

DISCUSSION PAPER SERIES

DP16254

**Static and Dynamic Mirrleesian Taxation
with Non-separable Preferences: A
Unified Approach**

Christian Hellwig

MACROECONOMICS AND GROWTH

PUBLIC ECONOMICS

CEPR

Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach

Christian Hellwig

Discussion Paper DP16254

Published 14 June 2021

Submitted 10 June 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Macroeconomics and Growth
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Christian Hellwig

Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach

Abstract

Abstract I analyze dynamic Mirrlees taxation with preferences that are non-separable between consumption, leisure and type, which determines both ability and consumption needs. I show how to account for non-separable preferences through a simple change in probability measures. I generalize the existing Inverse Euler Equation and optimal static labor tax formulae and provide a unified intuition based on a set of perturbations around the optimal allocations that preserve expected utility and incentive compatibility. Non-separability in preferences gives rise to a new tradeoff between current and future redistribution that is internalized by the planner's solution but not by private savings decisions. This leads to a novel rationale to subsidize (tax) savings and make labor taxes more (less) persistent, when more productive agents also have higher (lower) consumption needs.

JEL Classification: N/A

Keywords: N/A

Christian Hellwig - christian.hellwig@tse-fr.eu
Toulouse School of Economics and CEPR

Acknowledgements

I thank Fabrice Collard, Tuuli Vanhapelto and Nicolas Werquin for useful discussions and seminar participants at Toulouse School of Economics for comments. Funding from the French National Research Agency (ANR) under grant ANR-17-EURE-0010 (Investissements d'Avenir) is gratefully acknowledged.

Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach*

Christian Hellwig
Toulouse School of Economics

June 10, 2021

Abstract

I analyze dynamic Mirrlees taxation with preferences that are non-separable between consumption, leisure and type, which determines both ability and consumption needs. I show how to account for non-separable preferences through a simple change in probability measures. I generalize the existing Inverse Euler Equation and optimal static labor tax formulae and provide a unified intuition based on a set of perturbations around the optimal allocations that preserve expected utility and incentive compatibility. Non-separability in preferences gives rise to a new tradeoff between current and future redistribution that is internalized by the planner's solution but not by private savings decisions. This leads to a novel rationale to subsidize (tax) savings and make labor taxes more (less) persistent, when more productive agents also have higher (lower) consumption needs.

1 Introduction

How should optimal taxes be structured to balance redistribution motives and efficiency distortions? Starting with the seminal work of Mirrlees (1971), the existing literature on optimal income taxes captures the efficiency-redistribution tradeoff through a basic asymmetric information friction: workers have private information about their ability or inclination to work.¹ Dynamic life cycle economies augment the tax design problem by a dynamic insurance motive through stochastic evolution of individual types (i.e. abilities or preferences) and aggregate savings.

*I thank Fabrice Collard, Tuuli Vanhapelto and Nicolas Werquin for useful discussions and seminar participants at Toulouse School of Economics for comments. Funding from the French National Research Agency (ANR) under grant ANR-17-EURE-0010 (Investissements d'Avenir) is gratefully acknowledged.

¹see Mirrlees (1971), Diamond (1998), Saez (2001), among many others.

Existing results on dynamic optimal tax design leverage an assumption that preferences are separable, and in particular that the marginal utility of consumption is independent of leisure or ability types, into a sharp characterization of optimal labor and savings distortions.² Optimal savings distortions are characterized by the Inverse Euler Equation, which captures the idea that wealth accumulation has adverse effects on future incentives to work (see e.g. Diamond and Mirrlees, 1978, and Rogerson, 1985, and Golosov, Kocherlakota and Tsyvinski, 2003). In parallel, Farhi and Werning (2013) and Golosov, Troshkin and Tsyvinski (2016) have characterized optimal labor wedges emphasizing smoothing and backloading of tax distortions.

The separability assumption plays a key role in limiting the interaction between static redistribution and dynamic savings margins, which greatly facilitates the characterization of each of the two margins separately. At the same time it imposes strong restrictions on substitution between home production and market work, across periods or wealth effects on labor supply, which in turn limit the scope of applicability for many questions of applied interest involving life-cycle savings, home production, taxation of couples, or insurance against important life cycle risks. Aguiar and Hurst (2005, 2007) document patterns of substitution between market and non-market work and consumption in both the time series and cross section which directly contradict this hypothesis.³

I characterize optimal labor and savings wedges in a dynamic Mirrleesian economy with *arbitrary* non-separability in preferences between consumption, types and leisure.⁴ I show three main results:

First, I identify a new rationale for savings taxes or subsidies based on the effect of savings on *current* redistribution. This rationale is conceptually distinct from the wealth effect of savings on future incentives that was captured by the original Inverse Euler Equation, and it disappears when preferences are separable. Furthermore, when it is optimal to subsidize savings, the optimal allocation features mean-reverting social mobility and overturns the immiseration property that often characterizes optimal social mobility in dynamic private information economies.⁵

²see e.g. Kocherlakota, Golosov and Tsyvinski (2003), Kocherlakota (2005), Albanesi and Sleet (2006), Golosov and Tsyvinski (2006), Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2016), or Stantcheva (2018) for important contributions and Stantcheva (2020) for a recent review of the literature.

³see also Aguiar, Hurst and Karabarbounis (2012) for a review and further discussion of consumption patterns and intra- and inter-temporal substitution between time use and consumption goods at the micro level. Benhabib, Rogerson and Wright (1991) have shown how allowing for substitution between home and market production improves upon some key predictions of standard real business cycle models.

⁴Grochulski and Kocherlakota (2010), Köhne and Kuhn (2015) and Köhne (2018) extend the dynamic Mirrlees model to allow for non-time separable preferences through the introduction of habits or durable consumption goods, but they retain the separability of consumption and effort.

⁵See for example Thomas and Worrall (1990), or Atkeson and Lucas (1992).

Savings affect current information rents and redistribution both directly and indirectly because a change in future redistribution also feeds back into current incentives. These two effects go in opposite directions but the direct effect dominates unless types are highly persistent, and in expectation, the effect of savings on current redistribution always dominates the wealth effect of savings on future incentives. If consumption needs or marginal utilities are increasing with the worker's ability type, savings reduce current information rents and facilitate current redistribution: this makes it optimal to subsidize savings. When instead low-ability types have higher consumption needs, it becomes optimal to tax savings. The magnitude of this savings wedge is increasing in the optimal concurrent labor tax and decreasing in the persistence of types.

Second, I provide a new characterization of optimal labor wedges based on an arbitrage between redistribution of consumption and redistribution of leisure. This arbitrage highlights that the optimal allocation doesn't just trade off between incentives and redistribution, but also between the channels through which redistribution takes place: is it more efficient to tax top income earners by asking them to work harder, or by asking them to consume less?

The optimal allocation equalizes the marginal cost of redistribution to the marginal benefit of redistribution through consumption and to the marginal benefit of redistribution through leisure, for each type, i.e. at the optimum, the planner is indifferent between increasing the work load, or reducing the consumption, of the most productive types. This simple observation leads to a natural complement of existing formulae for optimal income taxes from the static Mirrlees model (e.g. Diamond 1998, Saez, 2001) or the dynamic tax formula in Golosov, Troshkin and Tsyvinski (2016), which applies in both static and dynamic economies, and is easy to link to sufficient statistics from the distribution of earnings and consumption.

The marginal cost of redistribution is given by the ratio of the marginal efficiency loss from the labor distortion to the reduction in marginal information rents for a given type. The marginal benefits of redistribution consist of a static component and a dynamic component. The static component is determined by perturbations that transfer consumption or leisure from higher to lower types while preserving expected utility and incentive compatibility. The additional dynamic benefit of tax distortions arises because a commitment to future labor taxes and redistribution also reduces current information rents, thus facilitating current redistribution. The optimal labor wedge arbitrages distortions over time, while internalizing that future labor distortions generate an extra benefit from raising future redistribution. As in Farhi and Werning (2013), the optimal labor taxes are persistent and backloaded over time or with age.

Third, when preferences are non-separable, labor tax-smoothing interacts with the optimal

savings wedge, since the latter determines the planner’s discount rate between current and future costs and benefits fo redistribution. Taxing or subsidizing savings to increase current redistribution goes hand in hand with increasing or reducing the persistence of labor tax distortions. In addition, the feedback from future to current incentives and redistribution also affects optimal persistence of labor taxes directly.

All my results derive from the following key technical insight: Non-separability in preferences enters the characterization of optimal allocations through a change in the probability measure of types, or Radon-Nikodym derivative, that can be interpreted as an “incentive adjustment” to Pareto weights that is required to respect incentive compatibility. If a local perturbation of consumption allocations relaxes (tightens) incentive compatibility constraints for higher types, then the perturbation increases (reduces) the scope for further redistribution. The change in probability measures captures this through a first-order shift of probability weights towards lower (higher) types. This effect disappears with separable preferences because local perturbations of consumption impact incentive constraints only locally, so the original and incentive-adjusted Pareto weights are the same. A similar change in probabilities results from perturbations to leisure, but the latter always shifts weights towards lower ability types. The change in probability measure appears in some existing characterizations of the optimal labor wedge, but doesn’t appear to have been interpreted as such.⁶

This characterization of optimal allocations in turn allows me to interpret optimal labor and savings wedges through elementary perturbations that capture the marginal costs and benefits of redistribution, or the private and social marginal returns to savings, thus providing a unified intuition and generalization of many existing results to arbitrary preferences: Optimal savings wedges are characterized by a generalized Inverse Euler Equation which incorporates the wealth effect of savings on future incentives through the change in probabilities, and adds a static savings wedge to account for the effect of savings on current period incentives. Optimal labor wedges are captured by a dual representation of the costs and benefits of redistribution through either consumption and leisure, which are each governed by their respective change in probabilities. Their combination highlights a novel two-way interaction that arises from concerns for redistribution that are internalized by the planner’s solution but not by private decisions. The magnitude of the optimal savings wedge is linked to the magnitude of current labor wedges since the latter is used to “price” current redistribution, and the persistence of labor wedges is linked to the sign and level of the

⁶See, e.g. Hellwig (2007) or for the static Mirrlees model with non-separable preferences, or Golosov, Troskin and Tsyvinski (2016) for a similar characterization in the dynamic model with non-separable preferences.

savings wedge since the latter determines the optimal tradeoff between current and future resources and efficiency distortions.

From an applied perspective, these results substantially expand the applicability of optimal taxation results. Whenever households tradeoff between market and non-market work or consumption, or when labor supply decisions coordinate multiple household members across different activities and interact with life-cycle savings, it is important to know whether non-separability of preferences makes it easier or harder to provide insurance without distorting incentives. The present analysis answers this question by showing that the tradeoff shifts towards positive savings taxes, lower but more persistent labor taxes, and more redistribution, when needs-based and ability-based redistribution motives are aligned (low types have low ability and high needs), but towards savings subsidies, higher but less persistent labor taxes, and less redistribution in the opposite case (low ability types also have low consumption needs). Understanding how different redistribution motives interact thus becomes crucially important for any quantitative assessment of dynamic tax and insurance policies.

Section 2 introduces the model. Section 3 develops the core idea of incentive-adjusted probabilities. Section 4 reviews the benchmark results with separable preferences. Section 5 discusses optimal savings distortions, section 6 optimal labor wedges. Section 7 concludes with examples.

2 The Model

There are $T < \infty$ periods, and there is a measure-1 continuum of agents. Each agent draws a type sequence $\theta^T \in [\underline{\theta}, \bar{\theta}]^T$ iid across agents according to cdf $F(\cdot)$ (and pdf $f(\cdot)$) at date 0 (prior to the first period), and $\theta_t \in [\underline{\theta}, \bar{\theta}]$ is then revealed to the agent at the beginning of period t ; in period t , the agent thus knows θ^t . Agents' preferences over consumption and earnings are defined by

$$\mathbb{E}_0 \left\{ \sum_{t=1}^T \beta^{t-1} U(C_t, Y_t; \theta_t) \right\}$$

i.e. an agent's utility in period t is a function of his current type θ_t , current consumption C_t and current earnings Y_t . The function U is twice continuously differentiable, $U_C > 0$, $U_{CC} < 0$, $U_Y < 0$, $U_{YY} < 0$, $U_\theta > 0$ and U otherwise satisfies the usual Inada conditions as C or Y approach 0 or ∞ . Consumption and earnings are assumed to be observable but individual types are the agents' private information. Types are private information, but the social planner keeps a record of the agents' past announcements and can borrow or save with a return $R > 0$.

I assume that $f(\cdot)$ is positive over the interior of its support, that $f(\theta_t | \theta^{t-1})$ only depends on

t and θ_{t-1} (i.e. θ^t follows a time-dependent Markov chain), conditional densities $f(\theta_t|\theta^{t-1})$ are differentiable w.r.t. θ_{t-1} , and that "types are persistent", i.e.

$$\mathcal{J}(\theta_t, \theta^{t-1}) \equiv \frac{\partial f(\theta_t|\theta^{t-1})}{\partial \theta_{t-1}} \frac{1}{f(\theta_t|\theta^{t-1})}$$

is non-decreasing in θ_t . This formulation includes types that follow a generic linear AR(1) process in which $\frac{\partial f(\theta_t|\theta^{t-1})}{\partial \theta_{t-1}} = -\rho \frac{\partial f(\theta_t|\theta^{t-1})}{\partial \theta_t}$, with useful polar cases $\rho = 0$ (types are independent across time, no persistence) and $\rho = 1$ (types follow a random walk, full persistence).

I make a few additional assumptions about $U(C, Y; \theta)$. The first is the *Strict Single-Crossing Condition*: $-U_Y(C, Y; \theta)/U_C(C, Y; \theta)$ is strictly decreasing in θ for all $(C, Y; \theta)$, or

$$\frac{U_{C\theta}}{U_C} - \frac{U_{Y\theta}}{U_Y} > 0.$$

This assumption guarantees monotonicity of any incentive-compatible allocation in the static Mirrlees screening problem: on the margin, higher types are more willing to work.

Next, I assume that both consumption and leisure (or earnings) are normal goods, or that the elasticities

$$\begin{aligned} \mathcal{E}_C &\equiv \frac{\partial \ln(-U_Y/U_C)}{\partial \ln C} = - \left(\frac{U_{CC}}{U_C} - \frac{U_{CY}}{U_Y} \right) C \\ \mathcal{E}_Y &\equiv \frac{\partial \ln(-U_Y/U_C)}{\partial \ln Y} = \left(\frac{U_{YY}}{U_Y} - \frac{U_{CY}}{U_C} \right) Y. \end{aligned}$$

are both non-negative for all $(C, Y; \theta)$. These two elasticities will play a key role in my analysis.

Finally, I make assumptions about the planner's motives for redistribution. The signs of $U_{C\theta}$ and $U_{Y\theta}$ play a key role in determining how non-separabilities affect the optimal tax design problem. For illustrative purposes, I will occasionally use the following class of "weakly separable" preferences:

$$U(C, Y; \theta) = \mathcal{U}(\gamma(C, \underline{C}(\theta)) - n(Y, \theta))$$

The utility aggregator $\mathcal{U}(\cdot)$ is strictly increasing and concave. The function $n(Y, \theta)$ can be interpreted as (a disutility of) hours worked to generate earnings Y .⁷ I assume that $n_Y > 0$, $n_\theta < 0$, and $n_{\theta Y} < 0$, so that θ can be interpreted as the agent's labor productivity or disutility of effort, with higher types being more productive. This formulation captures redistribution of effort from less to more productive agents, or equivalently, redistribution of leisure towards less productive agents, i.e. redistribution "*from each according to his ability*".

⁷While it is convenient for the analysis to define preferences in terms of the observables C and Y , it is straightforward to map the type-contingent preference over earnings into a preference over leisure or hours worked. I will thus refer interchangeably to redistribution of Y as redistribution of earnings or leisure.

The function $\gamma(C, \underline{C}(\theta))$ can be interpreted as a consumption index that includes a type-dependent subsistence consumption $\underline{C}(\theta)$. I assume that $\gamma_C > 0$, $\gamma_{\underline{C}} \leq 0$, and $\gamma_{\underline{C}C} \geq 0$. An increase in subsistence consumption $\underline{C}(\theta)$ then lowers the agent's utility and increases their marginal utility. This formulation thus introduces redistribution of consumption towards those with the highest marginal utilities, i.e. redistribution “to each according to his needs”.

With these assumptions, the cross-partials $U_{C\theta}$ and $U_{Y\theta}$ take the following form:

$$\begin{aligned}\frac{U_{C\theta}}{U_C} &= \frac{\mathcal{U}''}{\mathcal{U}'} (\gamma_{\underline{C}} \underline{C}'(\cdot) - n_\theta) + \frac{\gamma_{\underline{C}C}}{\gamma_C} \underline{C}'(\cdot) \\ \frac{U_{Y\theta}}{U_Y} &= \frac{\mathcal{U}''}{\mathcal{U}'} (\gamma_{\underline{C}} \underline{C}'(\cdot) - n_\theta) + \frac{n_{\theta Y}}{n_Y}.\end{aligned}$$

If preferences are separable between consumption and earnings ($\mathcal{U}'' = 0$), $U_{Y\theta}/U_Y$ captures the motive for redistribution based on ability, and $U_{C\theta}/U_C$ the needs-based redistribution motive. With the ordering assumption that higher types are more productive, it follows that $U_{Y\theta}/U_Y < 0$.

The slope of $\underline{C}(\cdot)$ in turn determines how consumption needs are aligned with ability. If consumption needs are decreasing in type, then the two redistribution motives reinforce each other towards increasing consumption and lowering effort of the lowest types, but they generate opposite incentives effects: higher ability types have less consumption needs which reduces incentives to work. In this case $U_{C\theta}/U_C \leq 0$. The Single-Crossing Condition imposes that $0 \geq U_{C\theta}/U_C > U_{Y\theta}/U_Y$, i.e. that ability has a stronger impact on work incentives than needs.

If consumption needs are instead increasing in type, $U_{C\theta}/U_C > 0$. In this case, the Single-Crossing Condition holds automatically since needs- and ability-based incentives are both stronger for higher types. In this case, the planner has a motive of demanding higher effort from, and offering higher consumption to high types.

Non-separability in preferences between consumption and earnings ($\mathcal{U}'' < 0$) does not affect incentives, but strengthens both redistribution motives towards lower types, since $U_\theta = \mathcal{U}' \cdot (\gamma_{\underline{C}} \underline{C}'(\cdot) - n_\theta) > 0$, i.e. higher types obtain higher utility from any pair (C, Y) .

Throughout the paper I assume that $U_{Y\theta}/U_Y < 0$, and I distinguish between the case where need-based redistribution is aligned with redistribution based on ability ($0 \geq U_{C\theta}/U_C > U_{Y\theta}/U_Y$ for all $(C, Y; \theta)$) and the case where it is not ($U_{C\theta}/U_C \geq 0 > U_{Y\theta}/U_Y$ for all $(C, Y; \theta)$).⁸ In both cases I assume that $U_{C\theta}$ doesn't change sign, i.e. either $U_{C\theta} \geq 0$ for all $(C, Y; \theta)$, or $U_{C\theta} \leq 0$ for all $(C, Y; \theta)$. The standard model of ability-based redistribution ($U(C, Y; \theta) = U(C, N)$ with $N = Y/a(\theta)$ and $U_{CN} \geq 0$, where $a(\theta)$ denotes labor productivity) belong to the first case.

⁸It is straight-forward to extend the results to the alternative case in which need-based redistribution is the dominant channel for incentives, adapting the corresponding preference assumptions and their interpretation.

2.1 Social planner's problem

The utilitarian planner designs a sequence $\{C(\theta^t), Y(\theta^t)\}$ to maximize the agents' expected utility, subject to incentive compatibility and break-even conditions. I reformulate this problem as a problem of minimizing the cost of the planner's allocation, subject to promise-keeping and incentive compatibility:

$$K(v_0) = \min_{\{C(\theta^t), Y(\theta^t)\}} \mathbb{E}_0 \left\{ \sum_{t=1}^T R^{-t} (C(\theta^t) - Y(\theta^t)) \right\}, \text{ s.t.}$$

$$\mathbb{E}_0 \left\{ \sum_{t=1}^T \beta^{t-1} U(C(\theta^t), Y(\theta^t); \theta_t) \right\} \geq v_0$$

and

$$\mathbb{E}_0 \left\{ \sum_{t=1}^T \beta^{t-1} U(C(\theta^t), Y(\theta^t); \theta_t) \right\} \geq \mathbb{E}_0 \left\{ \sum_{t=1}^T \beta^{t-1} U(C(\vec{\theta}^t(\theta^t)), Y(\vec{\theta}^t(\theta^t)); \theta_t) \right\}$$

for all types $\theta^T \in [\underline{\theta}, \bar{\theta}]^T$ and announcement strategies $\vec{\theta}^T : [\underline{\theta}, \bar{\theta}]^T \rightarrow [\underline{\theta}, \bar{\theta}]^T$ that are measurable w.r.t. θ^t in period t .

To define the problem recursively, let

$$v_{t-1}(\theta^{t-1}) = \mathbb{E} \left\{ \sum_{\tau=t}^T \beta^{\tau-t} U(C(\theta^\tau), Y(\theta^\tau); \theta_\tau) \mid \theta^{t-1} \right\}$$

$$w_t(\theta^t) = U(C(\theta^t), Y(\theta^t); \theta_t) + \beta v_t(\theta^t)$$

I relax incentive compatibility to local IC by which this inequality must only hold for θ' sufficiently close to θ . Following Pavan, Segal and Toikka (2014), Farhi and Werning (2013) or Kapicka (2013), the local IC constraint is

$$\frac{\partial w_t(\theta^t)}{\partial \theta_t} \equiv \dot{w}_t(\theta^t) = U_\theta(C(\theta^t), Y(\theta^t); \theta_t) + \beta \Delta_t(\theta^t), \text{ where}$$

$$\Delta_t(\theta^t) = \int w_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1} \mid \theta^t)}{\partial \theta_t} d\theta^{t+1} = \mathbb{E} \{ w_{t+1}(\theta^{t+1}) \mathcal{J}(\theta_{t+1}, \theta^t) \mid \theta^t \}.$$

The term $\Delta_t(\theta^t)$ constitutes a promise or commitment towards future redistribution so as to limit dynamic marginal information rents, i.e. the part of the marginal information rent $\dot{w}_t(\theta^t)$ that an agent obtains at type θ^t , which is not embedded in the current utility U_θ . With time-independent types, $\mathcal{J}(\theta_t, \theta^{t-1}) = 0$, so the local IC condition reduces to its static counterpart. In

that case, incentive provision occurs only within, but not across periods, so the planner incurs no dynamic incentive commitments and Δ can be dropped from the set of state variables.⁹

Using $v = v_{t-1}(\theta^{t-1})$ and $\Delta = \Delta_{t-1}(\theta^{t-1})$ as state variables, the optimal allocation then minimizes the planner's net present value of transfers $K_t(v, \Delta, \theta^{t-1})$ subject to promise keeping and local incentive compatibility constraints and satisfies the following recursive characterization:

$$K_t(v, \Delta, \theta^{t-1}) = \min \int R^{-1} \{C_t(\theta^t) - Y_t(\theta^t) + K_{t+1}(v_t(\theta^t), \Delta_t(\theta^t), \theta^t)\} f(\theta_t | \theta^{t-1}) d\theta_t, \text{ s.t.}$$

$$v = \int w_t(\theta^t) f(\theta_t | \theta^{t-1}) d\theta_t$$

$$\Delta = \int w_t(\theta^t) \mathcal{J}(\theta_t, \theta^{t-1}) f(\theta_t | \theta^{t-1}) d\theta_t$$

$$w_t(\theta^t) = U(C_t(\theta^t), Y_t(\theta^t); \theta_t) + \beta v_t(\theta^t)$$

$$\dot{w}_t(\theta^t) = U_\theta(C_t(\theta^t), Y_t(\theta^t); \theta_t) + \beta \Delta_t(\theta^t)$$

The initial state variables v_0 and Δ_0 are then chosen so that $\Delta_0 \in \arg \min_{\Delta} K_1(v_0, \Delta, \theta^{-1})$ and $K(v_0) = K_1(v_0, \Delta_0, \theta^{-1}) = 0$, i.e. Δ_0 is set to minimize $K_1(v_0, \Delta, \theta^{-1})$, meaning that at date 0 the planner's break-even constraint is satisfied and the planner is free of prior commitments.

This recursive control problem decomposes into a static optimal control problem using $w_t(\theta^t)$ as the state and $C_t(\theta^t)$, $Y_t(\theta^t)$ and $\Delta_t(\theta^t)$ as control variables and a dynamic programming problem to keep track of utility promises $v_t(\theta^t)$ and incentive commitments $\Delta_t(\theta^t)$. The Hamiltonian is

$$\begin{aligned} \mathcal{H}_t &= R^{-1} \{C_t(\theta^t) - Y_t(\theta^t) + K(v_t(\theta^t), \Delta_t(\theta^t), \theta^t, t+1)\} f(\theta_t | \theta^{t-1}) \\ &+ \{\lambda_t(v - w_t(\theta^t)) + \eta_t(w_t(\theta^t) \mathcal{J}(\theta_t, \theta^{t-1}) - \Delta)\} f(\theta_t | \theta^{t-1}) \\ &+ \psi_t(\theta^t) \{w_t(\theta^t) - U(C_t(\theta^t), Y_t(\theta^t); \theta_t) - \beta v_t(\theta^t)\} \\ &+ \mu_t(\theta^t) \{U_\theta(C_t(\theta^t), Y_t(\theta^t); \theta_t) + \beta \Delta_t(\theta^t)\} \end{aligned}$$

The first-order conditions w.r.t. $C_t(\cdot)$ and $Y_t(\cdot)$ yield

$$\psi_t(\theta^t) = \frac{R^{-1}}{U_C(\theta^t)} f(\theta_t | \theta^{t-1}) + \mu_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} = \frac{R^{-1}}{-U_Y(\theta^t)} f(\theta_t | \theta^{t-1}) + \mu_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}$$

where I use the notation $G(\theta^t) \equiv G(C_t(\theta^t), Y_t(\theta^t); \theta_t)$ for any function G of both the allocation $(C_t(\theta^t), Y_t(\theta^t))$ and the current type θ_t , to ease notation. These first order conditions define a

⁹As a corollary, it follows that unless types are time-independent, the planner's solution is generically time-inconsistent, since the planner must keep track of prior incentive commitments that are summarized by Δ : at date t agents' incentives are influenced by how the planner promises to reallocate resources at future dates. But once those future dates arrive those prior incentive effects are "sunk".

shadow cost of increasing the utility of agents with type θ^t , $\psi_t(\theta^t)$, which consists of two components. The first term $f(\theta_t|\theta^{t-1})R^{-1}/U_C(\theta^t)$ or $f(\theta_t|\theta^{t-1})R^{-1}/(-U_Y(\theta^t))$ represents a direct shadow cost of increasing type θ utility through higher consumption or lower earnings. The second term measures how a direct consumption or earnings increase affects $U_\theta(\theta^t)$ and thereby tightens or relaxes the local incentive compatibility constraint at θ^t . The latter is weighted by the multiplier $\mu_t(\theta^t)$ and added to the former.

The first-order conditions for $v_t(\theta^t)$ and $\Delta_t(\theta^t)$, along with the envelope conditions at $t+1$, yield

$$\begin{aligned}\beta R \frac{\psi_t(\theta^t)}{f(\theta_t|\theta^{t-1})} &= \frac{\partial K_{t+1}}{\partial v_t(\theta^t)} = \lambda_{t+1} \\ \beta R \frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} &= -\frac{\partial K_{t+1}}{\partial \Delta_t(\theta^t)} = \eta_{t+1}\end{aligned}$$

These conditions link the marginal costs of promised utility λ_{t+1} and incentive commitments η_{t+1} back in time to marginal costs of promised utility and shadow price of redistribution at θ^t . The initial incentive commitment Δ_0 is freely chosen, which results in $\eta_1 = -\frac{\partial K_1}{\partial \Delta_0} = 0$.

Finally, the multiplier $\mu_t(\theta^t)$ is given by the solution to the following linear ODE:

$$\begin{aligned}\dot{\mu}_t(\theta^t) &= -\frac{\partial \mathcal{H}_t}{\partial w_t(\theta^t)} = \lambda_t f(\theta_t|\theta^{t-1}) - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) f(\theta_t|\theta^{t-1}) - \psi_t(\theta^t) \\ &= \lambda_t f(\theta_t|\theta^{t-1}) - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) f(\theta_t|\theta^{t-1}) - \frac{R^{-1}}{U_C(\theta^t)} f(\theta_t|\theta^{t-1}) - \mu_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)}\end{aligned}$$

along with the boundary conditions $\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0$, which imply that there are no distortions at the top and the bottom of the type distribution. The FOC for earnings yields an analogous ODE.

Let $\tau_t(\theta^t) = 1 + U_Y(\theta^t)/U_C(\theta^t)$ denote the labor wedge at θ^t , i.e. the distortion between the marginal product and the marginal rate of substitution between consumption and earnings. Combining the first-order conditions for $C_t(\cdot)$ and $Y_t(\cdot)$ yields the following static optimality condition:

$$\frac{R^{-1}}{U_C(\theta^t)} \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} = \frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} \left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right).$$

This static optimality condition succinctly summarizes the static trade-off between efficiency and redistribution at date 1. To substantiate this interpretation, consider a perturbation that leaves a given type θ^t agent's utility unchanged ($\Delta C_1(\theta^t) = \delta/U_C(\theta^t)$ and $\Delta Y_1(\theta^t) = \delta/(-U_Y(\theta^t))$). The planner raises resources

$$\Delta Y_1(\theta^t) - \Delta C_1(\theta^t) = \left(\frac{1}{-U_Y(\theta^t)} - \frac{1}{U_C(\theta^t)} \right) \delta = \frac{1}{U_C(\theta^t)} \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \delta$$

through this perturbation, which have to be weighted by the density $f(\theta_t|\theta^{t-1})$. The LHS in the static optimality condition thus represents the planner's resource gain from this perturbation discounted by R^{-1} , or in other words, the marginal efficiency cost of distorting labor supply at θ^t .

The RHS describes how this resource gain can be re-distributed. The perturbation changes the marginal information rent $U_\theta(\theta^t)$ at θ^t by

$$\Delta U_\theta(\theta^t) = U_{\theta C}(\theta^t) \Delta C_1 + U_{\theta Y}(\theta^t) \Delta Y_1 = \left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right) \delta.$$

From the single-crossing condition, $\Delta U_\theta(\theta^t) > 0$. Therefore, upon reducing the distortion at θ^t , the planner cannot freely re-distribute the extra resources across all types, but must raise the utility of all types $\theta' > \theta_t$ by an extra $\Delta U_\theta(\theta^t)$, relative to all types $\theta' < \theta_t$. Hence, at each θ^t , the planner faces a simple tradeoff between an efficiency motive on the LHS and a redistribution motive on the RHS: More redistribution around θ^t must come at the cost of lower efficiency at θ^t , and vice versa. I can thus re-state the static optimality condition as follows:

$$\frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} = MC_t(\theta^t) \equiv R^{-1} \frac{\frac{1}{-U_Y(\theta^t)} - \frac{1}{U_C(\theta^t)}}{\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}}.$$

In other words, the optimal allocation equates the normalized multiplier $\mu_t(\theta^t)/f(\theta_t|\theta^{t-1})$, or *shadow price of redistribution*, to a static *marginal cost of redistribution at θ* that is given by the ratio of the resource loss of marginally increasing the tax distortion, to the reduction in marginal information rents and hence the increase in redistribution that this entails.

3 Incentive-adjusted Probability Measures

In this section, I introduce the key technical idea that the solution to the multipliers $\mu_t(\theta^t)$ and λ_t can be defined by means of a change in the probability measure that governs the Markov chain of types. This change captures the notion that incentive compatibility imposes restrictions on the planner's ability to redistribute utility from higher to lower types. Specifically, by combining the linear ODE for $\mu_t(\theta^t)$ with the first order conditions for consumption and earnings and the boundary conditions $\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0$, I obtain two separate, but equivalent characterizations of $\mu_t(\theta^t)$ and λ_t , once based on $\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)}$ and marginal utilities of consumption, and once based on $\frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}$ and marginal disutilities of earnings. These characterizations are summarized below in Proposition 1:

Proposition 1 : The multipliers $\mu_t(\theta^t)$ and λ_t can be represented in one of the following two (equivalent) forms:

(i) In terms of marginal benefits of redistribution through consumption:

$$\frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} = \widehat{MB}_t(\theta^t) + \hat{\varrho}_t(\theta^t) \eta_t, \text{ where}$$

$$\widehat{MB}_t(\theta^t) \equiv R^{-1} \frac{1 - \hat{F}(\theta_t|\theta^{t-1})}{\hat{f}(\theta_t|\theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta', \theta^{t-1})} | \theta' \geq \theta_t, \theta^{t-1} \right) - \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^t)} | \theta^{t-1} \right) \right\}$$

$$\hat{f}(\theta_t|\theta^{t-1}) \equiv \frac{f(\theta_t|\theta^{t-1}) \hat{m}(\theta^t)}{\int_{\underline{\theta}}^{\bar{\theta}} f(\theta_t|\theta^{t-1}) \hat{m}(\theta^t) d\theta_t} \text{ with } \hat{m}(\theta^t) = e^{-\int_{\underline{\theta}}^{\bar{\theta}} \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} d\theta_t}$$

$$\hat{\varrho}_t(\theta^t) = \varrho_t(\theta^t) + \frac{1 - \hat{F}(\theta_t|\theta^{t-1})}{\hat{f}(\theta_t|\theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\varrho_t(\theta, \theta^{t-1}) \frac{U_{\theta C}(\theta, \theta^{t-1})}{U_C(\theta, \theta^{t-1})} | \theta \geq \theta_t, \theta^{t-1} \right) - \hat{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} | \theta^{t-1} \right) \right\}$$

(ii) In terms of marginal benefits of redistribution through earnings:

$$\frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} = \widetilde{MB}_t(\theta^t) + \tilde{\varrho}_t(\theta^t) \eta_t, \text{ where}$$

$$\widetilde{MB}_t(\theta^t) \equiv R^{-1} \frac{1 - \tilde{F}(\theta_t|\theta^{t-1})}{\tilde{f}(\theta_t|\theta^{t-1})} \left\{ \tilde{\mathbb{E}} \left(\frac{1}{-U_Y(\theta', \theta^{t-1})} | \theta' \geq \theta_t, \theta^{t-1} \right) - \tilde{\mathbb{E}} \left(\frac{1}{-U_Y(\theta^t)} | \theta^{t-1} \right) \right\}$$

$$\tilde{f}(\theta_t|\theta^{t-1}) \equiv \frac{f(\theta_t|\theta^{t-1}) \tilde{m}(\theta^t)}{\int_{\underline{\theta}}^{\bar{\theta}} f(\theta_t|\theta^{t-1}) \tilde{m}(\theta^t) d\theta_t} \text{ with } \tilde{m}(\theta^t) = e^{-\int_{\underline{\theta}}^{\bar{\theta}} \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} d\theta_t}$$

$$\tilde{\varrho}_t(\theta^t) = \varrho_t(\theta^t) + \frac{1 - \tilde{F}(\theta_t|\theta^{t-1})}{\tilde{f}(\theta_t|\theta^{t-1})} \left\{ \tilde{\mathbb{E}} \left(\varrho_t(\theta, \theta^{t-1}) \frac{U_{\theta Y}(\theta, \theta^{t-1})}{U_Y(\theta, \theta^{t-1})} | \theta \geq \theta_t, \theta^{t-1} \right) - \tilde{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} | \theta^{t-1} \right) \right\}$$

In addition,

$$\begin{aligned} \lambda_t &= R^{-1} \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^t)} | \theta^{t-1} \right) + \eta_t \hat{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} | \theta^{t-1} \right) \\ &= R^{-1} \tilde{\mathbb{E}} \left(\frac{1}{-U_Y(\theta^t)} | \theta^{t-1} \right) + \eta_t \tilde{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} | \theta^{t-1} \right) \end{aligned}$$

In this proposition $\hat{\mathbb{E}}(\cdot)$ denotes the expectations operator associated with cdf $\hat{F}(\cdot)$ and $\tilde{\mathbb{E}}(\cdot)$ the expectations operator associated with cdf $\tilde{F}(\cdot)$. The main message from Proposition 1 is that non-separability of preferences enters the characterization of multipliers by adjusting the probability measure over types from $F(\cdot)$ to $\hat{F}(\cdot)$ when $U_{\theta C} \neq 0$ or from $F(\cdot)$ to $\tilde{F}(\cdot)$ when $U_{\theta Y} \neq 0$.

The direction of these incentive adjustments depends on the signs of $U_{\theta C}$ and $U_{\theta Y}$. Since $U_{\theta Y}/U_Y < \min\{0, U_{\theta C}/U_C\}$, $\tilde{F}(\cdot|\theta^{t-1})$ is first-order stochastically dominated by $F(\cdot|\theta^{t-1})$ and

$\hat{F}(\cdot|\theta^{t-1})$. In addition, $\hat{F}(\cdot|\theta^{t-1})$ first-order stochastically dominates $F(\cdot|\theta^{t-1})$ when $U_{\theta C} > 0$, and $F(\cdot|\theta^{t-1})$ first-order stochastically dominates $\hat{F}(\cdot|\theta^{t-1})$ when $U_{\theta C} < 0$. The characterization subsumes the benchmark case with separable preferences ($U_{\theta C} = 0$) as a special case in which $\hat{F}(\cdot|\theta^{t-1}) = F(\cdot|\theta^{t-1})$, i.e. incentive-adjusted and original probability measures coincide.

The ranking of these distributions in terms of first-order dominance is naturally linked to the interpretation of $U_{\theta C}/U_C$ and $U_{\theta Y}/U_Y$ as needs- and ability-based redistribution motives. Since $U_{\theta Y}/U_Y < 0$, increasing utilities by reducing earnings lowers information rents, and thus allows for more redistribution of utility towards lower types. This explains why $\tilde{F}(\cdot|\theta^{t-1})$ overweighs the lower types relative to $F(\cdot|\theta^{t-1})$.

Likewise, $\hat{F}(\cdot|\theta^{t-1})$ is shifted towards higher (lower) types, whenever higher types have higher (lower) consumption needs. When higher types have lower consumption needs ($U_{\theta C} < 0$), increasing consumption reduces information rents and thus allows for more redistribution towards lower types. In contrast, when consumption needs are increasing in type, extra consumption increases information rents and reduces redistribution towards lower types, in which case $\hat{F}(\cdot|\theta^{t-1})$ first-order stochastically dominates $F(\cdot|\theta^{t-1})$.

Redistributive Perturbations: These changes in probability measures are naturally linked to a perturbation that identifies the *marginal benefit of redistribution* $\widehat{MB}_t(\theta^t)$ as the resource gain of a local utility transfer from $\theta' > \theta_t$ to $\theta' < \theta_t$ that preserves incentive compatibility and expected utility.¹⁰ When $\hat{\varrho}_t(\theta^t)\eta_t = 0$, the optimal allocation reduces to a simple comparison of static costs and benefits of redistribution. When $\hat{\varrho}_t(\theta^t)\eta_t > 0$, the static marginal benefit is augmented by a backward-looking dynamic component.

Consider the following class of perturbations which transfer consumption from all agents with types $\theta' > \theta$ to all agents with types $\theta' < \theta$

$$\Delta C_1(\theta', \theta^{t-1}) = \begin{cases} \frac{1}{U_C(\theta', \theta^{t-1})} \cdot \Delta V(\theta', \theta^{t-1}) < 0 & \text{if } \theta' > \theta_t \\ \frac{1}{U_C(\theta', \theta^{t-1})} \cdot \Delta V(\theta', \theta^{t-1}) > 0 & \text{if } \theta' < \theta_t \end{cases}$$

This perturbation decreases realized utility by $-\Delta V(\theta', \theta^{t-1})$ for all agents with types $\theta' > \theta$ and increases realized utility by $\Delta V(\theta', \theta^{t-1})$ for all agents with types $\theta' < \theta$. Set $\Delta V(\theta', \theta^{t-1})$ so that $\lim_{\theta' \uparrow \theta} \Delta V(\theta', \theta^{t-1}) = \hat{\zeta}_- > 0$ and $\lim_{\theta' \downarrow \theta} \Delta V(\theta', \theta^{t-1}) = \hat{\zeta}_+ < 0$, with $\hat{\zeta}_- - \hat{\zeta}_+ = \delta > 0$, i.e. the change in utility around θ is of size δ .

¹⁰See Brendon (2013) and Farhi and Werning (2012) for related perturbation arguments in static and dynamic Mirrlees models with separable preferences.

The perturbation changes marginal information rents $U_\theta(\theta', \theta^{t-1})$ by

$$\Delta U_\theta(\theta', \theta^{t-1}) = \frac{U_{\theta C}(\theta', \theta^{t-1})}{U_C(\theta', \theta^{t-1})} \cdot \Delta V(\theta', \theta^{t-1}),$$

and it therefore preserves local incentive compatibility if and only if

$$\Delta V'(\theta', \theta^{t-1}) = \Delta U_\theta(\theta', \theta^{t-1}) = \frac{U_{\theta C}(\theta', \theta^{t-1})}{U_C(\theta', \theta^{t-1})} \Delta V(\theta', \theta^{t-1}).$$

The unique perturbation that preserves local IC for all $\theta' \neq \theta$ is then given by

$$\Delta V(\theta', \theta^{t-1}) = \begin{cases} e^{\int_\theta^{\theta'} \frac{U_{\theta C}(\theta'', \theta^{t-1})}{U_C(\theta'', \theta^{t-1})} d\theta''} \cdot \hat{\zeta}_+ = \frac{\hat{m}(\theta', \theta^{t-1})}{\hat{m}(\theta_t, \theta^{t-1})} \hat{\zeta}_+ < 0 \text{ if } \theta' > \theta_t \\ e^{\int_\theta^{\theta'} \frac{U_{\theta C}(\theta'', \theta^{t-1})}{U_C(\theta'', \theta^{t-1})} d\theta''} \cdot \hat{\zeta}_- = \frac{\hat{m}(\theta', \theta^{t-1})}{\hat{m}(\theta_t, \theta^{t-1})} \hat{\zeta}_- > 0 \text{ if } \theta' < \theta_t \end{cases}$$

For θ' just above θ , the perturbation changes $U_\theta(\theta', \theta^{t-1})$ by $(U_{\theta C}(\theta', \theta^{t-1}) / U_C(\theta', \theta^{t-1})) \cdot \hat{\zeta}_+$, and this adjustment to the slope must be offered to all $\theta'' > \theta'$ to restore local incentive compatibility at θ' . But these modifications for any $\theta'' > \theta'$ then generate further adjustments for all $\theta''' > \theta''$, and so on. By re-weighting the utility perturbations according to $\hat{m}(\cdot)$, the perturbation preserves local incentive compatibility for all types: the ODE is solved by “integrating up” the cumulative utility changes for higher types that are required as a result of preserving local IC at all lower types.

I complete the perturbation by requiring that the net change to expected utility is 0:

$$0 = \int_{\underline{\theta}}^{\bar{\theta}} \Delta V(\theta', \theta^{t-1}) f(\theta' | \theta^{t-1}) d\theta' = \int_{\underline{\theta}}^{\bar{\theta}} \frac{\hat{m}(\theta')}{\hat{m}(\theta)} \hat{\zeta}_+ f(\theta' | \theta^{t-1}) d\theta' + \int_{\underline{\theta}}^{\theta} \frac{\hat{m}(\theta')}{\hat{m}(\theta)} \hat{\zeta}_- f(\theta' | \theta^{t-1}) d\theta',$$

which yields $\hat{\zeta}_- = \delta \left(1 - \hat{F}(\theta_t | \theta^{t-1})\right)$ and $\hat{\zeta}_+ = -\delta \hat{F}(\theta_t | \theta^{t-1})$. The discounted resource gain of this perturbation is then given by

$$\begin{aligned} & -R^{-1} \int_{\underline{\theta}}^{\bar{\theta}} \Delta C_1(\theta', \theta^{t-1}) f(\theta' | \theta^{t-1}) d\theta' \\ &= \delta R^{-1} \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{m}(\theta_t, \theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta', \theta^{t-1})} \middle| \theta' \geq \theta_t, \theta^{t-1} \right) - \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^t)} \middle| \theta^{t-1} \right) \right\} \\ &= \delta f(\theta_t | \theta^{t-1}) \widehat{MB}_t(\theta^t), \end{aligned}$$

which justifies the interpretation of $\widehat{MB}_t(\theta^t)$ as the marginal benefit of redistribution around θ^t via consumption, i.e. the marginal benefit of transferring consumption around θ^t while maintaining incentive compatibility and keeping expected utility unchanged. This transfer saves resources because higher types have lower marginal utilities.

In the benchmark with separable preferences ($U_{\theta C}(\cdot) = 0$), this perturbation transfers utility uniformly from all types $\theta' > \theta_t$ to all type $\theta' < \theta_t$. With separable preferences, this uniform

transfer of utilities does not affect local incentive compatibility constraints other than at θ_t , which explains why the original and incentive-adjusted probability measures coincide.

An analogous perturbation argument explains why a local IC and expected-utility preserving perturbation that redistributes leisure or earnings around θ must reweight states according to $\tilde{m}(\cdot)$, resulting in the above expression $\widetilde{MB}_t(\theta^t)$ as the marginal benefit of redistribution via earnings.

Tying incentive adjustments to observables: The next proposition decomposes the changes in probability measures $\hat{m}(\cdot)$ and $\tilde{m}(\cdot)$ into a risk component $U_C(\theta^t)$ or $-U_Y(\theta^t)$ that captures the redistribution motive, and an incentive component $\hat{M}(\theta^t)$ or $\widetilde{M}(\theta^t)$ that represents the adjustment to the marginal redistribution of resources: preserving incentive compatibility requires that on the margin consumption must be redistributed in proportion to $\hat{M}(\cdot)$ and earnings in proportion to $\widetilde{M}(\cdot)$. Furthermore, the latter only depends on the elasticities \mathcal{E}_C and \mathcal{E}_Y and the distribution of consumption and earnings.

Proposition 2 *The incentive adjustments $\hat{m}(\cdot)$ and $\tilde{m}(\cdot)$ admit the following representations:*

(i) *Consumption-based incentive adjustment: $\hat{m}(\theta^t) = U_C(\theta^t) \hat{M}(\theta^t)$, where*

$$\hat{M}(\theta^t) = e^{\int_{\underline{\theta}}^{\theta^t} \mathcal{E}_C(\theta', \theta^{t-1}) d \ln C(\theta', \theta^{t-1})}$$

(ii) *Earnings-based incentive adjustment: $\tilde{m}(\theta^t) = -U_Y(\theta^t) \widetilde{M}(\theta^t)$, where*

$$\widetilde{M}(\theta^t) = e^{-\int_{\underline{\theta}}^{\theta^t} \mathcal{E}_Y(\theta', \theta^{t-1}) d \ln C(\theta', \theta^{t-1})}$$

The incentive component $\hat{M}(\cdot)$ is increasing in θ_t so consumption must be redistributed regressively on the margin except when $\mathcal{E}_C = 0$. Moreover, $\hat{M}(\cdot)$ only depends on the distribution of consumption and the elasticity $\mathcal{E}_C(\cdot)$, and when the latter is constant, $\hat{M}(\theta^t) = C(\theta^t)^{\mathcal{E}_C}$. On the other hand, $\widetilde{M}(\theta^t)$ is decreasing in θ_t , unless $\mathcal{E}_Y = 0$, so the marginal incentive-compatible redistribution of utility via earnings is progressive. It only depends on the distribution of earnings and the earnings elasticity $\mathcal{E}_Y(\cdot)$, and when the latter is constant, $\widetilde{M}(\theta^t) = Y(\theta^t)^{-\mathcal{E}_Y}$.

The elasticities $\mathcal{E}_Y(\theta^t)$ and $\mathcal{E}_C(\theta^t)$ were defined in section 2. They determine how much the marginal rate of substitution responds to increases in consumption or earnings. These two parameters have natural counter-parts in terms of income and substitution effects of labor supply, which are fully discussed in a static application of this framework in Hellwig and Werquin (2021). Intuitively speaking, $1/\mathcal{E}_Y(\theta^t)$ represents an elasticity of labor supply to wages, holding consumption constant, and $\mathcal{E}_C(\theta^t)$ governs the relative strength of income and substitution effects.

When preferences are weakly separable ($U(C, Y; \theta) = \mathcal{U}(\gamma(C, \underline{C}(\theta)) - n(Y, \theta))$), the risk and incentive components both include a factor $\gamma_C(\cdot)$ or $n_Y(\cdot)$ which then cancels out from the product:

$\hat{m}(\theta^t) = \mathcal{U}'(\theta^t) \chi(\theta^t)$ where $\chi(\theta^t) = e^{\int_{\underline{\theta}}^{\theta^t} \frac{\gamma_{C\underline{C}}(\theta', \theta^{t-1})}{\gamma_C(\theta', \theta^{t-1})} \underline{C}'(\theta', \theta^{t-1}) d\theta'}$ and $\tilde{m}(\theta^t) = \mathcal{U}'(\theta^t) p(\theta^t)$ where $p(\theta^t) = e^{\int_{\underline{\theta}}^{\theta^t} \frac{n_{Y\theta}(\theta', \theta^{t-1})}{n_Y(\theta', \theta^{t-1})} d\theta'}$. Hence $\hat{m}(\theta^t)$ and $\tilde{m}(\theta^t)$ decompose into a risk component from the curvature in \mathcal{U} and an incentive component from the non-separabilities between the consumption utility $\gamma(C, \underline{C}(\cdot))$ or effort $n(Y, \cdot)$ and type θ . When preferences are additively separable, then $\mathcal{U}'(\cdot) = \chi(\cdot) = 1$, and the risk and incentive components in the adjusted probability measure just offset each other, $\hat{M}(\theta^t) = 1/U_C(\theta^t)$. Also, with weakly separable preferences, the elasticity $\mathcal{E}_C(\cdot)$ only depends on the curvature of $\gamma_C(\cdot)$, while the elasticity $\mathcal{E}_Y(\cdot)$ only depends on the curvature in $n(Y, \theta)$, in which case $1/\mathcal{E}_Y(\cdot)$ represents the Frisch elasticity of labor supply.

4 Benchmark: separable preferences

Before developing the general characterization of optimal labor and savings wedges, it will be useful to pause and review the well-known benchmark with separable preferences that has been extensively analyzed in Golosov, Kocherlakota and Tsyvinski (2003), Farhi and Werning (2013), Golosov, Troshkin and Tsyvinski (2016) and summarized by Stantcheva (2020). This can be done simply by setting $U_{\theta C} = U_{CY} = 0$ in the characterization of multipliers given by Proposition 1.

Inverse Euler Equation: With separable preferences, the optimal savings wedge is given by the Inverse Euler Equation

$$U_C(\theta^t) = \beta R \left\{ \mathbb{E} \left(\frac{1}{U_C(\theta^{t+1})} \middle| \theta^t \right) \right\}^{-1} < \beta R \mathbb{E}(U_C(\theta^{t+1}) | \theta^t).$$

The incentive-compatibility and expected utility preserving perturbation in the previous section generalizes a well-known perturbation argument which explains why, with separable preferences, the returns to savings must be distributed in proportion to $1/U_C(\theta^{t+1})$ and therefore the discounted shadow cost of utility at date $t + 1$ equals $\lambda_{t+1} = R^{-1} \cdot \mathbb{E}(1/U_C(\theta^{t+1}))$. This expression arises because the planner must redistribute resources at date $t + 1$ in such a manner that utility changes are uniform across all types, so as to preserve incentive compatibility. As a result the resources from savings are distributed regressively: higher θ_{t+1} types enjoy higher returns from savings than lower θ_{t+1} types. This observation then leads to the Inverse Euler Equation (IEE) for optimal savings distortions, which augments the standard inter-temporal substitution channel by the wealth effect of savings on future incentives and redistribution. This effect reduces the social returns to savings relative to the private ones and is captured by the savings wedge implied by the IEE.

Smoothing and back-loading of labor taxes: The optimal labor wedge satisfies

$$\frac{R^{-1}}{U_C(\theta^t)} \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} = \frac{\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} = MC_t(\theta^t) = MB_t(\theta^t) + \beta R \varrho_t(\theta^t) \cdot MC_{t-1}(\theta^{t-1})$$

$$\begin{aligned} \text{where } MB_t(\theta^t) &\equiv R^{-1} \frac{1 - F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})} \left\{ \mathbb{E} \left(\frac{1}{U_C(\theta', \theta^{t-1})} \mid \theta' \geq \theta_t, \theta^{t-1} \right) - \mathbb{E} \left(\frac{1}{U_C(\theta^t)} \mid \theta^{t-1} \right) \right\} \\ \varrho_t(\theta^t) &= \frac{1 - F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})} \mathbb{E} \left(\mathcal{J}(\theta', \theta^{t-1}) \mid \theta' \geq \theta_t, \theta^{t-1} \right) = \frac{\partial(1 - F(\theta_t|\theta^{t-1}))}{\partial \theta_{t-1}} > 0 \end{aligned}$$

i.e. the planner's solution equates the static marginal cost of redistribution for each θ^t , to the static marginal benefit of redistribution around θ^t , $MB_t(\theta^t)$, and a backwards-looking dynamic marginal benefit $\beta R \varrho_t(\theta^t) \cdot MC_{t-1}(\theta^{t-1})$, which captures the marginal impact of prior incentive commitments on the shadow price of redistribution in the current period. These prior incentive commitments constitute a promise to limit future information rents. They appear in the recursive optimality condition through the term $\beta R \varrho_t(\theta^t) \cdot MC_{t-1}(\theta^{t-1})$, where $\varrho_t(\theta^t) > 0$ denotes the rate of decay in information rents. One can rewrite the incentive commitment constraint as

$$\Delta = \int (U_\theta(\theta^t) + \beta \Delta_t(\theta^t)) \varrho_t(\theta^t) f(\theta_t|\theta^{t-1}) d\theta_t.$$

Therefore, $\varrho_t(\theta^t)$ describes the rate at which marginal information rents at θ^t feed back into incentive commitments at θ^{t-1} . $\varrho_t(\theta^t)$ naturally maps into the persistence of the type process at θ^{t-1} : If $\varrho_t(\theta^t) = 0$ for all $\theta_t|\theta^{t-1}$, $\Delta = 0$ and the recursive optimality condition reduces to its static counterpart, since there is no scope for dynamic incentive commitments. When types are perfectly persistent, $\varrho_t(\theta^t) = 1$ and information rents at θ^t are fully passed through to incentive commitments at θ^{t-1} .

The optimal labor distortion trades off the marginal cost of efficiency distortions at θ^t against the marginal redistributive gain. Consider a perturbation $(\Delta C_t(\theta^t), \Delta Y_t(\theta^t))$ that marginally increases distortions at θ^t , but leaves a type θ^t agent's utility unchanged. The marginal efficiency loss associated with this perturbation is $MC_t(\theta^t) \cdot \delta$. The marginal benefit of redistribution at θ^t is $MB_t(\theta^t) \cdot \delta$. In addition, the same reduction in marginal information rents at θ^t also feeds back into marginal information rent at θ^{t-1} at a rate $\varrho_t(\theta^t)$. This reduction in marginal information rents can be transformed into a marginal efficiency gain at a rate $\beta MC_{t-1}(\theta^{t-1})$, which must in turn be adjusted to period t by the return R . Hence the dynamic benefit from additional redistribution at t is given by $\delta \cdot \varrho_t(\theta^t) \cdot \beta R \cdot MC_{t-1}(\theta^{t-1})$. Combining these terms yields the recursive characterization of $MC_t(\theta^t)$ given above.

The recursive optimality condition thus smoothes tax distortions optimally between θ^{t-1} and θ^t , where backloading distortions to θ^t generates additional marginal benefits from redistribution around θ^t . Solving the optimality condition backwards and expressing $MC_t(\theta^t)$ in terms of the labor tax $\tau_t(\theta^t)$ yields the following expression:

$$\frac{\tau_t(\theta^t)}{U_{\theta Y}(\theta^t)} = \sum_{\tau=0}^{t-1} (\beta R)^\tau \prod_{s=0}^{\tau-1} \varrho_{t-s}(\theta^{t-s}) MB_{t-\tau}(\theta^{t-\tau})$$

i.e. at the optimum the current tax distortion is equal to a weighted sum of current and past marginal benefits of redistribution, with weights that take into account the impact of efficiency distortions at θ^t on prior information rents. Hence, the optimal cost-benefit trade-off in the dynamic model sets the cost of current tax distortions against the benefits of higher current and past redistribution: more redistribution in period t reduces prior information rents and thus allows for more redistribution in earlier periods without adversely affecting incentives. The incentive commitment thus constitutes a commitment to future taxes and redistribution to limit current information rents.

Alternative representations: The recursive optimality condition can be related in a straightforward manner to the characterizations provided by Golosov, Troshkin and Tsyvinski (2016) and Farhi and Werning (2013). Multiplying both sides of the optimality condition by $\frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} U_C(\theta^t)$ gives the characterization provided by Golosov, Troshkin and Tsyvinski (2016):

$$\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} = \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} R U_C(\theta^t) MB_t(\theta^t) + \beta R \varrho_t(\theta^t) \frac{U_{Y\theta}(\theta^t)/U_Y(\theta^t)}{U_{Y\theta}(\theta^{t-1})/U_Y(\theta^{t-1})} \frac{U_C(\theta^t)}{U_C(\theta^{t-1})} \frac{\tau_{t-1}(\theta^{t-1})}{1 - \tau_{t-1}(\theta^{t-1})}$$

Based on this expression, they show that the impact of past labor wedges on the current one increases with $U_C(\theta^t)$, and therefore decreases with the current type realization.

One obtains the characterization provided by Farhi and Werning (2013) by multiplying both sides of the recursive optimality condition by $\frac{U_{Y\theta}(\theta^t)}{-U_Y(\theta^t)}$ and then taking expectations:¹¹

$$\mathbb{E} \left(\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \frac{1/U_C(\theta^t)}{\mathbb{E}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) = Cov \left(\int_{\underline{\theta}}^{\theta^t} \frac{U_{Y\theta}(\theta', \theta^{t-1})}{-U_Y(\theta', \theta^{t-1})} d\theta', \frac{1/U_C(\theta^t)}{\mathbb{E}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) \\ + \mathcal{R}(\theta^{t-1}) \frac{\tau_{t-1}(\theta^{t-1})}{1 - \tau_{t-1}(\theta^{t-1})}$$

where $\mathcal{R}(\theta^{t-1}) = \frac{\mathbb{E}(\varrho_t(\theta^t) U_{Y\theta}(\theta^t)/U_Y(\theta^t) | \theta^{t-1})}{U_{Y\theta}(\theta^{t-1})/U_Y(\theta^{t-1})}$.

Hence the recursive optimality condition takes expectations of future tax distortions and future marginal benefits of redistribution under an incentive-adjusted probability measure that re-weights

¹¹Their representation is more general in that marginal costs and benefits can be multiplied with arbitrary weighting functions before taking expectations. Hence their result also implies the present one by applying a weighting function $\pi(\cdot)$ that places weight only on a singleton of types. The two characterizations are thus equivalent.

returns by $1/U_C(\theta^t)$, and labor taxes mean-revert with a rate $\mathcal{R}(\theta^{t-1})$. The term $\frac{U_{Y\theta}(\theta^{t-1})}{U_Y(\theta^{t-1})}$ measures the impact of current labor taxes on current information rents, and $\mathbb{E}(\varrho_t(\theta^t) \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} | \theta^{t-1})$ measures the impact of expected future labor taxes on current information rents. Therefore $\mathcal{R}(\theta^{t-1})$ trades off between the two in a such a way that a marginal transfer of tax distortions over time does not change current information rents and redistribution.

With isoelastic preferences ($U(C, Y; \theta) = \gamma(C) - \kappa(Y/A(\theta))^{1+\mathcal{E}_Y}$) and $A(\theta) = e^{\phi\theta}$ with $\phi > 0$, the expected marginal benefit of future redistribution is $(1 + \mathcal{E}_Y) Cov\left(\log(A(\theta_t)), \frac{\hat{M}(\theta^t)}{\mathbb{E}(\hat{M}(\theta^t) | \theta^{t-1})} | \theta^{t-1}\right)$, while labor taxes mean-revert with a rate $\mathcal{R}(\theta^{t-1}) = \mathbb{E}(\varrho_t(\theta^t) | \theta^{t-1})$ that only depends on the stochastic process of types.

Trading off efficiency and redistribution in this manner relies on the fact that an intervention at or around θ only affects incentives locally when preferences are separable. This assumption completely separates the incentive-compatible utility transfers across time from efficient redistribution within each period. Moreover, the characterization of optimal allocations remains incomplete since it provides a single optimality condition to characterize two allocation variables ($C_t(\theta^t), Y_t(\theta^t)$) for each type. In the remainder of this paper I show how these two results generalize to non-separable preferences, and I complete the characterization of optimal allocations by also considering the benefits of redistribution via earnings or leisure.

5 Optimal Savings Wedges

I now discuss the implications of Proposition 1 for optimal savings wedges. Two changes appear in the expression for λ_t , relative to the separable benchmark: the use of incentive-adjusted probabilities, and the additional term $\eta_t \hat{\mathbb{E}}\left(\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} | \theta^{t-1}\right)$ or $\eta_t \tilde{\mathbb{E}}\left(\varrho_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} | \theta^{t-1}\right)$.

The use of incentive-adjusted probabilities follows from the fact that marginal returns to savings must be redistributed in proportion to $\hat{M}(\cdot)$ for consumption and $\tilde{M}(\cdot)$ for earnings, in order to preserve local incentive compatibility.

The additional term in the expression for λ_t identifies a novel feedback effect from savings to current marginal information rents. A marginal increase in savings at θ^{t-1} , which is redistributed at $\theta^t | \theta^{t-1}$ in proportion to $1/U_C(\theta^t) \cdot \hat{m}(\theta^t) \cdot \delta = \hat{M}(\theta^t)$ to preserve incentive compatibility, changes marginal information rents by an agent of type θ^t by $U_{\theta C}(\theta^t) / U_C(\theta^t) \cdot \hat{m}(\theta^t) \cdot \delta$, and as discussed above, this change in marginal information rents is passed through to θ^{t-1} at a rate $\varrho_t(\theta^t)$. Hence the term $\hat{\mathbb{E}}\left(\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} | \theta^{t-1}\right)$ measures the feedback from future returns to savings into current information rents, and these information rents are “priced” by η_t , which represents the shadow cost

of incentive commitments at θ^{t-1} . Along the same lines, $\tilde{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} | \theta^{t-1} \right)$ represents the feedback from a labor deferral to current information rents.

Combining the dynamic optimality conditions with the characterization for λ_{t+1} and rearranging terms then yields the following expression for the planner's consumption-savings tradeoff:

$$\begin{aligned} \frac{1}{U_C(\theta^t)} + \frac{R\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} &= \frac{R\psi_t(\theta^t)}{f(\theta_t|\theta^{t-1})} \\ &= (\beta R)^{-1} \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^{t+1})} | \theta^t \right) + \frac{R\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} \hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} | \theta^t \right) \end{aligned}$$

This expression generalizes the IEE to non-separable preferences. The terms $1/U_C(\theta^t)$ on the LHS and $(\beta R)^{-1} \hat{\mathbb{E}}(1/U_C(\theta^{t+1}) | \theta^t)$ on the RHS describe the inter-temporal tradeoff between saving resources at θ^t and redistributing the returns to savings in an incentive-compatible manner at date $t+1$. As before savings reduce future redistribution. The incentive-adjusted expectation amplifies the impact of savings on future incentives when $U_{\theta C} > 0$ and dampens it when $U_{\theta C} < 0$.

The two new terms that appear in this expression are related to how savings affect information rents and redistribution at θ^t . Savings affect current redistribution through (i) the direct effect of savings on redistribution at θ^t , and (ii) the feedback from future information rents at $\theta^{t+1} | \theta^t$ to redistribution at θ^t . These two effects determine together how savings affect current redistribution, which in turn increases or reduces the resources generated by savings at θ^t . They are both weighted by the normalized shadow price of redistribution at θ^t , $R\mu_t(\theta^t) / f(\theta_t | \theta^{t-1})$.

The direct effect captures the fact that increasing savings and lowering consumption of type θ^t by $\delta/U_C(\theta^t)$ changes $U_\theta(\theta^t)$ and therefore redistribution around θ^t by $U_{\theta C}(\theta^t) / U_C(\theta^t)$. The resource gain (if $U_{\theta C}(\theta^t) > 0$) or cost (if $U_{\theta C}(\theta^t) < 0$) associated with this change in redistribution appears in the $U_{\theta C}(\theta^t) / U_C(\theta^t)$ term on the LHS.

However, this direct effect must be offset against the feedback of savings from future to current information rents, which is given by $\hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} | \theta^t \right)$ on the RHS. This term inherits the sign of $U_{\theta C}(\cdot)$, i.e. the effect nets out when preferences are separable, it reduces current redistribution when consumption needs are increasing in type ($U_{\theta C}(\cdot) > 0$) and increases current redistribution when consumption needs are decreasing in type ($U_{\theta C}(\cdot) < 0$). This force works against the direct effect, and its strength depends on the decay of future information rents, $\varrho_{t+1}(\theta^{t+1})$.

Theorem 1 shows how these effects combine to determine the overall savings wedge.

Theorem 1 : *An interior solution to the multi-period dynamic optimal taxation problem satisfies*

the following **Generalized Inverse Euler Equation** in consumption:

$$U_C(\theta^t) = \beta R(1 + \hat{s}_C(\theta^t)) \mathbb{E} \left(U_C(\theta^{t+1}) \frac{\hat{M}(\theta^{t+1})}{\mathbb{E}(\hat{M}(\theta^{t+1})|\theta^t)} \middle| \theta^t \right)$$

where

$$\hat{s}_C(\theta^t) = \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \frac{U_{\theta C}(\theta^t)/U_C(\theta^t) - \hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} \middle| \theta^t \right)}{U_{\theta C}(\theta^t)/U_C(\theta^t) - U_{\theta Y}(\theta^t)/U_Y(\theta^t)}.$$

The optimal type-contingent savings wedge $\hat{S}_C(\theta^{t+1}) \equiv (1 + \hat{s}_C(\theta^t)) \frac{\hat{M}(\theta^{t+1})}{\mathbb{E}(\hat{M}(\theta^{t+1})|\theta^t)} - 1$ has expectation $\mathbb{E}(\hat{S}_C(\theta^{t+1})|\theta^t) = \hat{s}_C(\theta^t)$. Finally, $\hat{s}_C(\theta^t) \gtrless 0$ if and only if

$$\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} \gtrless \hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} \middle| \theta^t \right)$$

Theorem 1 highlights the existence of a new intertemporal tradeoff between redistribution at date t and redistribution at date $t + 1$. The dynamic Mirrlees problem augments the standard consumption-savings tradeoff in resources by an additional inter-temporal tradeoff between present and future redistribution that is internalized by the planner, but not by private savings decisions. This tradeoff in turn can lead to optimal savings taxes or subsidies, depending on whether the planner is more concerned about present or future redistribution, as well as by the direction in which savings taxes or subsidies facilitate redistribution.

The optimal type-contingent savings subsidy $\hat{S}_C(\theta^{t+1})$ decomposes into a factor $1 + \hat{s}_C(\theta^t)$ that internalizes the combined effect of savings on current redistribution, and a factor $\frac{\hat{M}(\theta^{t+1})}{\mathbb{E}(\hat{M}(\theta^{t+1})|\theta^t)}$ that redistributes future marginal returns regressively to preserve incentive compatibility at θ^{t+1} . This second factor is derived from the IC-preserving change in probability measure and averages to 1. It therefore doesn't affect the average savings subsidy, generalizing the result of zero expected wealth taxes from Kocherlakota (2005) to arbitrary non-separable preferences. The first factor internalizes the effect of savings on current redistribution. This factor determines whether savings should, on average, be taxed or subsidized.

Effect of savings on current redistribution: If types are not too persistent (the direct effect dominates), savings should be subsidized if consumption needs are increasing in type ($U_{\theta C} > 0$) and taxed if consumption needs are decreasing in type ($U_{\theta C} < 0$). These prescriptions are reversed if the indirect effect dominates. The magnitude of the wedge is decreasing in the degree of persistence of information rents, concretely the condition for $\hat{s}_C(\theta^t) \gtrless 0$ say that savings should be subsidized (taxed) if the impact of savings on current information rents, $U_{\theta C}(\theta^t)/U_C(\theta^t)$, is larger than the

expected feedback from future information rents, $\hat{\mathbb{E}} \left(\varrho_{t+1} (\theta^{t+1}) \frac{U_{\theta C} (\theta^{t+1})}{U_C (\theta^{t+1})} | \theta^t \right)$. This condition can be equivalently re-stated in terms of the marginal impact of the returns on information rents as

$$U_{\theta C} (\theta^t) \gtrless \beta R \mathbb{E} \left(\varrho_{t+1} (\theta^{t+1}) U_{\theta C} (\theta^{t+1}) \frac{\hat{M} (\theta^{t+1})}{\mathbb{E} (\hat{M} (\theta^{t+1}) | \theta^t)} | \theta^t \right).$$

The magnitude of the savings wedge thus relates to the current labor tax $\tau_t (\theta^t) / (1 - \tau_t (\theta^t))$, which is used to “price” redistribution at θ^t , and the expected rate of decay $\varrho_{t+1} (\theta^{t+1})$ which governs the relative strength of the indirect effect. If types are highly persistent, one reverts to a “benchmark” in which it is optimal to leave savings decisions on average undistorted. In this case the direct and indirect effect on current redistribution just offset so savings have no impact on current information rents and redistribution. But this knife-edge result arises for reasons very different from Kocherlakota (2005), since it is based on two competing forces for savings taxes or subsidies just canceling each other.

Effect of savings on future incentives: The change of probability measure $\frac{\hat{M} (\theta^{t+1})}{\mathbb{E} (\hat{M} (\theta^{t+1}) | \theta^t)}$ summarizes how returns to savings must be distributed to preserve incentive compatibility. Since $\hat{M} (\cdot)$ is increasing in θ_{t+1} unless $\mathcal{E}_C = 0$, the returns on savings are regressive, i.e. higher types must earn higher returns to preserve incentive compatibility. This in turn implies that $\mathbb{E} \left(U_C (\theta^{t+1}) \frac{\hat{M} (\theta^{t+1})}{\mathbb{E} (\hat{M} (\theta^{t+1}) | \theta^t)} | \theta^t \right) < \mathbb{E} (U_C (\theta^{t+1}) | \theta^t)$, i.e. the state-contingent savings subsidy reduces incentives to save by increasing returns for high realizations of θ^{t+1} .

When preferences are separable, $\hat{M} (\theta^{t+1}) = 1/U_C (\theta^{t+1})$, and the representation collapses to the standard IEE. In general, the adjustment to preserve incentive compatibility can be either more or less regressive, depending on whether $\hat{M} (\theta^{t+1})$ is steeper or less steep than $1/U_C (\theta^{t+1})$, or equivalently whether $U_{\theta C} > 0$ or $U_{\theta C} < 0$.

Finally, if $\mathcal{E}_C = 0$, the generalized IEE coincides with the normal Euler Equation and it is optimal not to distort savings. This case corresponds to a generalized form of GHH preferences.

With weakly separable preferences ($U (C, Y; \theta) = \mathcal{U} (\gamma (C, \underline{C} (\theta)) - n (Y, \theta))$), the inverse Euler equation can be rewritten as

$$U_C (\theta^t) = \beta R (1 + \hat{s}_C (\theta^t)) \mathbb{E} (\mathcal{U}' (\theta^{t+1}) \cdot \chi (\theta^{t+1}) | \theta^t) \left\{ \mathbb{E} \left(\frac{\chi (\theta^{t+1})}{\gamma_C (\theta^{t+1})} | \theta^t \right) \right\}^{-1}$$

This representation illustrates that the marginal returns to savings, after adjusting for incentive compatibility, depend on three elements: (i) the arithmetic expectation of the “outer” marginal utility $\mathcal{U}' (\cdot)$ which does not affect incentives, (ii) the harmonic expectation of the “inner” marginal

utility $\gamma_C(\cdot)$ which does, and (iii) an incentive adjustment to the expectations given by $\chi(\cdot)$ which factors in the non-separability in γ . It is thus important to distinguish between non-separability coming from curvature in the aggregator function $\mathcal{U}(\cdot)$ which is immaterial for incentives and thus treated according to the standard Euler Equation, and curvature and non-separability in the inner function γ , which matters for incentives and is thus treated according to the Inverse Euler Equation, with an incentive adjustment for the non-separability in γ . The result is a “mixed” Euler Equation that combines elements of both the standard and the Inverse Euler Equation in one.

Optimal Social Mobility: Theorem 1 has direct implications for the question of optimal social mobility, originally posed in the work of Thomas and Worrall (1990) and Atkeson and Lucas (1992) who presented dynamic private information economies in which marginal utilities tended to infinity and consumption to 0 with probability 1 for almost all agents. To this end, I set $\beta R = 1$, so that long-run dynamics are not determined by discounting and returns. Then the optimal savings wedge implies the following about the process of Marginal utilities:

Corollary 1 : *Marginal utilities satisfy the following dynamics:*

(i) *If $\hat{s}_C(\theta^t) \leq 0$ for all θ^t , then $\frac{1}{U_C(\theta^t)} \geq \hat{\mathbb{E}}\left(\frac{1}{U_C(\theta^{t+1})}|\theta^t\right)$ and $U_C(\theta^t) \leq \mathbb{E}(U_C(\theta^{t+1})|\theta^t)$ for all θ^t .*

(ii) *If $\hat{s}_C(\theta^t) > 0$ for all θ^t , then $\frac{1}{U_C(\theta^t)} < \hat{\mathbb{E}}\left(\frac{1}{U_C(\theta^{t+1})}|\theta^t\right)$ for all θ^t .*

The first part of the corollary states that if it is optimal to always restrict savings, the inverse marginal utility process (and by extension, consumption) follow a super-martingale under the incentive-adjusted probability measure, and marginal utility follows a sub-martingale under the original measure. This in turn generalizes the force towards immiseration that is present when preferences are separable and inverse Marginal utilities follow a Martingale.

When it is instead optimal to subsidize savings, then the inequalities are reversed and inverse marginal utilities and consumption drift up under the incentive-adjusted probability measure. While this cannot be directly mapped into corresponding dynamics under the original measure, using the representation of $\hat{M}(\cdot)$ from proposition 2 when \mathcal{E}_C is constant and preferences weakly separable yields

$$\mathbb{E}\left(\left(\frac{C(\theta^{t+1})}{C(\theta^t)}\right)^{\mathcal{E}_C}|\theta^t\right) \geq \mathbb{E}\left(\frac{\mathcal{U}'(\theta^{t+1}) \cdot \chi(\theta^{t+1})}{\mathcal{U}'(\theta^t) \cdot \chi(\theta^t)}|\theta^t\right)$$

and when $\chi(\cdot)$ is increasing in θ^{t+1} and types are sufficiently mean-reverting, then it is optimal to have a savings subsidy and a consumption process that inherits the mean-reversion from $\mathcal{U}'(\theta^{t+1}) \cdot \chi(\theta^{t+1})$, thus breaking the forces towards immiseration. As discussed for example by Farhi and

Werning (2007), the existing immiseration results rely on the fact that at the optimal allocation, temporary shocks to labor productivity have permanent effects on future marginal utilities as a consequence of smoothing incentives over time. Here the temporary co-movement of consumption needs or marginal utilities with types makes it optimal to build some “foregiveness” into optimal consumption allocations: more productive types also have higher consumption needs, which makes it optimal to front-load consumption increases for incentive reasons. This restores mean reversion in the long-run without differential discounting or altruism towards future generations as in Farhi and Werning (2007) or Phelan (2006).

Inter-temporal Earnings Wedge: Along the same lines, one can re-cast optimal inter-temporal allocations of effort or labor by a generalized IEE which yields the following expression:

$$-U_Y(\theta^t) = \beta R (1 + \tilde{s}_Y(\theta^t)) \mathbb{E} \left(-U_Y(\theta^{t+1}) \frac{\tilde{M}(\theta^{t+1})}{\mathbb{E}(\tilde{M}(\theta^{t+1})|\theta^t)} \middle| \theta^t \right)$$

where

$$\tilde{s}_Y(\theta^t) = \tau_t(\theta^t) \frac{U_{\theta Y}(\theta^t)/U_Y(\theta^t) - \tilde{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta Y}(\theta^{t+1})}{U_Y(\theta^{t+1})} \middle| \theta^t \right)}{U_{\theta C}(\theta^t)/U_C(\theta^t) - U_{\theta Y}(\theta^t)/U_Y(\theta^t)}.$$

The generalized IEE for earnings augments the tradeoff in intertemporal allocation of earnings by a tradeoff between current and future redistribution that is internalized by the planner but not by private decisions. This second tradeoff decomposes into the effect on future incentives through the incentive-adjusted probability or returns $\tilde{M}(\theta^{t+1})$, a direct effect on current redistribution, and the feedback from future to current redistribution. The latter two are captured by $\tilde{s}_Y(\theta^t)$.

Since redistribution based on ability lowers current information rents ($U_{\theta Y}/U_Y < 0$), $\tilde{M}(\cdot)$ is decreasing in θ_{t+1} . Therefore, if $-U_Y(\cdot)$ is decreasing in θ_{t+1} ,¹² $\mathbb{E} \left(-U_Y(\theta^{t+1}) \frac{\tilde{M}(\theta^{t+1})}{\mathbb{E}(\tilde{M}(\theta^{t+1})|\theta^t)} \middle| \theta^t \right) > \mathbb{E}(-U_Y(\theta^{t+1})|\theta^t)$, so the incentive adjustment favors backloading leisure, unless $n(Y, \theta)$ is linear in Y , or $\mathcal{E}_Y = 0$. On the other hand, if the direct effect dominates the indirect, then $\tilde{s}_Y(\theta^t) < 0$ which implies that it is optimal to backload earnings, and subsidize current leisure. Since the latter dominates on average, earnings become mean-reverting at the optimal allocation.

¹²Under the weak separability assumption, $-U_Y(\cdot)$ is decreasing in θ_2 whenever $n_Y(\cdot)$ is decreasing in θ_2 . This holds automatically whenever $Y(\cdot)$ is not too strongly increasing.

6 Optimal Labor Wedges

I now turn to the characterization of optimal labor wedges. Combining Proposition 1 with the dynamic optimality conditions yields the following recursive characterization of static labor wedges:

$$MC_t(\theta^t) = \widehat{MB}_t(\theta^t) + \hat{\varrho}_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1}) = \widetilde{MB}_t(\theta^t) + \tilde{\varrho}_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1})$$

The equation combines the recursive characterization of labor wedges from the case with separable preferences with the observation that the optimal allocation equates the marginal cost of efficiency distortions to the marginal benefit of redistribution through either consumption or earnings - and the two marginal benefits of redistribution must therefore also be equal to each other. In other words, the planner equates marginal costs and benefits of redistribution, and is indifferent between redistribution via consumption or via earnings at each realization of θ . The marginal benefits of redistribution include both a static and a dynamic component, which are based on the same arbitrage between smoothing tax distortions and backloading to internalize the future marginal benefit of redistribution as in the separable model.

Non-separability in preferences modifies the static marginal benefits $\widehat{MB}_t(\theta^t)$ and $\widetilde{MB}_t(\theta^t)$, and changes the rate of decay in information rents, $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$. The rate of decay of information rents need not be the same for redistribution via consumption or via leisure. Hence the optimal solution also trades off between current and past marginal benefits: if for example $\hat{\varrho}_t(\theta^t) > \tilde{\varrho}_t(\theta^t)$, then the optimal allocation will generate more redistribution via consumption than via leisure implying $\widehat{MB}_t(\theta^t) < \widetilde{MB}_t(\theta^t)$, and $\widehat{MB}_t(\theta^t)$ has more persistent effects on current taxes than $\widetilde{MB}_t(\theta^t)$. Solving the optimality condition backwards and expressing $MC_t(\theta^t)$ in terms of the labor tax $\tau_t(\theta^t)$ yields the following expressions:

$$\begin{aligned} \frac{\tau_t(\theta^t)}{U_{\theta Y}(\theta^t) + (1 - \tau_t(\theta^t)) U_{\theta C}(\theta^t)} &= \sum_{\tau=0}^{t-1} (\beta R)^\tau \prod_{s=0}^{\tau-1} \hat{\varrho}_{t-s}(\theta^{t-s}) \widehat{MB}_{t-\tau}(\theta^{t-\tau}) \\ &= \sum_{\tau=0}^{t-1} (\beta R)^\tau \prod_{s=0}^{\tau-1} \tilde{\varrho}_{t-s}(\theta^{t-s}) \widetilde{MB}_{t-\tau}(\theta^{t-\tau}). \end{aligned}$$

Hence, how the non-separability affects the rate of decay of information rents is key to understanding the persistence of labor wedges and the respective impact of marginal benefits of redistribution via consumption and earnings.

Decay of information rents: $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$ can be decomposed into (i) the feedback $\varrho_t(\theta^t)$ from future distortions $MC_t(\theta^t)$ to information rents at θ^{t-1} and (ii) the feedback from $\widehat{MB}_t(\theta^t)$ or $\widetilde{MB}_t(\theta^t)$ to information rents at θ^{t-1} .

Consider the same perturbation as in section 3, which marginally increases distortions and redistribution at θ^t while preserving expected utility and incentive compatibility. As before this perturbation generates a static cost $MC_t(\theta^t)$ and benefit $\widehat{MB}_t(\theta^t)$, and $MC_t(\theta^t)$ feeds into marginal information rents at θ^{t-1} at a rate $\varrho_t(\theta^t)$. But in addition, $\widehat{MB}_t(\theta^t)$ also feeds back into information rents at θ^{t-1} . If consumption is redistributed around θ_t in proportion to $\hat{m}(\theta', \theta^{t-1})/U_C(\theta', \theta^{t-1})$ to preserve incentive compatibility and expected utility, then marginal information rents for $\theta' \neq \theta_t$ change in proportion to $\hat{m}(\theta', \theta^{t-1})U_{\theta C}(\theta', \theta^{t-1})/U_C(\theta', \theta^{t-1})$, and these changes in turn feed back into information rents at θ^{t-1} at a rate $\varrho_t(\theta', \theta^{t-1})$. Hence the second term in the expression of $\hat{\varrho}_t(\theta^t)$ represents the feedback from $\widehat{MB}_t(\theta^t)$ to information rents at θ^{t-1} . Along the same lines, the second term in $\tilde{\varrho}_t(\theta^t)$ represents the feedback from $\widetilde{MB}_t(\theta^t)$ to information rents at θ^{t-1} .

The representation of $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$ in Proposition 1 leads to the following comparative statics result:

Proposition 3 : *Comparative statics for $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$:*

- (i) $\hat{\varrho}_t(\theta^t) \gtrless \varrho_t(\theta^t)$ if $\varrho_t(\theta, \theta^{t-1}) \frac{U_{\theta C}(\theta, \theta^{t-1})}{U_C(\theta, \theta^{t-1})}$ is increasing/constant/decreasing in θ .
- (ii) $\tilde{\varrho}_t(\theta^t) \gtrless \varrho_t(\theta^t)$ if $\varrho_t(\theta, \theta^{t-1}) \frac{U_{\theta Y}(\theta, \theta^{t-1})}{U_Y(\theta, \theta^{t-1})}$ is increasing/constant/decreasing in θ .
- (iii) $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$ converge to $\varrho_t(\theta^t)$ as θ_t converges to $\bar{\theta}$ or $\underline{\theta}$.
- (iv) $\hat{\varrho}_t(\theta^t) = \tilde{\varrho}_t(\theta^t) = \varrho_t(\theta^t)$ in the special case where $\frac{\partial}{\partial \theta} \varrho_t(\theta, \theta^{t-1}) = 0$ for all θ and $U_{\theta C}/U_C$ and $U_{\theta Y}/U_Y$ are constant.

Proposition 3 shows how the persistence of labor taxes $\hat{\varrho}_t(\theta^t)$ or $\tilde{\varrho}_t(\theta^t)$ differs from the decay of information rents $\varrho_t(\theta^t)$. The terms $\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)}$ and $\varrho_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}$ measure the feedback from consumption or earnings changes at θ^t to marginal information rents at θ^{t-1} . If this feedback is increasing in θ_t , then the feedback from $\widehat{MB}_t(\theta^t)$ to information rents at θ^{t-1} is positive and results in an optimal persistence of labor taxes that is higher than the decay of information rents. If instead the feedback from consumption or earnings changes to prior information rents is decreasing in the type realization, then the feedback from $\widetilde{MB}_t(\theta^t)$ to information rents at θ^{t-1} is negative and the optimal persistence of labor taxes that is less than the decay of information rents.

Two special cases deserve to be mentioned: first, at the top and the bottom of the distribution the optimal persistence of labor taxes converges to the decay of information rents since marginal benefits of redistribution vanish at the top and bottom of the distribution. Second, if the decay rate is uniform in θ_t ($\frac{\partial}{\partial \theta} \varrho_t(\theta, \theta^{t-1}) = 0$ for all θ), and $U_{\theta C}/U_C$ and $U_{\theta Y}/U_Y$ are constant, then the feedback is constant for all types, and therefore the optimal persistence of labor taxes is also uniform and equal to the decay of information rents. In these cases, future redistribution has

no effect on current information rents, the optimal allocation satisfies $\widehat{MB}_t(\theta^t) = \widetilde{MB}_t(\theta^t)$ and $MC_t(\theta^t) = \widehat{MB}_t(\theta^t) + \varrho_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1})$, i.e. the planner's solution equalizes static marginal benefits of redistribution and the persistence of labor tax distortions only depends on the exogenous decay of information rents.

Marginal Benefits of Redistribution: The marginal benefits of redistribution can be re-written as $\widehat{MB}_t(\theta^t) = \frac{1-F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})U_C(\theta^t)} \cdot \widehat{B}(\theta^t)$ and $\widetilde{MB}_t(\theta^t) = \frac{1-F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})(-U_Y(\theta^t))} \cdot \widetilde{B}(\theta^t)$, where

$$\begin{aligned}\widehat{B}(\theta^t) &= \mathbb{E} \left(\frac{\widehat{M}(\theta')}{\widehat{M}(\theta_t)} \mid \theta' \geq \theta_t, \theta^{t-1} \right) - \mathbb{E} \left(\frac{\widehat{M}(\theta')}{\widehat{M}(\theta_t)} \mid \theta^{t-1} \right) \frac{1 - \widehat{F}(\theta_t | \theta^{t-1})}{1 - F(\theta_t | \theta^{t-1})} \\ \widetilde{B}(\theta^t) &= \mathbb{E} \left(\frac{\widetilde{M}(\theta')}{\widetilde{M}(\theta_t)} \mid \theta' \geq \theta_t, \theta^{t-1} \right) - \mathbb{E} \left(\frac{\widetilde{M}(\theta')}{\widetilde{M}(\theta_t)} \mid \theta^{t-1} \right) \frac{1 - \widetilde{F}(\theta_t | \theta^{t-1})}{1 - F(\theta_t | \theta^{t-1})}\end{aligned}$$

When the two marginal benefits have equal persistence, the optimal labor wedge satisfies the following simple characterization:

Theorem 2 : *If $\hat{\varrho}_t(\theta^t) = \tilde{\varrho}_t(\theta^t) = \varrho_t(\theta^t)$, then the optimal labor wedge satisfies*

$$\frac{-U_Y(\theta^t)}{U_C(\theta^t)} = 1 - \tau_t(\theta^t) = \frac{\widetilde{B}(\theta^t)}{\widehat{B}(\theta^t)}$$

If instead $\hat{\varrho}_t(\theta^t) > \tilde{\varrho}_t(\theta^t)$, then $1 - \tau_t(\theta^t) < \widetilde{B}(\theta^t) / \widehat{B}(\theta^t)$, while if $\hat{\varrho}_t(\theta^t) < \tilde{\varrho}_t(\theta^t)$, then $1 - \tau_t(\theta^t) > \widetilde{B}(\theta^t) / \widehat{B}(\theta^t)$.

Theorem 2 shows that with equal persistence, the optimal labor wedge corresponds to the ratio of the marginal benefit of redistribution through earnings to the marginal benefit of redistribution through consumption. This condition captures the basic intuition that the planner arbitrages between asking the high income types to work more vs. asking them to consume less, and on the margin the planner's solution must be indifferent between the two. When persistence is not equal, then the ratio of the two marginal benefits serves as an upper or lower bound, given that the planner's solution trades off more static redistribution against higher persistence.

A special case of the theorem arises when types are iid ($\varrho_t(\theta^t) = 0$). In this case the optimal labor wedge only depends on static costs and benefits of redistribution, and the two static optimality conditions can be re-stated as follows:

$$\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} = A(\theta^t) \cdot \widehat{B}(\theta^t) \quad \text{and} \quad \tau_t(\theta^t) = A(\theta^t) \cdot \widetilde{B}(\theta^t)$$

$$\text{where } A(\theta^t) = \left(\frac{U_{C\theta}(\theta^t)}{U_C(\theta^t)} - \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} \right) \cdot \frac{1 - F(\theta_t | \theta^{t-1})}{\theta_t f(\theta_t | \theta^{t-1})}$$

The first optimality condition restates and generalizes the well-known ABC formula of Diamond (1998) and Saez (2001) to the present model which allows for redistribution based on both needs and abilities. The second optimality condition complements the first by obtaining the equivalent representation from marginal benefits of redistribution via earnings. The two optimality conditions offer a new interpretation to this representation as equating the marginal efficiency cost of redistribution to the marginal benefit of redistribution via consumption, since the terms $\frac{\tau_t(\theta^t)}{1-\tau_t(\theta^t)}/A(\theta^t)$ and $\tau_t(\theta^t)/A(\theta^t)$ map directly to the marginal cost of redistribution, rescaled by $U_C(\theta^t)$ or $-U_Y(\theta^t)$, and $\widehat{B}(\theta^t)$ and $\widetilde{B}(\theta^t)$ correspond to the re-scaled marginal benefits of redistribution via consumption or earnings.

By taking the ratio between the two marginal benefit terms I obtain the alternative representation given by Theorem 2, which allows me to unify the representation of optimal labor wedges in the static and dynamic Mirrlees model.

What's more, recall that in the iso-elastic case $\widehat{M}(\cdot)$ and $\widetilde{M}(\cdot)$ directly derive from the distribution of consumption and income. Hence $\widehat{B}(\theta^t)$ and $\widetilde{B}(\theta^t)$ are robustly tied to observable statistics and estimates of the elasticities $\mathcal{E}_C(\cdot)$ and $\mathcal{E}_Y(\cdot)$ which relate to income and substitution effects of labor supply. In Hellwig and Werquin (2021), we tie this alternative representation of optimal income taxes to the sufficient statistics approach of Saez (2001).

The above representation also determines how optimal taxes and marginal costs and benefits of redistribution vary with primitive parameters. If $U(C, Y; \theta) = \mathcal{U}(\gamma(C, \underline{C}(\theta)) - n(Y, \theta))$, then $\frac{U_{C\theta}}{U_C} = \frac{\mathcal{U}''}{(\mathcal{U}')^2}U_\theta + \frac{\chi'}{\chi}$ and $\frac{U_{Y\theta}}{U_Y} = \frac{\mathcal{U}''}{(\mathcal{U}')^2}U_\theta - \frac{-p'}{p}$, where $\frac{\chi'(\cdot)}{\chi(\cdot)}$ and $\frac{p'}{p}$ were defined in the text following proposition 2. Therefore fixing the allocation, an increase in $\frac{\chi'(\cdot)}{\chi(\cdot)}$ increases $\frac{U_{C\theta}}{U_C}$, making consumption needs more regressive or less progressive. An increase in $\frac{-p'(\cdot)}{p(\cdot)}$ makes productivities or disutilities of effort more sensitive to type, which makes $\frac{U_{Y\theta}}{U_Y}$ more negative. An increase in $\frac{-\mathcal{U}''}{\mathcal{U}'}$ increases the outer curvature or overall risk aversion and reduces both $\frac{U_{C\theta}}{U_C}$ and $\frac{U_{Y\theta}}{U_Y}$ by the same magnitude. I also consider comparative statics w.r.t. $\mathcal{E}_C(\cdot)$ and $\mathcal{E}_Y(\cdot)$ which only enter through their respective marginal benefits.

Proposition 4 *For a given allocation of consumption $C(\cdot)$ or earnings $Y(\cdot)$:*

(i) *The (rescaled) marginal costs $A(\theta)^{-1}$ is decreasing in $\frac{\chi'(\cdot)}{\chi(\cdot)}$ and $\frac{-p'(\cdot)}{p(\cdot)}$, and independent of $\mathcal{E}_C(\cdot)$, $\mathcal{E}_Y(\cdot)$, and $\frac{-\mathcal{U}''}{\mathcal{U}'}$.*

(ii) *The (rescaled) marginal benefit of redistribution via consumption $\widehat{B}(\cdot)$ is increasing in $\mathcal{E}_C(\cdot)$ and $\frac{-\mathcal{U}''}{\mathcal{U}'}$, decreasing in $\frac{\chi'(\cdot)}{\chi(\cdot)}$, and independent of $\mathcal{E}_Y(\cdot)$ and $\frac{-p'(\cdot)}{p(\cdot)}$.*

(iii) *The (rescaled) marginal benefit of redistribution via earnings $\widetilde{B}(\cdot)$ is increasing in $\frac{-\mathcal{U}''}{\mathcal{U}'}$ and $\frac{-p'(\cdot)}{p(\cdot)}$, decreasing in $\mathcal{E}_Y(\cdot)$, and independent of $\mathcal{E}_C(\cdot)$ and $\frac{\chi'(\cdot)}{\chi(\cdot)}$.*

Proposition 4 shows how non-separabilities in preferences alter the tradeoff between efficiency and redistribution. It translates the stochastic dominance ordering of $\hat{F}(\cdot)$, $F(\cdot)$ and $\tilde{F}(\cdot)$, as well as the elasticities $\mathcal{E}_C(\cdot)$ and $\mathcal{E}_Y(\cdot)$ into comparative statics of re-scaled marginal costs and benefits of redistribution for a given type-contingent consumption or earnings profile. These comparative statics then identify welfare-improving changes to the optimal allocation that can be implemented through a combination of the perturbations that were presented above:

If $\mathcal{E}_C(\cdot)$ increases, the marginal benefits of redistributing consumption increase, and the optimal allocation therefore shifts towards higher labor taxes, more redistribution via consumption and less redistribution via earnings. If $\mathcal{E}_Y(\cdot)$ increases, the marginal benefits of redistributing earnings are reduced, so the optimal allocation shifts towards less redistribution of earnings and more redistribution of consumption, while also opting for labor taxes.¹³

If $\chi'(\cdot)/\chi(\cdot)$ increases, it is optimal to increase redistribution via earnings and compensate with a combination of higher labor taxes and/or less redistribution via consumption, in line with the intuition that an increase in $\chi'(\cdot)/\chi(\cdot)$ results in a less progressive or more regressive motive of redistribution based on consumption needs.

If $-p'(\cdot)/p(\cdot)$ increases, productivities or disutilities of effort are more dispersed so the gains from redistribution via earnings increase. It's then optimal to increase labor taxes and redistribution via earnings. Redistribution via consumption may increase or decrease depending on whether the combined effect of the tax increase and the increase in $a'(\cdot)/a(\cdot)$ results in higher or lower marginal costs of efficiency distortions.

If $-\mathcal{U}''/\mathcal{U}'$ increases, additional curvature in utility strengthens both redistribution motives. The planner gains from increasing labor taxes and redistribution via both consumption and earnings.

To recap, increases in $\mathcal{E}_C(\cdot)$, $-\mathcal{U}''/\mathcal{U}'$ and $-p'(\cdot)/p(\cdot)$ or a reduction in $\mathcal{E}_Y(\cdot)$ unambiguously shift the tradeoff between efficiency and redistribution towards higher taxes and more redistribution, but the effect of $\chi'(\cdot)/\chi(\cdot)$, which captures changes in the motive for redistribution based on consumption needs, is more subtle: An increase in $\chi'(\cdot)/\chi(\cdot)$ reduces both the marginal benefits of redistribution based on consumption and the marginal cost of efficiency distortions. This leads to more redistribution via earnings, but whether this is compensated by a reduction of redistribution via consumption or an increase in taxes, or a combination of both, is ambiguous and depends on the relative effect of $\chi'(\cdot)/\chi(\cdot)$ on the marginal cost of efficiency distortions and the marginal benefit of consumption-based redistribution.

¹³This comparative static of the labor supply elasticity $\mathcal{E}_Y(\cdot)$ holds $\frac{-p'(\cdot)}{p(\cdot)}$ constant. In practice, an increase in the Frisch elasticity would also change productivities $\frac{-p'(\cdot)}{p(\cdot)}$.

Optimal Persistence of Labor wedges: My final result links optimal tax-smoothing to the sign and magnitude of the savings wedge, which completes the connection between the two wedges. Theorem 3 generalizes the tax-smoothing representation of Farhi and Werning (2013) to non-separable preferences and provides an analogous representation based on redistribution through earnings.¹⁴

Theorem 3 : *The optimal labor wedge satisfies the following two recursive characterizations:*

(i) *Based on marginal benefits of redistribution through consumption:*

$$\hat{\mathbb{E}} \left(\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \frac{1/U_C(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) = \widehat{Cov} \left(\log \left(\frac{\hat{m}(\theta^t)}{\hat{m}(\theta^t)} \right), \frac{1/U_C(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) \\ + \frac{1}{1 + \hat{s}_C(\theta^{t-1})} \hat{\mathcal{R}}(\theta^{t-1}) \frac{\tau_{t-1}(\theta^{t-1})}{1 - \tau_{t-1}(\theta^{t-1})}$$

$$\text{where } \hat{\mathcal{R}}(\theta^{t-1}) = \frac{\hat{\mathbb{E}} \left(\hat{\varrho}_t(\theta^t) \cdot \left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right) | \theta^{t-1} \right)}{\frac{U_{\theta C}(\theta^{t-1})}{U_C(\theta^{t-1})} - \frac{U_{\theta Y}(\theta^{t-1})}{U_Y(\theta^{t-1})}}$$

(ii) *Based on marginal benefits of redistribution through earnings:*

$$\tilde{\mathbb{E}} \left(\tau_t(\theta^t) \frac{1/(-U_Y(\theta^t))}{\tilde{\mathbb{E}}(1/(-U_Y(\theta^t)) | \theta^{t-1})} | \theta^{t-1} \right) = \widetilde{Cov} \left(\log \left(\frac{\hat{m}(\theta^t)}{\hat{m}(\theta^t)} \right), \frac{1/(-U_Y(\theta^t))}{\tilde{\mathbb{E}}(1/(-U_Y(\theta^t)) | \theta^{t-1})} | \theta^{t-1} \right) \\ + \frac{1}{1 + \tilde{s}_Y(\theta^{t-1})} \tilde{\mathcal{R}}(\theta^{t-1}) \tau_t(\theta^{t-1})$$

$$\text{where } \tilde{\mathcal{R}}(\theta^{t-1}) = \frac{\tilde{\mathbb{E}} \left(\tilde{\varrho}_t(\theta^t) \cdot \left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right) | \theta^{t-1} \right)}{\frac{U_{\theta C}(\theta^{t-1})}{U_C(\theta^{t-1})} - \frac{U_{\theta Y}(\theta^{t-1})}{U_Y(\theta^{t-1})}}$$

This representation incorporates two changes relative to the benchmark with separable preferences. First, expected future tax distortions and marginal benefits of redistribution are computed using incentive-adjusted probability measures. Expected future tax distortions can be re-stated as $\hat{\mathbb{E}} \left(\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \frac{1/U_C(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) = \mathbb{E} \left(\frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \frac{\hat{M}(\theta^t)}{\mathbb{E}(\hat{M}(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right)$ and $\tilde{\mathbb{E}} \left(\tau_t(\theta^t) \frac{1/(-U_Y(\theta^t))}{\tilde{\mathbb{E}}(1/(-U_Y(\theta^t)) | \theta^{t-1})} | \theta^{t-1} \right) = \mathbb{E} \left(\tau_t(\theta^t) \frac{\tilde{M}(\theta^t)}{\mathbb{E}(\tilde{M}(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right)$, which re-weights the distortions according to the incentive-adjusted returns $\hat{M}(\theta^t)$ or $\tilde{M}(\theta^t)$. The expected marginal benefits of redistribution can be re-stated along similar lines.

Second the persistence of the labor wedge is determined not just by the tradeoff between future and current information rents that is captured by $\hat{\mathcal{R}}(\theta^{t-1})$ and $\tilde{\mathcal{R}}(\theta^{t-1})$, but also by the

¹⁴The dynamic tax formula of Golosov, Troshkin and Tsyvinski (2016) follows from the recursive optimality condition along the same lines as in the separable case.

optimal savings wedge $1 + \hat{s}_C(\theta^{t-1})$ or $1 + \tilde{s}_Y(\theta^{t-1})$. The savings wedge enters optimal tax-smoothing because the planner's solution discounts consumption between θ^t and θ^{t-1} by a factor $R(1 + \hat{s}_C(\theta^{t-1})) \frac{\hat{M}(\theta^t)}{\mathbb{E}(\hat{M}(\theta^t)|\theta^{t-1})}$ and earnings by a factor $R(1 + \tilde{s}_Y(\theta^{t-1})) \frac{\tilde{M}(\theta^t)}{\mathbb{E}(\tilde{M}(\theta^t)|\theta^{t-1})}$, while the private return on savings is given by R . The planner's solution internalizes dynamic tradeoffs between redistribution at θ^t and θ^{t-1} that are not internalized by private savings decisions. The same wedge enters the resource tradeoff between current and future tax distortions. The savings wedge can thus either reinforce or dampen labor-tax smoothing: a savings tax increases the persistence of labor taxes, while a savings subsidy reduces it. Which one arises is determined as before by the net impact of savings on current information rents.

As before, the terms $\hat{\mathcal{R}}(\theta^{t-1})$ and $\tilde{\mathcal{R}}(\theta^{t-1})$ smooth labor wedges inter-temporally in such a way that a marginal transfer of distortions leaves current information rents unchanged. As shown in proposition 3 labor wedges become more (less) persistent when feedback from future to current information rents is increasing (decreasing) in θ_t .

Through $1 + \hat{s}_C(\theta^{t-1})$ and $\hat{\mathcal{R}}(\theta^{t-1})$, as well as $1 + \tilde{s}_Y(\theta^{t-1})$ and $\tilde{\mathcal{R}}(\theta^{t-1})$, the persistence of labor taxes thus depends on the sign, level and slope of $U_{\theta C}/U_C$ and $U_{\theta Y}/U_Y$, which govern the strength and direction of the need-based and ability based redistribution motives.

7 Examples

In this section I present four examples to illustrate the different possibilities that the main results allow for: (i) with GHH preferences, the wealth effect on incentives disappears, and the characterization of savings and labor distortions are uniquely determined from static concerns about redistribution. (ii) With preferences that incorporate type dependent marginal utilities or consumption needs (subsistence consumption), the optimal savings tax or subsidy is linked to how consumption needs vary with type. (iii) In an extension with type-dependent volatility, the persistence of labor wedges becomes state-dependent.

1. GHH Preferences: Suppose that $U(C, Y; \theta) = \mathcal{U}\left(C - \kappa(Y/A(\theta))^{1+\mathcal{E}_Y}\right)$, $A(\theta) = e^{\phi\theta}$ with $\phi > 0$, and $\frac{-\mathcal{U}''(\cdot)}{\mathcal{U}'(\cdot)}(C - n) = \psi > 0$. Then, the incentive-adjusted probability takes the form $\hat{m}(\theta^t) = \mathcal{U}'(\theta^t)$, and $\hat{M}(\theta^t) = 1$, i.e. the wealth effect on incentives disappears and the adjustment only incorporates risk aversion due to $\mathcal{U}(\cdot)$. It is straight-forward to check that $\frac{U_{\theta C}}{U_C} - \frac{U_{\theta Y}}{U_Y} = (1 + \mathcal{E}_Y)\phi > 0$ and $\frac{U_{\theta C}}{U_C} = -Z(\theta^t)\psi(1 + \mathcal{E}_Y)\phi$, where $Z(\theta^t) = n(\theta^t)/(C(\theta^t) - n(\theta^t))$. Then the Generalized IEE reduces to a standard Euler equation with a savings wedge:

$$\mathcal{U}'(\theta^t) = \beta R(1 + \hat{s}_C(\theta^t)) \mathbb{E}(\mathcal{U}'(\theta^{t+1}) | \theta^t),$$

where

$$\hat{s}_C(\theta^{t-1}) = -\frac{\tau_t(\theta^t)}{1-\tau_t(\theta^t)}\psi\left(Z(\theta^t) - \hat{\mathbb{E}}(\varrho_{t+1}(\theta^{t+1})Z(\theta^{t+1})|\theta^t)\right).$$

If $Z(\theta^t) > \hat{\mathbb{E}}(\varrho_{t+1}(\theta^{t+1})Z(\theta^{t+1})|\theta^t)$, it is optimal to tax savings. The recursive tax-smoothing equation yields

$$\frac{\tau_t(\theta^t)}{1-\tau_t(\theta^t)} = (1+\mathcal{E}_Y)\phi\frac{\hat{F}(\theta_t|\theta^{t-1})-F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})} + \frac{\mathcal{U}'(\theta^t)}{\mathbb{E}(\mathcal{U}'(\theta^t)|\theta^{t-1})}\frac{\hat{\varrho}_t(\theta^t)}{1+\hat{s}_C(\theta^{t-1})}\frac{\tau_{t-1}(\theta^{t-1})}{1-\tau_{t-1}(\theta^{t-1})},$$

where the optimal decay of information rents is

$$\hat{\varrho}_t(\theta^t) = \varrho_t(\theta^t) - \frac{(1+\mathcal{E}_Y)\phi\psi}{f(\theta_t|\theta^{t-1})}\int_{\theta_t}^{\bar{\theta}}\left(\varrho_t(\theta',\theta^{t-1})Z(\theta',\theta^{t-1}) - \hat{\mathbb{E}}(\varrho_t(\theta^t)Z(\theta^t)|\theta^{t-1})\right)\hat{f}(\theta'|\theta^{t-1})d\theta'.$$

Taking expectations of the recursive tax-smoothing equation yields the Farhi-Werning representation where $\mathbb{E}\left(\frac{\hat{F}(\theta_t|\theta^{t-1})-F(\theta_t|\theta^{t-1})}{f(\theta_t|\theta^{t-1})}|\theta^{t-1}\right) = -Cov\left(\theta_t, \frac{\mathcal{U}'(\theta^t)}{\mathbb{E}(\mathcal{U}'(\theta^t)|\theta^{t-1})}|\theta^{t-1}\right)$.

As these expressions show, with GHH preferences the adjusted probability measure only includes a risk component $\frac{\mathcal{U}'(\theta^t)}{\mathbb{E}(\mathcal{U}'(\theta^t)|\theta^{t-1})}$, but the incentive component $\hat{M}(\cdot)$ disappears. If types are not too persistent, risk aversion introduces a motive for taxing savings and making labor taxes more persistent than at the separable benchmark. If in addition $Z(\theta^t)$ is decreasing and $\varrho_t(\theta^t)$ independent of θ_t , the GHH preferences result in higher feedback of information rents and higher optimal persistence of the labor wedge.¹⁵

2. Isoelastic preferences: Suppose preferences take the form

$$U(C, Y; \theta) = \frac{1}{1-\mathcal{E}_C}\left(\frac{C}{\underline{C}(\theta)}\right)^{1-\mathcal{E}_C} - \kappa(Y/A(\theta))^{1+\mathcal{E}_Y},$$

with $\mathcal{E}_C \geq 1$ and $\mathcal{E}_Y > 0$.¹⁶ Suppose also that $A(\theta) = e^{\phi\theta}$, $\phi > 0$, and $\underline{C}(\theta) = \underline{C}e^{\Gamma\theta}$ where Γ denotes the sensitivity of consumption needs to types, and Γ can be positive or negative. It follows that $\hat{m}(\theta^t) = \underline{C}(\theta_t)^{\mathcal{E}_C-1}$, $\tilde{m}(\theta^t) = A(\theta_t)^{-(1+\mathcal{E}_Y)}$, $\hat{M}(\theta^t) = C(\theta^t)^{\mathcal{E}_C}$, and $\tilde{M}(\theta^t) = Y(\theta^t)^{-\mathcal{E}_Y}$. The Generalized Inverse Euler Equation take the form

$$C(\theta^t)^{-\mathcal{E}_C}\underline{C}(\theta_t)^{\mathcal{E}_C-1} = \beta R(1+\hat{s}_C(\theta^t))\frac{\mathbb{E}\left(\underline{C}(\theta_{t+1})^{\mathcal{E}_C-1}|\theta^t\right)}{\mathbb{E}\left(C(\theta_{t+1})^{\mathcal{E}_C}|\theta^t\right)},$$

where the savings wedges satisfy

$$\hat{s}_C(\theta^t) = \frac{\tau_t(\theta^t)}{1-\tau_t(\theta^t)}\frac{(\mathcal{E}_C-1)\Gamma\left(1-\hat{\mathbb{E}}(\varrho_{t+1}(\theta^{t+1})|\theta^t)\right)}{(1+\mathcal{E}_Y)\phi+(\mathcal{E}_C-1)\Gamma}$$

¹⁵ $Z(\theta^t)$ is decreasing in θ_t whenever $n(\theta^t)$ is decreasing in θ_t .

¹⁶ It is possible to rewrite these preferences as $U(C, Y; \theta) = \underline{C}(\theta)^{\mathcal{E}_C-1}\left\{\frac{1}{1-\mathcal{E}_C}C^{1-\mathcal{E}_C} - \kappa(Y/\bar{A}(\theta))^{1+\mathcal{E}_Y}\right\}$, where $\bar{A}(\theta) = \underline{C}(\theta)^{\frac{\mathcal{E}_C-1}{1+\mathcal{E}_Y}}A(\theta)$. Hence the iso-elastic preference model has an equivalent reinterpretation as including a shock to time preference rates.

Hence it is optimal to subsidize savings if $1 > \hat{\mathbb{E}}(\varrho_{t+1}(\theta^{t+1})|\theta^t)$, and to tax them otherwise. Rewriting the IEEs yields:

$$\mathbb{E}\left(\left(\frac{C(\theta^{t+1})}{C(\theta^t)}\right)^{\mathcal{E}_C}|\theta^t\right) = \beta R(1 + \hat{s}_C(\theta^t))\mathbb{E}\left(\left(\frac{\underline{C}(\theta_{t+1})}{\underline{C}(\theta_t)}\right)^{\mathcal{E}_C-1}|\theta^t\right),$$

Therefore expected consumption growth is an increasing function of the savings subsidy $\hat{s}_C(\theta^t)$ and the expected growth in consumption needs $\underline{C}(\theta_{t+1})/\underline{C}(\theta_t)$. Then, if types are fully persistent ($\varrho_{t+1}(\theta^{t+1}) = 1$) or consumption needs independent of type ($\Gamma = 0$), the savings subsidy is 0 and consumption growth is independent of the current type, resulting in divergence of consumption profiles. These dynamics are reinforced when $\Gamma < 0$, which leads to a savings tax that frontloads consumption and reduces future consumption growth. In addition consumption growth is increasing in θ_t resulting in even more divergence and polarization of consumption over time.

If instead $\Gamma > 0$ and $\hat{\mathbb{E}}(\varrho_{t+1}(\theta^{t+1})|\theta^t) < 1$, then it is optimal to subsidize savings and consumption needs introduce a force towards mean-reversion. The resulting consumption process is mean-reverting around a positive growth trend. With $\Gamma < 0$ on the other hand, the savings tax frontloads consumption and lowers average consumption growth. Furthermore, consumption growth is increasing in θ_t resulting in more divergence of consumption over time.

3. Subsistence consumption needs: Suppose preferences take the form

$$U(C, Y; \theta) = \frac{1}{1-\chi} (C - \underline{C}(\theta))^{1-\chi} - \kappa (Y/A(\theta))^{1+\mathcal{E}_Y},$$

with $\chi \geq 0$, $A(\theta) = e^{\phi\theta}$, $\phi > 0$, and $\underline{C}(\theta) = \underline{C}e^{\Gamma\theta}$ where Γ denotes as before the sensitivity of consumption needs to types, and Γ can be positive or negative. It then follows that $\frac{U_{\theta C}}{U_C} = \Gamma(\mathcal{E}_C(\theta^t) - \chi)$, $\frac{U_{\theta Y}}{U_Y} = (1 + \mathcal{E}_Y)\phi$, $\hat{m}(\theta^t) = e^{\Gamma \int_{\underline{\theta}}^{\theta} (\mathcal{E}_C(\theta', \theta^{t-1}) - \chi) d\theta'}$ and $\hat{M}(\theta^t) = e^{\int_{\underline{\theta}}^{\theta} \mathcal{E}_C(\theta', \theta^{t-1}) d\theta'}$, where $\mathcal{E}_C(\theta^t) = \chi \frac{C(\theta^t)}{C(\theta^t) - \underline{C}(\theta^t)}$. The Generalized IEE takes the form

$$(C(\theta^t) - \underline{C}(\theta^t))^{-\chi} = \beta R(1 + \hat{s}_C(\theta^t)) \left\{ \hat{\mathbb{E}}\left(\left(C(\theta^{t+1}) - \underline{C}(\theta^{t+1})\right)^\chi |\theta^t\right) \right\}^{-1}$$

where the savings wedge takes the form

$$\hat{s}_C(\theta^t) = \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \Gamma \frac{\mathcal{E}_C(\theta^t) - \chi - \hat{\mathbb{E}}(\varrho_t(\theta^t)(\mathcal{E}_C(\theta^{t+1}) - \chi) |\theta^t)}{(1 + \mathcal{E}_Y)\phi + \Gamma(\mathcal{E}_C(\theta^t) - \chi)}$$

Hence if $\Gamma(\mathcal{E}_C(\theta^t) - \chi - \hat{\mathbb{E}}(\varrho_t(\theta^t)(\mathcal{E}_C(\theta^{t+1}) - \chi) |\theta^t)) < 0$ it is optimal to tax savings, in the opposite case it is optimal to subsidize them.

Suppose that $C(\cdot)/\underline{C}(\cdot)$ is increasing in θ_t and hence $\mathcal{E}_C(\theta^t)$ is decreasing in θ_t . This holds whenever Γ is negative or at most small and positive, i.e. consumption needs are decreasing or not

too strongly increasing in θ_t . Then, if $\Gamma < 0$, the persistence parameters satisfy $\hat{\varrho}_t(\theta^t) > \tilde{\varrho}_t(\theta^t) = \varrho_t(\theta^t)$, and therefore $\widehat{MB}_t(\theta^t) < \widetilde{MB}_t(\theta^t)$, i.e. at the optimal allocation, marginal benefits of redistribution through consumption are smaller, but more persistent than marginal benefits of redistribution through earnings - the decreasing consumption needs strengthen the planner's motive to redistribute consumption towards lower types, along with a positive tax on savings. If instead $\Gamma > 0$, the opposite conclusion holds ($\hat{\varrho}_t(\theta^t) < \tilde{\varrho}_t(\theta^t) = \varrho_t(\theta^t)$ and $\widehat{MB}_t(\theta^t) > \widetilde{MB}_t(\theta^t)$), i.e. the optimal allocation shifts towards more redistribution via earnings, and it becomes optimal to subsidize savings.

As in the case with GHH preferences, the endogenous persistence and the savings wedge are mutually reinforcing, leading to higher $\hat{\varrho}_t(\theta^t)$ and positive savings taxes when $\Gamma < 0$. But the two effects can also result in savings subsidies and less persistence than at the separable benchmark if $\Gamma > 0$ (consumption needs are increasing in type).

4. Type-dependent volatility: Suppose as in example 2 that preferences are isoelastic,

$$U(C, Y; \theta) = \frac{1}{1 - \mathcal{E}_C} \left(\frac{C}{\underline{C}(\theta)} \right)^{1 - \mathcal{E}_C} - \kappa (Y/A(\theta))^{1 + \mathcal{E}_Y},$$

with $\mathcal{E}_C \geq 1$, $A(\theta) = e^{\phi\theta}$, $\phi > 0$, and $\underline{C}(\theta) = \underline{C}e^{\Gamma\theta}$ where Γ can be positive or negative. Suppose that $\theta_t = \mu(\theta_{t-1}) + \sigma(\theta_{t-1})v_t$, where v_t is iid over time, and $\mu'(\theta_t) \in [0, 1]$. In this case, $\varrho_t(\theta^t)$ is no longer uniform for all θ_t , but instead

$$\varrho_t(\theta^t) = \mu'(\theta_{t-1}) + \sigma'(\theta_{t-1}) \frac{\theta_t - \mu(\theta_{t-1})}{\sigma(\theta_{t-1})}$$

Therefore, $\varrho_t(\theta^t)$ is increasing in θ_t if $\sigma'(\theta_{t-1}) > 0$ (higher types face higher uncertainty), and $\varrho_t(\theta^t)$ is decreasing in θ_t if $\sigma'(\theta_{t-1}) < 0$ (higher types face less uncertainty). We then have

$$\begin{aligned} \hat{\varrho}_t(\theta^t) &= \varrho_t(\theta^t) + \mathcal{E}_C \Gamma \frac{1 - \hat{F}(\theta_t|\theta^{t-1})}{\hat{f}(\theta_t|\theta^{t-1})} \frac{\sigma'(\theta_{t-1})}{\sigma(\theta_{t-1})} \left\{ \hat{\mathbb{E}}(\theta_t|\theta \geq \theta_t, \theta^{t-1}) - \hat{\mathbb{E}}(\theta_t|\theta^{t-1}) \right\} \\ \tilde{\varrho}_t(\theta^t) &= \varrho_t(\theta^t) + (1 + \mathcal{E}_Y) \phi \frac{1 - \tilde{F}(\theta_t|\theta^{t-1})}{\tilde{f}(\theta_t|\theta^{t-1})} \frac{\sigma'(\theta_{t-1})}{\sigma(\theta_{t-1})} \left\{ \tilde{\mathbb{E}}(\theta_t|\theta \geq \theta_t, \theta^{t-1}) - \tilde{\mathbb{E}}(\theta_t|\theta^{t-1}) \right\} \end{aligned}$$

Proposition 3 then implies that $\tilde{\varrho}_t(\theta^t) < \varrho_t(\theta^t)$ if $\sigma'(\theta_{t-1}) < 0$ and $\tilde{\varrho}_t(\theta^t) > \varrho_t(\theta^t)$ if $\sigma'(\theta_{t-1}) > 0$, while the sign of $\hat{\varrho}_t(\theta^t) - \varrho_t(\theta^t)$ depends on $\Gamma\sigma'(\theta_{t-1})$. This example shows that type dependent volatilities also contribute the persistence of the labor wedge: the feedback of information rents is increasing with the uncertainty about the current type realization, which increases persistence in labor wedges when $\Gamma > 0$ and $\sigma'(\theta_{t-1}) > 0$.

8 Conclusion

This paper aims to provide a general analysis of optimal dynamic taxation with non-separable preferences, and explore to what extent the core insights from the existing literature survive or generalize. First, I have shown how to incorporate non-separability in the analysis by means of a simple change in probability measures that captures the need for regressive redistribution of resources, and/or either progressive or regressive redistribution of utilities, as a means to preserve incentive compatibility. This incentive adjustment depends on preference parameters such as risk aversion, labor supply elasticities and consumption needs.

Second, by applying this incentive adjustment to the optimality condition for redistribution through both consumption and earnings, I obtain a double representation of the optimal labor and savings wedges, which completes the existing representations that are solely based on redistribution through consumption. This double representation captures the basic intuition that the planner can transfer utility from high to low types by transferring either consumption or leisure, and at the optimal allocation, the planner is indifferent between the two.

Third, I have shown how non-separability generates feedback from future allocations to current information rents. This feedback generates a new inter-temporal tradeoff between current and future redistribution and tax distortions which in turn modifies the main results from the existing literature on optimal savings and labor tax distortions. Optimal savings taxes may be positive or negative, depending on how savings internalize a tradeoff between current and future redistribution. The dynamics of labor taxes inherit the same savings wedge that determines how the planner's solution discounts resources from one period to the next. Finally the optimal persistence of labor taxes also depends on how tax-smoothing feeds back into current information rents.

One stark benchmark result stands out: when preferences are "sufficiently" isoelastic and types are highly persistent then the rationale for savings taxes or subsidies again disappears, and labor taxes are highly persistent with an upwards drift that is independent of assumptions about agent's preferences. Whether or not the theoretical results presented here also provide a strong quantitative case for taxing or subsidizing savings then depends on whether this benchmark comes close to reality or not. Answering this question, or related ones that bring the present results closer to applications, is an important direction for future work.

Finally, my results do not address the issue of tax implementation. In Hellwig (2021), I propose an implementation using history-dependent labor taxes with no private savings, and show that private savings decisions do not matter for the implementation of the optimal (or more generally,

any) incentive compatible allocation, but only serve to fulfill future tax obligations, extending the principle of Ricardian Equivalence to dynamic Mirrlees models. Even when savings are allowed the resulting labor and savings taxes need not match the wedges of the present characterization one-for-one: with dynamic information rents, the static labor supply decision includes a forward-looking element that alters the mapping from static labor wedges to tax implementations. Future work will have to explore the importance of these dynamic information rents for optimal tax design.

References

- [1] Aguiar, Mark, and Eric Hurst (2005). “Consumption versus Expenditure.” *Journal of Political Economy*, Vol 113(5), 919-948.
- [2] Aguiar, Mark, and Eric Hurst (2007). “Life-Cycle Prices and Production.” *The American Economic Review*, Vol 97(5), 1533-1559.
- [3] Aguiar, Mark, Eric Hurst (2007), and Loukas Karabarbounis (2012). “Recent Developments in the Economics of Time Use.” *Annual Review of Economics*, Vol 4(1), 373-397.
- [4] Albanesi, Stefania, and Christopher Sleet (2006). “Dynamic Optimal Taxation With Private Information.” *The Review of Economic Studies*, Vol 73(1), 1-30.
- [5] Atkeson, Andrew, and Robert Lucas (1992). “On Efficient Distribution with Private Information.” *The Review of Economic Studies*, Vol 59(3), 427-453.
- [6] Benhabib, Jess, Richard Rogerson, and Randall Wright (1991). “Homework in Macroeconomics: Household Production and Aggregate Fluctuations.” *Journal of Political Economy*, Vol 99(6), 1166-1187.
- [7] Brendon, Charles (2013). “Efficiency, Equity, and Optimal Income Taxation”, working paper, Cambridge University.
- [8] Diamond, Peter (1998). “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates.” *The American Economic Review*, Vol 88(1), 83-95.
- [9] Diamond, Peter and James Mirrlees (1978). “A Model of Social Insurance With Variable Retirement.” *Journal of Public Economics*, Vol 10(3), 295-336.
- [10] Farhi, Emmanuel and Ivan Werning (2007). “Inequality and Social Discounting.” *Journal of Political Economy*, Vol 117(3), 365-402.

- [11] Farhi, Emmanuel and Ivan Werning (2011). “Capital Taxation: Quantitative Explorations of the Inverse Euler Equation.” *Journal of Political Economy*, Vol 120(3), 398-445.
- [12] Farhi, Emmanuel and Ivan Werning (2013). “Insurance and Taxation over the Life Cycle.” *The Review of Economic Studies*, Vol 80(2), 596-635.
- [13] Golosov, Mikhail, Narayana Kocherlakota and Aleh Tsyvinski (2003). “Optimal Indirect and Capital Taxation.” *The Review of Economic Studies*, Vol 70(3), 569-587.
- [14] Golosov, Mikhail, Maxim Troshkin and Aleh Tsyvinski (2016). “Redistribution and Social Insurance.” *The American Economic Review*, Vol 106(2), 359-386.
- [15] Golosov, Mikhail, and Aleh Tsyvinski (2006). “Designing Optimal Disability Insurance: A Case for Asset Testing.” *Journal of Political Economy*, Vol 114(2), 257-279.
- [16] Grochulski, Borys, and Narayana Kocherlakota (2010). “Nonseparable Preferences and Optimal Social Security Systems.” *Journal of Economic Theory*, Vol 145(6), 2055-2077.
- [17] Hellwig, Martin (2007). “A Contribution to the Theory of Optimal Utilitarian Income Taxation.” *Journal of Public Economics*, Vol 91(7-8), 1449-1477.
- [18] Hellwig, Christian (2021). “Barro meets Mirrlees: Ricardian Equivalence in Dynamic Private Information Economies”, working paper, Toulouse School of Economics.
- [19] Hellwig, Christian, and Nicolas Werquin (2021). “Sufficient statistics for optimal income taxes in dynamic economies”, work in progress, Toulouse School of Economics
- [20] Kapicka, Marek (2013). “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-order Approach.” *The Review of Economic Studies*, Vol 80(3), 1027-1054.
- [21] Kocherlakota, Narayana (2005). “Zero Expected Wealth Taxes: A Mirrlees Approach to Optimal Dynamic Taxation.” *Econometrica*, Vol 73(5), 1587-1621.
- [22] Koehne, Sebastian (2018). “On the Taxation of Durable Goods.” *International Economic Review*, Vol 59(2), 825-857.
- [23] Koehne, Sebastian and Moritz Kuhn (2015). “Optimal Taxation in a Habit Formation Economy.” *Journal of Public Economics*, Vol 122, 31-39.

- [24] Mirrlees, James A. (1971). “An Exploration in the Theory of Optimal Income Taxation.” *The Review of Economic Studies*, Vol 38(2), 175-208.
- [25] Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014). “Dynamic Mechanism Design; A Myersonian Approach.” *Econometrica*, Vol 82(2), 601-653.
- [26] Phelan, Christopher (2006). “Opportunity and Social Mobility.” *The Review of Economic Studies*, Vol 73(2), 487-504.
- [27] Rogerson, William (1985). “Repeated Moral Hazard.” *Econometrica*, Vol 53(1), 69-76.
- [28] Saez, Emmanuel (2001). “Using Elasticities to Derive Optimal Income Tax Rates.” *The Review of Economic Studies*, Vol 68(1), 205-229.
- [29] Stantcheva, Stefanie (2017). “Optimal Taxation and Human Capital Policies over the Life Cycle.” *Journal of Political Economy*, Vol 125(6), 1931-1990.
- [30] Stantcheva, Stefanie (2020). “Dynamic Taxation.” *Annual Review of Economics*, Vol 12, 801-831.
- [31] Thomas, Jonathan, and Tim Worrall (1990). “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem.” *Journal of Economic Theory*, Vol 51(2), 367-390.

9 Appendix: Proofs

Proof of Proposition 1:

The ODE that characterizes $\mu_t(\theta^t)$ satisfies

$$\dot{\mu}_t(\theta^t) + \mu_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} = \left(\lambda_t - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) - \frac{R^{-1}}{U_C(\theta^t)} \right) f(\theta_t | \theta^{t-1})$$

along with boundary conditions $\mu_t(\theta) = \mu_t(\bar{\theta}) = 0$. Define $\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} = \frac{\hat{m}_{\theta_t}(\theta^t)}{\hat{m}(\theta^t)}$, or $\hat{m}(\theta^t) = e^{-\int_{\theta_t}^{\bar{\theta}} \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} d\theta_t}$. Substituting into the above ODE and integrating out yields

$$\mu_t(\bar{\theta}) \hat{m}(\bar{\theta}) - \mu_t(\theta^t) \hat{m}(\theta^t) = \int_{\theta_t}^{\bar{\theta}} \left(\lambda_t - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) - \frac{R^{-1}}{U_C(\theta^t)} \right) f(\theta_t | \theta^{t-1}) \hat{m}(\theta^t) d\theta_t,$$

or

$$\frac{\mu_t(\theta^t)}{\hat{f}(\theta_t | \theta^{t-1})} = \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\frac{R^{-1}}{U_C(\theta', \theta^{t-1})} | \theta' \geq \theta_t, \theta^{t-1} \right) + \eta_t \hat{\mathbb{E}} \left(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta' \geq \theta_t, \theta^{t-1} \right) - \lambda_t \right\}.$$

From the boundary conditions it follows that

$$\int_{\underline{\theta}}^{\bar{\theta}} \left(\lambda_t - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) - \frac{R^{-1}}{U_C(\theta^t)} \right) f(\theta_t | \theta^{t-1}) d\theta_t = 0$$

or $\lambda_t = R^{-1} \hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1}) + \eta_t \hat{\mathbb{E}}(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta^{t-1})$. Rewrite $\hat{\mathbb{E}}(\mathcal{J}(\theta_{t+1}, \theta^t) | \theta^t)$ as

$$\begin{aligned} \hat{\mathbb{E}}(\mathcal{J}(\theta_{t+1}, \theta^t) | \theta^t) &= \int_{\underline{\theta}}^{\bar{\theta}} \mathcal{J}(\theta_{t+1}, \theta^t) \frac{\hat{m}(\theta^{t+1})}{\mathbb{E}(\hat{m}(\theta^{t+1}) | \theta^t)} f(\theta_{t+1} | \theta^t) d\theta_{t+1} \\ &= \int_{\underline{\theta}}^{\bar{\theta}} \int_{\theta_{t+1}}^{\bar{\theta}} \mathcal{J}(\theta', \theta^t) f(\theta' | \theta^t) d\theta' \frac{\hat{m}'(\theta_{t+1}, \theta^t)}{\hat{m}(\theta_{t+1}, \theta^t)} \frac{\hat{m}(\theta^{t+1})}{\mathbb{E}(\hat{m}(\theta^{t+1}) | \theta^t)} f(\theta_{t+1} | \theta^t) d\theta_{t+1} \\ &= \int_{\underline{\theta}}^{\bar{\theta}} \varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} \frac{\hat{m}(\theta^{t+1})}{\mathbb{E}(\hat{m}(\theta^{t+1}) | \theta^t)} f(\theta_{t+1} | \theta^t) d\theta_{t+1} \\ &= \hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} | \theta^t \right) \end{aligned}$$

Substituting into the solution for $\mu_t(\theta^t)$ and rearranging terms yields

$$\frac{\mu_t(\theta^t)}{\hat{f}(\theta_t | \theta^{t-1})} = R^{-1} \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta', \theta^{t-1})} | \theta' \geq \theta_t, \theta^{t-1} \right) - \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^t)} | \theta^{t-1} \right) \right\} + \eta_t \hat{\varrho}_t(\theta^t)$$

where

$$\begin{aligned} \hat{\varrho}_t(\theta^t) &= \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \left\{ \hat{\mathbb{E}}(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta' \geq \theta_t, \theta^{t-1}) - \hat{\mathbb{E}}(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta^{t-1}) \right\} \\ &= \frac{1}{\hat{f}(\theta_t | \theta^{t-1})} \int_{\theta_t}^{\bar{\theta}} \mathcal{J}(\theta', \theta^{t-1}) \frac{\hat{m}'(\theta')}{\mathbb{E}(\hat{m}(\theta') | \theta^{t-1})} f(\theta' | \theta^{t-1}) d\theta' - \hat{\mathbb{E}}(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta^{t-1}) \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \\ &= \frac{1}{\hat{f}(\theta_t | \theta^{t-1})} \int_{\theta_t}^{\bar{\theta}} \mathcal{J}(\theta', \theta^{t-1}) f(\theta' | \theta^{t-1}) d\theta' \frac{\hat{m}(\theta_t)}{\mathbb{E}(\hat{m}(\theta_t) | \theta^{t-1})} - \hat{\mathbb{E}}(\mathcal{J}(\theta_t, \theta^{t-1}) | \theta^{t-1}) \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \\ &\quad + \frac{1}{\hat{f}(\theta_t | \theta^{t-1})} \int_{\theta_t}^{\bar{\theta}} \int_{\theta}^{\bar{\theta}} \mathcal{J}(\theta', \theta^{t-1}) f(\theta' | \theta^{t-1}) d\theta' \frac{\hat{m}'(\theta)}{\mathbb{E}(\hat{m}(\theta) | \theta^{t-1})} d\theta \\ &= \varrho_t(\theta^t) + \frac{1 - \hat{F}(\theta_t | \theta^{t-1})}{\hat{f}(\theta_t | \theta^{t-1})} \left\{ \hat{\mathbb{E}} \left(\varrho_t(\theta, \theta^{t-1}) \frac{U_{\theta C}(\theta, \theta^{t-1})}{U_C(\theta, \theta^{t-1})} | \theta \geq \theta_t, \theta^{t-1} \right) - \hat{\mathbb{E}} \left(\varrho_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} | \theta^{t-1} \right) \right\} \end{aligned}$$

The same steps apply to the ODE

$$\dot{\mu}_t(\theta^t) + \mu_t(\theta^t) \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} = \left(\lambda_t - \eta_t \mathcal{J}(\theta_t, \theta^{t-1}) - \frac{R^{-1}}{U_Y(\theta^t)} \right) f(\theta_t | \theta^{t-1})$$

to arrive at the characterization linking $\mu_t(\theta^t)$ to the inverse marginal utilities of earnings.

Proof of Proposition 2:

For Part (i), notice that

$$\begin{aligned}
\frac{\frac{d}{d\theta^t} U_C(\theta^t)}{U_C(\theta^t)} &= \frac{U_{C\theta}(\theta^t)}{U_C(\theta^t)} + \frac{U_{CC}(\theta^t)}{U_C(\theta^t)} C'(\theta^t) + \frac{U_{CY}(\theta^t)}{U_C(\theta^t)} Y'(\theta^t) \\
&= \frac{U_{C\theta}(\theta^t)}{U_C(\theta^t)} + \left(\frac{U_{CC}(\theta^t)}{U_C(\theta^t)} - \frac{U_{CY}(\theta^t)}{U_Y(\theta^t)} \right) C'(\theta^t) \\
&= \frac{U_{C\theta}(\theta^t)}{U_C(\theta^t)} - \mathcal{E}_C(\theta^t) \frac{C'(\theta^t)}{C(\theta^t)}
\end{aligned}$$

where the second equality substituted the local IC constraint $U_C(\theta^t) C'(\theta^t) = -U_Y(\theta^t) Y'(\theta^t)$, the third equality the definition of $\mathcal{E}_C(\theta^t)$. Integrating out yields

$$\hat{m}(\theta^t) = K U_C(\theta^t) e^{\int_{\underline{\theta}}^{\theta^t} \mathcal{E}_C(\theta', \theta^{t-1}) \frac{C'(\theta', \theta^{t-1})}{C(\theta', \theta^{t-1})} d\theta'}$$

where K is a constant of proportionality that can be set equal to 1 as a normalization (given that the change in probability measure is defined by $\hat{m}(\theta) / \hat{\mathbb{E}}(\hat{m}(\theta))$).

The steps for Part (ii) are identical: combining

$$\begin{aligned}
\frac{\frac{d}{d\theta^t} (-U_Y(\theta^t))}{-U_Y(\theta^t)} &= \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} + \frac{U_{YY}(\theta^t)}{U_Y(\theta^t)} Y'(\theta^t) + \frac{U_{CY}(\theta^t)}{U_Y(\theta^t)} C'(\theta^t) \\
&= \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} + \left(\frac{U_{YY}(\theta^t)}{U_Y(\theta^t)} - \frac{U_{CY}(\theta^t)}{U_C(\theta^t)} \right) Y'(\theta^t) \\
&= \frac{U_{Y\theta}(\theta^t)}{U_Y(\theta^t)} + \mathcal{E}_Y(\theta^t) \frac{Y'(\theta^t)}{Y(\theta^t)}
\end{aligned}$$

from which it follows that $\tilde{m}(\theta^t) = -U_Y(\theta^t) e^{-\int_{\underline{\theta}}^{\theta^t} \mathcal{E}_Y(\theta', \theta^{t-1}) \frac{Y'(\theta', \theta^{t-1})}{Y(\theta', \theta^{t-1})} d\theta'}$.

Proof of Theorem 1:

Rearranging the inter-temporal optimality condition and substituting

$$\frac{R\mu_t(\theta^t)}{f(\theta_t|\theta^{t-1})} = \frac{1}{U_C(\theta)} \frac{\tau_t(\theta^t)}{1 - \tau_t(\theta^t)} \left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right)^{-1}$$

yields

$$\frac{1}{U_C(\theta)} \left(1 + \frac{\tau_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \hat{\mathbb{E}} \left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} | \theta^t \right)}{1 - \tau_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}} \right) = (\beta R)^{-1} \hat{\mathbb{E}} \left(\frac{1}{U_C(\theta^{t+1})} | \theta^t \right)$$

It follows that $\hat{s}_C(\theta^t) \geq 0$ if and only if $\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} \geq \hat{\mathbb{E}}\left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} \mid \theta^t\right)$. Re-arranging terms,

$$\begin{aligned} \hat{s}_C(\theta^t) &= \frac{\tau_t(\theta^t) \frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \hat{\mathbb{E}}\left(\varrho_{t+1}(\theta^{t+1}) \frac{U_{\theta C}(\theta^{t+1})}{U_C(\theta^{t+1})} \mid \theta^t\right)}{1 - \tau_t(\theta^t) \frac{\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}}{\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}}} \\ &= \frac{\tau_t(\theta^t) \frac{1}{U_C(\theta^t)} \left(U_{\theta C}(\theta^t) - \beta R (1 + \hat{s}_C(\theta^t)) \mathbb{E}\left(\varrho_{t+1}(\theta^{t+1}) U_{\theta C}(\theta^{t+1}) \frac{\hat{M}(\theta^{t+1})}{\mathbb{E}(\hat{M}(\theta^{t+1}) \mid \theta^t)} \mid \theta^t\right)\right)}{1 - \tau_t(\theta^t) \frac{\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}}{\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}}} \end{aligned}$$

which implies that $\hat{s}_C(\theta^t) \geq 0$ if and only if

$$U_{\theta C}(\theta^t) \geq \beta R \mathbb{E}\left(\varrho_{t+1}(\theta^{t+1}) U_{\theta C}(\theta^{t+1}) \frac{\hat{M}(\theta^{t+1})}{\mathbb{E}(\hat{M}(\theta^{t+1}) \mid \theta^t)} \mid \theta^t\right).$$

Proof of Proposition 3:

The result follows directly from the characterization of $\hat{\varrho}_t(\theta^t)$ and $\tilde{\varrho}_t(\theta^t)$ in proposition 1.

Proof of Theorem 2:

From the representation of $\widehat{MB}_t(\theta^t)$ and $\widetilde{MB}_t(\theta^t)$ it follows that

$$\frac{\widehat{MB}_t(\theta^t)}{\widetilde{MB}_t(\theta^t)} = \frac{\widetilde{B}(\theta^t)}{(1 - \tau_t(\theta^t)) \widehat{B}(\theta^t)}.$$

It follows from $\widehat{MB}_t(\theta^t) + \hat{\varrho}_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1}) = \widetilde{MB}_t(\theta^t) + \tilde{\varrho}_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1})$ that $\widehat{MB}_t(\theta^t) \geq \widetilde{MB}_t(\theta^t)$ or equivalently $1 - \tau_t(\theta^t) \geq \frac{\widetilde{B}(\theta^t)}{\widehat{B}(\theta^t)}$ if and only if $\hat{\varrho}_t(\theta^t) \leq \tilde{\varrho}_t(\theta^t)$.

Proof of Proposition 4:

(i) Notice that $\widehat{MB}(\cdot) U_C(\cdot)$ can be rewritten as

$$R\widehat{MB}(\theta) U_C(\theta) = \frac{1}{f(\theta)} \left\{ \hat{F}(\theta) \int_{\theta}^{\bar{\theta}} \frac{\hat{M}(\theta')}{\hat{M}(\theta)} f(\theta') d\theta' - (1 - \hat{F}(\theta)) \int_{\underline{\theta}}^{\theta} \frac{\hat{M}(\theta')}{\hat{M}(\theta)} f(\theta') d\theta' \right\}$$

Since $\frac{\hat{M}(\theta')}{\hat{M}(\theta)} = e^{\int_{\theta}^{\theta'} \mathcal{E}_C(\theta'') d \ln C(\theta'')}$, it follows that for a given allocation, $\hat{M}(\theta') / \hat{M}(\theta)$ is increasing in θ and becomes steeper when $\mathcal{E}_C(\cdot)$ goes up, but does not change with $\mathcal{E}_Y(\cdot)$. Therefore, holding fixed $\hat{F}(\theta)$, $\widehat{MB}(\theta) U_C(\theta)$ is increasing with $\mathcal{E}_C(\cdot)$ and bounded below by $R\widehat{MB}(\theta) U_C(\theta) \geq \frac{\hat{F}(\theta) - F(\theta)}{f(\theta)}$.

In addition, $\hat{F}(\theta)$ does not depend on $\mathcal{E}_C(\cdot)$ or $\mathcal{E}_Y(\cdot)$, but an increase in χ'/χ or a reduction in $-U''/U'$ both increase $U_{\theta C}/U_C$, which is independent of a'/a . The results then follow by noting that an increase in $U_{\theta C}/U_C$ results in a FOSD shift in $\hat{F}(\cdot)$ which lowers $\widehat{MB}(\cdot)$.

(ii) Rewrite $\widetilde{MB}(\cdot)(-U_Y(\cdot))$ as

$$R\widetilde{MB}(\theta)(-U_Y(\theta)) = \frac{1}{f(\theta)} \left\{ \widetilde{F}(\theta) \int_{\theta}^{\bar{\theta}} \frac{\widetilde{M}(\theta')}{\widetilde{M}(\theta)} f(\theta') d\theta' - \left(1 - \widetilde{F}(\theta)\right) \int_{\underline{\theta}}^{\theta} \frac{\widetilde{M}(\theta')}{\widetilde{M}(\theta)} f(\theta') d\theta' \right\}$$

Since $\frac{\widetilde{M}(\theta')}{\widetilde{M}(\theta)} = e^{-\int_{\theta}^{\theta'} \mathcal{E}_Y(\theta'') d\ln C(\theta'')}$, it follows that for a given allocation $\widetilde{M}(\theta')/\widetilde{M}(\theta)$ is decreasing in θ and becomes steeper when $\mathcal{E}_Y(\cdot)$ goes up, but does not change with $\mathcal{E}_C(\cdot)$. Therefore, holding fixed $\widetilde{F}(\theta)$, $\widetilde{MB}(\cdot)(-U_Y(\cdot))$ is decreasing with $\mathcal{E}_Y(\cdot)$ and bounded above by $R\widetilde{MB}(\cdot)(-U_Y(\cdot)) \leq \frac{\widetilde{F}(\theta) - F(\theta)}{f(\theta)}$.

In addition, $\widetilde{F}(\theta)$ does not depend on $\mathcal{E}_C(\cdot)$ or $\mathcal{E}_Y(\cdot)$, but an increase in a'/a or a reduction in $-U''/U'$ both increase $U_{\theta Y}/U_Y$, which is independent of χ'/χ . The results then follow by noting that an increase in $U_{\theta Y}/U_Y$ results in a FOSD shift in $\widetilde{F}(\cdot)$ which lowers $\widetilde{MB}(\cdot)$.

(iii) Rewrite $\widehat{MB}(\cdot)U_C(\cdot)$ and $\widetilde{MB}(\cdot)(-U_Y(\cdot))$ as

$$R\widehat{MB}(\cdot)U_C(\cdot) = \frac{\frac{\tau(\theta)}{1-\tau(\theta)}}{U_{\theta C}/U_C - U_{\theta Y}/U_Y} = \frac{R\widetilde{MB}(\theta)(-U_Y(\theta))}{1-\tau(\theta)}$$

and which only depend on $U_{\theta C}/U_C - U_{\theta Y}/U_Y = \chi'/\chi - a'/a$.

Proof of Theorem 3:

Multiply both sides of $MC_t(\theta^t) = \widehat{MB}_t(\theta^t) + \hat{\varrho}_t(\theta^t) \beta R \cdot MC_{t-1}(\theta^{t-1})$ by $\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)}$, divide by $\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})$, and take expectations with respect to the incentive-adjusted probability measure to find

$$\hat{\mathbb{E}} \left(\frac{\tau_t(\theta^t)}{1-\tau_t(\theta^t)} \frac{1/U_C(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) = \hat{\mathbb{E}} \left(\left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right) \frac{\widehat{MB}_t(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) \\ = \frac{\beta R}{U_C(\theta^{t-1})} \frac{1}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} \cdot \hat{\mathcal{R}}(\theta^{t-1}) \frac{\tau_{t-1}(\theta^{t-1})}{1-\tau_{t-1}(\theta^{t-1})}$$

By Theorem 1, $\frac{\beta R}{U_C(\theta^{t-1})} \frac{1}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} = \frac{1}{1+\hat{s}_C(\theta^{t-1})}$. At the same time,

$$\hat{\mathbb{E}} \left(\left(\frac{U_{\theta C}(\theta^t)}{U_C(\theta^t)} - \frac{U_{\theta Y}(\theta^t)}{U_Y(\theta^t)} \right) \frac{\widehat{MB}_t(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) = \hat{\mathbb{E}} \left(\left(\frac{\frac{\partial \hat{m}(\theta^t)}{\partial \theta_t}}{\hat{m}(\theta^t)} - \frac{\frac{\partial \widetilde{m}(\theta^t)}{\partial \theta_t}}{\widetilde{m}(\theta^t)} \right) \frac{\widehat{MB}_t(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right) \\ = \widehat{Cov} \left(\log \left(\frac{\hat{m}(\theta^t)}{\widetilde{m}(\theta^t)} \right), \frac{1/U_C(\theta^t)}{\hat{\mathbb{E}}(1/U_C(\theta^t) | \theta^{t-1})} | \theta^{t-1} \right)$$

Applying the same steps to $MC_t(\theta^t) = \widetilde{MB}_t(\theta^t) + \tilde{\varrho}_t(\theta^t)\beta R \cdot MC_{t-1}(\theta^{t-1})$ leads to the analogous representation based on redistribution through earnings.