{t-2}\right]} \tag{111}$$

We have that the object:

$$\frac{\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\frac{u_{c,s}(\theta^s)}{u_{c,t-1}(\theta_{t-1})}d^*_s(\theta^s)\,\middle|\,\theta^{t-2},\theta'_{t-1}\right]}{d'_{t-1}(\theta_{t-1})} \tag{112}$$

is weakly decreasing in θ_{t-1}, and so the function:

$$\left(\mathbb{E}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\frac{u_{c,s}(\theta^s)}{u_{c,t-1}(\theta_{t-1})}d^*_s(\theta^s)\,\middle|\,\theta^{t-2},\theta_{t-1}\right]-d'_{t-1}(\theta_{t-1})\right) \tag{113}$$

crosses zero once from above in θ_{t-1}, and has zero expectation given θ^{t-2}. Thus, since $u_{c,t-1}(\theta_{t-1})$ is decreasing in θ_{t-1}:

$$\mathbb{E}\left[\frac{\beta u_{c,t-1}(\theta_{t-1})}{u_{c,t-2}}\left(\sum_{s=t-1}^{\tau}\beta^{s-t+1}\frac{u_{c,s}(\theta^s)}{u_{c,t-1}(\theta_{t-1})}d^*_s(\theta^s)-d'_{t-1}(\theta_{t-1})\right)\,\middle|\,\theta^{t-2}\right]\geq 0 \tag{114}$$

So:

$$\begin{aligned}
&\mathbb{E}_{t-2}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+2}\frac{u_{c,s}\left(\theta^{s}\right)}{u_{c,t-2}\left(\theta_{t-1}\right)}d_{s}^{*}\left(\theta^{s}\right)\right] \\
&\geq \mathbb{E}_{t-2}\left[\frac{\beta u_{c,t-1}\left(\theta_{t-1}\right)}{u_{c,t-2}}d_{t-1}'\left(\theta_{t-1}\right)\right] \\
&= R^{-1}\mathbb{E}_{t-2}\left[\sum_{s=t-1}^{\tau}\beta^{s-t+1}\frac{u_{c,s}\left(\theta^{s}\right)}{u_{c,t-1}\left(\theta_{t-1}\right)}d_{s}^{*}\left(\theta^{s}\right)\right] \\
&= R^{-1}\mathbb{E}_{t-2}\left[d_{t-1}^{*}\left(\theta_{t-1}\right)+\mathbb{E}_{t-1}\left[\sum_{s=t}^{\tau}\beta^{s-t+1}\frac{u_{c,s}\left(\theta^{s}\right)}{u_{c,t-1}\left(\theta_{t-1}\right)}d_{s}^{*}\left(\theta^{s}\right)\right]\right] \\
&\geq R^{-1}\mathbb{E}_{t-2}\left[d_{t-1}^{*}\left(\theta_{t-1}\right)+\mathbb{E}_{t-1}\left[\sum_{s=t}^{\tau}R^{t-1-s}d_{s}^{*}\left(\theta^{s}\right)\right]\right] \\
&= \mathbb{E}_{t-2}\left[\sum_{s=t-1}^{\tau}R^{t-2-s}d_{s}^{*}\left(\theta^{s}\right)\right]
\end{aligned} \tag{115}$$

where the second line makes use of the inverse Euler equation between $t-2$ and $t-1$, and the penultimate line applies (108). Thus when (108) holds at all nodes in $t-1$, it also holds at all nodes in $t-2$. By induction, it thus holds in period 0, and directly implies $S_{c}'\left(\theta_{0}\right)\geq 0$.

A.6 Proof of Lemma 3

The expenditure minimisation problem yields the following optimality conditions:

$$u_{c,0}\left(\theta_{0}\right)\left(a-T'\left(y_{0}\right)\right)+u_{y,0}\left(\theta_{0}\right)=0 \tag{116}$$

$$u_{c,0}\left(\theta_{0}\right)\left(p+S'\left(m_{0}\right)\right)-\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^{t}\left[u_{c,t}\left(\theta^{t}\right)c_{m,t}^{*}+u_{y,t}\left(\theta^{t}\right)y_{m,t}^{*}\right]\middle|\theta_{0}\right]=0 \tag{117}$$

Coombine a differential change to θ_0 by $\frac{dy_0}{dp}$ units and to p by $-\frac{dy_0}{d\theta_0}$ units (all derivatives with respect to prices are compensated demand changes, as defined in the main text). Differentiating the condition (116), this joint change gives:

$$\left[u_{cc,0}\left(a-T'\left(y_{0}\right)\right)+u_{cy,0}\right]\left[\frac{dc_0}{d\theta_0}\frac{dy_0}{dp}-\frac{dc_0}{dp}\frac{dy_0}{d\theta_0}\right]+\left[u_{c\theta,0}\left(a-T'\left(y_{0}\right)\right)+u_{y\theta,0}\right]\frac{dy_0}{dp}=0 \tag{118}$$

Now apply similar analysis to the savings optimality condition (117), assuming a change to θ_0 by $\frac{dm_0}{dp}$

units and p by $-\frac{dm_0}{d\theta_0}$:

$$0 = -u_{c,0}\frac{dm_0}{d\theta_0} + \left\{ u_{c\theta,0}\left(p + S'\left(m_0\right)\right) - \frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]\right\}\frac{dm_0}{dp} \tag{119}$$
$$+ u_{cc,0}\left(p + S'\left(m_0\right)\right)\left[\frac{dc_0}{d\theta_0}\frac{dm_0}{dp} - \frac{dc_0}{dp}\frac{dm_0}{d\theta_0}\right]$$
$$+ u_{cy,0}\left(p + S'\left(m_0\right)\right)\left[\frac{dy_0}{d\theta_0}\frac{dm_0}{dp} - \frac{dy_0}{dp}\frac{dm_0}{d\theta_0}\right]$$

Combining the two previous conditions:

$$\begin{aligned}
&\left\{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right] - u_{c\theta,0}\left(p + S'\left(m_0\right)\right)\right\}\frac{dm_0}{dp} \qquad (120)\\
&+ \left[u_{c\theta,0}\left(a - T'\left(y_0\right)\right) + u_{y\theta,0}\right]\frac{dy_0}{dp} + u_{c,0}\frac{dm_0}{d\theta_0}\\
=&u_{cc,0}\left\{\left(p + S'\left(m_0\right)\right)\left[\frac{dc_0}{d\theta_0}\frac{dm_0}{dp} - \frac{dc_0}{dp}\frac{dm_0}{d\theta_0}\right] - \left(a - T'\left(y_0\right)\right)\left[\frac{dc_0}{d\theta_0}\frac{dy_0}{dp} - \frac{dc_0}{dp}\frac{dy_0}{d\theta_0}\right]\right\}\\
&+ u_{cy,0}\left\{\left(p + S'\left(m_0\right)\right)\left[\frac{dy_0}{d\theta_0}\frac{dm_0}{dp} - \frac{dy_0}{dp}\frac{dm_0}{d\theta_0}\right] - \left[\frac{dc_0}{d\theta_0}\frac{dy_0}{dp} - \frac{dc_0}{dp}\frac{dy_0}{d\theta_0}\right]\right\}
\end{aligned}$$

Differentiating the budget constraint gives:

$$\frac{dy_0}{d\theta_0}\left(a - T'\left(y_0\right)\right) = \frac{dc_0}{d\theta_0} + \left(p + S'\left(m_0\right)\right)\frac{dm_0}{d\theta_0} \tag{121}$$
$$\frac{dy_0}{dR^{-1}}\left(a - T'\left(y_0\right)\right) = \frac{dc_0}{dp} + \left(p + S'\left(m_0\right)\right)\frac{dm_0}{dp} \tag{122}$$

(recall the change in p is compensated). So, eliminating derivatives in c_0, the coefficient on $u_{cc,0}$ can be rewritten as:

$$\begin{aligned}
&\left(p + S'\left(m_0\right)\right)\left(\frac{dy_0}{d\theta_0}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{d\theta_0}\right)\frac{dm_0}{dp} \qquad (123)\\
-&\left(p + S'\left(m_0\right)\right)\left(\frac{dy_0}{dp}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{dp}\right)\frac{dm_0}{d\theta_0}\\
-&\left(a - T'\left(y_0\right)\right)\left(\frac{dy_0}{d\theta_0}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{d\theta_0}\right)\frac{dy_0}{dp}\\
+&\left(a - T'\left(y_0\right)\right)\left(\frac{dy_0}{dp}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{dp}\right)\frac{dy_0}{d\theta_0}\\
&=0
\end{aligned}$$

(the equality by simple cancellation of terms). The coefficient on $u_{cy,0}$ is:

$$\begin{aligned}
&\left(p + S'\left(m_0\right)\right)\left[\frac{dy_0}{d\theta_0}\frac{dm_0}{dp} - \frac{dy_0}{dp}\frac{dm_0}{d\theta_0}\right] \\
&- \left(\frac{dy_0}{d\theta_0}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{d\theta_0}\right)\frac{dy_0}{dp} \\
&+ \left(\frac{dy_0}{dR^{-1}}\left(a - T'\left(y_0\right)\right) - \left(p + S'\left(m_0\right)\right)\frac{dm_0}{dp}\right)\frac{dy_0}{d\theta_0} \\
&= 0
\end{aligned} \tag{124}$$

Hence the result:

$$\begin{aligned}
-u_{c,0}\frac{dm_0}{d\theta_0} &= \left[u_{c\theta,0}\left(a - T'\left(y_0\right)\right) + u_{y\theta,0}\right]\frac{dy_0}{dp} \\
&+ \left\{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right] - u_{c\theta,0}\left(p + S'\left(m_0\right)\right)\right\}\frac{dm_0}{dp}
\end{aligned} \tag{125}$$

Similar steps for the derivative with respect to a give the two conditions:

$$-u_{c,0}\frac{dy_0}{d\theta_0} + \left[u_{cc,0}\left(a - T'\left(y_0\right)\right) + u_{cy,0}\right]\left[\frac{dc_0}{d\theta_0}\frac{dy_0}{da} - \frac{dc_0}{da}\frac{dy_0}{d\theta_0}\right] + \left[u_{c\theta,0}\left(a - T'\left(y_0\right)\right) + u_{y\theta,0}\right]\frac{dy_0}{da} = 0 \tag{126}$$

$$\begin{aligned}
0 &= \left\{u_{c\theta,0}\left(p + S'\left(m_0\right)\right) - \frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]\right\}\frac{dm_0}{da} \\
&+ u_{cc,0}\left(p + S'\left(m_0\right)\right)\left[\frac{dc_0}{d\theta_0}\frac{dm_0}{da} - \frac{dc_0}{da}\frac{dm_0}{d\theta_0}\right] + u_{cy,0}\left(p + S'\left(m_0\right)\right)\left[\frac{dy_0}{d\theta_0}\frac{dm_0}{da} - \frac{dy_0}{da}\frac{dm_0}{d\theta_0}\right]
\end{aligned} \tag{127}$$

These sum and cancel as before, to give:

$$\begin{aligned}
u_{c,0}\frac{dy_0}{d\theta_0} &= \left[u_{c\theta,0}\left(a - T'\left(y_0\right)\right) + u_{y\theta,0}\right]\frac{dy_0}{da} \\
&+ \left\{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right] - u_{c\theta,0}\left(p + S'\left(m_0\right)\right)\right\}\frac{dm_0}{da}
\end{aligned} \tag{128}$$

By using Slutsky symmetry, the definition of the compensated elasticities and the consumer first-order conditions, conditions (125) and (128) can be rearranged to give the result in the text – for which $p = a = 1$ has been imposed.

A.7 Proof of Proposition 5

Equations (52) and (53) can be written in matrix form and (given strict concavity) inverted to give:

$$\frac{1}{\varepsilon_{yy}\varepsilon_{mm}-\varepsilon_{ym}\varepsilon_{my}}\left[\varepsilon_{mm}\frac{1}{y_0}\frac{dy_0}{d\theta_0}-\varepsilon_{ym}\frac{1}{m_0}\frac{dm_0}{d\theta_0}\right]=\left[\frac{u_{c\theta,0}}{u_{c,0}}-\frac{u_{y\theta,0}}{u_{y,0}}\right] \tag{129}$$

$$\frac{1}{\varepsilon_{yy}\varepsilon_{mm}-\varepsilon_{ym}\varepsilon_{my}}\left[\varepsilon_{yy}\frac{1}{m_0}\frac{dm_0}{d\theta_0}-\varepsilon_{my}\frac{1}{y_0}\frac{dy_0}{d\theta_0}\right]=\left[\frac{u_{c\theta,0}}{u_{c,0}}-\frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t}+u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t}+u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}\right] \tag{130}$$

Strict monotonicity in choices implies the derivatives can be replaced with ratios of densities a.e.. We have:

$$\left[\varepsilon_{yy}\frac{1}{m_0}\frac{dm_0}{d\theta_0}-\varepsilon_{my}\frac{1}{y_0}\frac{dy_0}{d\theta_0}\right]=\frac{\pi\left(\theta_0\right)}{1-\Pi\left(\theta_0\right)}\frac{1}{\alpha^y\left(y_0\right)\alpha^m\left(m_0\right)}\left[\varepsilon_{yy}\alpha^y\left(y_0\right)-\varepsilon_{my}\alpha^m\left(m_0\right)\right] \tag{131}$$

$$\left[\varepsilon_{yy}\frac{1}{m_0}\frac{dm_0}{d\theta_0}-\varepsilon_{my}\frac{1}{y_0}\frac{dy_0}{d\theta_0}\right]=\frac{\pi\left(\theta_0\right)}{1-\Pi\left(\theta_0\right)}\frac{1}{\alpha^y\left(y_0\right)\alpha^m\left(m_0\right)}\left[\varepsilon_{mm}\alpha^m\left(m_0\right)-\varepsilon_{ym}\alpha^y\left(u_0\right)\right] \tag{132}$$

and:

$$\frac{\pi\left(\theta_0\right)}{1-\Pi\left(\theta_0\right)}\frac{1}{\alpha^y\left(y_0\right)\alpha^m\left(m_0\right)}>0 \tag{133}$$

$$\frac{1}{\varepsilon_{yy}\varepsilon_{mm}-\varepsilon_{ym}\varepsilon_{my}}<0 \tag{134}$$

the latter by strict concavity of the expenditure function. The results follow directly.

A.8 Proof of Proposition 6

Optimal tax arbitrage gives:

$$\frac{T'\left(y_0\right)}{u_{y\theta,0}\left(\theta_0\right)+\left(1-T'\left(y_0\right)\right)u_{c\theta,0}\left(\theta_0\right)} \tag{135}$$

$$=\frac{S'\left(m_0\right)}{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t}+u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]-\left(1+S'\left(m_0\right)\right)u_{c\theta,0}\left(\theta_0\right)} \tag{136}$$

$$\equiv\Gamma\left(\theta_0\right) \tag{137}$$

(defining a constant $\Gamma(\theta_0)$). Using this in (52) and (53), optimality requires:

$$\frac{1}{y_0}\frac{dy_0}{d\theta_0} = \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{u_{y\theta,0}}{u_{y,0}}\right]\varepsilon_{yy} + \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t} + u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\middle|\,\theta_0\right]}\right]\varepsilon_{ym} \tag{138}$$

$$\frac{1}{m_0}\frac{dm_0}{d\theta_0} = \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{u_{y\theta,0}}{u_{y,0}}\right]\varepsilon_{my} + \left[\frac{u_{c\theta,0}}{u_{c,0}} - \frac{\frac{1}{\theta_0}\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t D_{0,t}\theta_t\left[u_{c\theta,t}\left(\theta^t\right)c^*_{m,t} + u_{y\theta,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}{\mathbb{E}_0\left[\sum_{t=1}^{\tau}\beta^t\left[u_{c,t}\left(\theta^t\right)c^*_{m,t} + u_{y,t}\left(\theta^t\right)y^*_{m,t}\right]\right]}\right]\varepsilon_{mm} \tag{139}$$

$$\frac{1}{y_0}\frac{dy_0}{d\theta_0}u_{c,0}\Gamma(\theta_0) = \frac{T'(y_0)}{(1-T'(y_0))}\varepsilon_{yy} - \frac{S'(m_0)}{(1+S'(m_0))}\varepsilon_{ym} \tag{140}$$

$$\frac{1}{m_0}\frac{dm_0}{d\theta_0}u_{c,0}\Gamma(\theta_0) = \frac{T'(y_0)}{(1-T'(y_0))}\varepsilon_{my} - \frac{S'(m_0)}{(1+S'(m_0))}\varepsilon_{mm} \tag{141}$$

Eliminating type derivatives as in the Proof of Proposition 5:

$$\frac{1-\Pi^y(y_0)}{y_0\cdot\pi^y(y_0)}\left[u_{c,0}\Gamma(\theta_0)\frac{\pi(\theta_0)}{1-\Pi(\theta_0)}\right] = \frac{T'(y_0)}{(1-T'(y_0))}\varepsilon_{yy} - \frac{S'(m_0)}{(1+S'(m_0))}\varepsilon_{ym} \tag{142}$$

$$\frac{1-\Pi^m(m_0)}{m_0\cdot\pi^m(m_0)}\left[u_{c,0}\Gamma(\theta_0)\frac{\pi(\theta_0)}{1-\Pi(\theta_0)}\right] = \frac{T'(y_0)}{(1-T'(y_0))}\varepsilon_{my} - \frac{S'(m_0)}{(1+S'(m_0))}\varepsilon_{mm} \tag{143}$$

and the result follows, with $\Phi(\theta_0) := u_{c,0}\Gamma(\theta_0)\frac{\pi(\theta_0)}{1-\Pi(\theta)}$.