

Wedge Dynamics with Evolving Private Information*

Miltiadis Makris
University of Kent

Alessandro Pavan
Northwestern University

April 27, 2020

PRELIMINARY AND INCOMPLETE

Abstract

This paper uses a recursive approach to arrive at a concise formula describing the forces responsible for the dynamics of wedges (i.e., distortions in the second-best allocations relative to their first-best counterparts) in a large class of economies with an arbitrary number of periods, and where the agents' private information evolves over time, possibly in an endogenous manner. The formula accommodates for a flexible specification of the planner's preferences for redistribution (captured by general non-linear Pareto weights on the agents' lifetime utilities) and of the agents' preferences for insurance (captured by the curvature of the agents' payoffs over consumption), as well as for rich specifications of the process governing the evolution of the agents' private information. The value of the formula is twofold. It helps us unify results in the macro new dynamic public finance literature and relate them to results in the micro dynamic mechanism design literature. It also permits us to shed new light on what drives the dynamics of distortions in various economies of interest. For example, we show how the formula can explain why, contrary to what suggested in the literature, distortions may increase over time in economies with constant, or slowly declining, impulse responses of future types to initial ones, even if risk (the variance of the agents' types) remains constant or declines with time. We also show that, when utility is non-transferable, the dynamics of wedges depend on the impulse responses of future types to types in all intermediate periods, as opposed to just the initial ones, as in standard dynamic mechanism design with transferable utility.

*For useful comments and suggestions we thank Laura Doval, Loren Fryxell, Nicola Pavoni, and Stefanie Stantcheva. Pavan also thanks Bocconi University for its hospitality during the 2017-2018 academic year. The usual disclaimer applies. Email addresses: mmakris.econ@gmail.com (Makris); alepavan@northwestern.edu (Pavan).

1 Introduction

The last fifteen years have witnessed great interest in extending mechanism-design techniques from static problems to dynamic ones, where the agents' private information evolves over time and decisions have to be made over multiple periods (see Bergemann and Pavan (2015), Pavan (2017), and Bergemann and Valimaki (2018) for overviews). While most of the micro literature has confined attention to settings with transferable utility (i.e., to economies in which the agents' payoffs are linear in consumption), the fast-growing new Dynamic Public Finance (DPF) and macro literatures have been interested primarily in economies with non-transferable utility (see, for example, Farhi and Werning (2013), Golosov et al. (2016), Stantcheva (2017), Makris and Pavan (2018) and the references therein).

With non-transferable payoffs, distortions in the second-best allocations (relative to their first-best counterparts) are best described in terms of wedges, i.e., discrepancies between marginal rates of substitution and marginal rates of transformation. The formulas for such wedges summarize all the key forces responsible for the distortions in the second-best allocations due to the agents' private information. In static economies, examples of such formulas can be found in Mirrlees' (1971) seminal work on optimal non-linear taxation and in Diamond's (1998) and Saez's (2001) subsequent generalizations of Mirrlees' work. In dynamic economies, examples of such formulas have been documented in various works in the fast-growing new DPF literature (see, for example, Albanesi and Sleet (2006), Golosov et al. (2006), Kocherlakota (2010), Gorry and Oberfield (2012), Kapicka (2013), Farhi and Werning (2013), and Golosov et al. (2016)).

Recent years have also witnessed growing interest in extending the theory of dynamic mechanism design to economies in which the agents' private information evolves endogenously over time, for example because of learning-by-doing, other investments in human capital (see, for example, Krause (2009), Best and Kleven (2013), Kapicka (2006, 2015a,b), Kapicka and Neira (2016), Stantcheva (2015, 2017), and Makris and Pavan (2018)), habit formation (see, for example, Bose and Makris (2016)), or experimentation (see, for example, Fershtman and Pavan (2017)).

This paper provides a general formula for the dynamics of the wedges under second-best allocations that unifies all special cases considered in both the micro and the macro literature. We consider a dynamic environment that accommodates for a flexible specification of (a) the planner's preferences for redistribution (captured by general non-linear Pareto weights on the agents' lifetime utilities), (b) the agents' preferences for insurance (captured by the curvature of the agents' utility function over consumption), and (c) the process governing the endogenous evolution of the agents' private information.

The formula is established through a recursive approach that controls for the endogeneity of the agents' private information. It summarizes all forces that are responsible for the intra- and inter-temporal distortions into four terms. The first term captures all the forces that are active in economies with exogenous private information, transferable utility, and Rawlsian preferences for

redistribution (equivalently, interim participation constraints, as in the micro dynamic mechanism design literature). The second term summarizes all the forces originating from the endogeneity of the agents' private information. These two terms interact linearly (i.e., in an additively separable way) in the wedge formula. The last two terms are correction terms that control for, respectively, the non transferability of the agents' payoffs (equivalently, the heterogeneity in the agents' marginal utility of consumption) and alternative specifications of the planner's preferences for redistribution. Importantly, these last two corrective terms interact multiplicatively with the first two terms, i.e., they operate as amplifiers or dampeners of the first two effects.

Letting W_t^{RRN} denote the wedge under risk neutrality and Rawlsian preferences for redistribution, Ω_t the component of the wedge controlling for the endogeneity of the agents' private information, RA_t the correction term controlling for the non-transferability of the agents' payoffs (equivalently, for the agents' risk aversion), and D the correction term controlling for the principal's preferences for redistribution, we have that the formula for the dynamics of the wedges can be described concisely as

$$W_t = [RA_t - D_t][W_t^{RRN} + \Omega_t].$$

Below, we comment in detail on each of the above four terms and relate them to the primitives of the problem. The formula applies to all economies for which the so-called first-order, Myersonian, approach is valid (i.e., for which the second-best allocations coincide with the solution to a relaxed problem where only local incentive compatibility constraints are imposed).¹

The value of the formula is twofold. It helps us identify, and isolate, the various forces that shape the dynamics of distortions under second-best allocations. It also helps us reconcile and unify various results in both the micro and the macro literature and bring them under a common conceptual umbrella.

The two forces responsible for the dynamics of the wedges (equivalently, of the distortions) identified in the literature are (i) the dynamics of the impulse responses of future types to the initial ones (this force is discussed primarily in the micro literature and is behind the dynamics of the term W_t^{RRN} in the above formula) and (ii) the dynamics of the interaction between the variance of the period- t types and the agents' risk aversion (this force is discussed primarily in the macro/DPF literature and is behind the dynamics of the term $RA_t - D_t$ in the above formula). Our analysis permits us to uncover two new forces. The first one is specific to economies in which the evolution of the agents' private information is endogenous, and is captured by the dynamics of the term Ω_t in the above formula. As we explain below, this term controls for the effects of (variations in) period- t allocations on future rents and plays an important role in economies with learning-by-doing, habit formation, and other sources of endogenous private information. In previous work (Makris and Pavan (2018)), we uncovered the role of this force in a special economy in which the agents' types change only once. In this paper, we show how such force operates in more general economies with an

¹See, among others, Pavan, Segal, and Toikka (2014) and Kapicka (2016) for a description of such an approach.

arbitrary number of periods and of (endogenous) shocks to the agents' private information, and how such force interacts with the other forces identified by the above formula in shaping the dynamics of the wedges.

The second new force identified in the present paper is the accumulation over time of the “transaction” costs to the planner of backward shifting the agents' information rents. Suppose the planner increases an agent's period- t compensation to induce him to reveal his period- t private information, and then adjusts the agent's period-1 compensation by deducting from the period-1 compensation the increase in the expected value of the period- t rent. When the marginal utility of consumption varies across types and time, such a double adjustment in compensation comes with extra costs to the planner that grow with the distance between the time at which the rent is given (period t) and the time at which the rent is partially recouped (period 1). This force is absent in economies in which the agents are risk neutral. As explained later in the paper, this force, which is also behind the dynamics of the term RA_t in our general formula, has not been identified in previous work. It is driven by the interaction of the agents' risk aversion with the dynamics of the impulse responses of the agents' future types to their types at *all* intermediate periods (i.e., the impulse responses of period- t types to each period- s types, with $1 \leq s < t$).

To appreciate the role of this new force, consider an economy with exogenous types in which the agents are risk averse, the impulse responses of the future types to the initial ones are either constant or decline slowly over time (i.e., $\Omega_t = 0$ and W_t^{RRN} constant or declining slowly with t for given effort by the agents), and in which the planner maximizes the sum of the agents' lifetime utilities by assigning equal weights to all types (i.e., has an utilitarian objective). Previous work (notably, Farhi and Werning (2013)) has noticed that, in such economies, wedges tend to increase over the lifecycle and has attributed such dynamics to the fact that the variance of the agents' types increases with time, when evaluated from the perspective of the initial period (which is the case when the agent's type follows a random walk, as in their calibrated economy). Our formula permits us to qualify that the intuition proposed in the literature for why wedges in such economies tend to increase over time is incomplete. This intuition is based on the idea that the planner benefits from shielding the agents from risk. To induce the agents to reveal their private information, the planner must let the agents' compensation vary with the agents' labor supply. The volatility of the agents' compensation, however, can be reduced by distorting their labor downwards. Because risk grows with time, at the optimum, distortions in labor supply then naturally increase over the lifecycle.

The reason why this intuition is incomplete is that, in such economies, wedges may increase over time even when the variance of the agents' types (and hence the risk the agents face) is constant, or even declines with time. As we show later in the paper, the dynamics noticed in the literature are, instead, largely due to the fact that the cost to the planner of recouping the informational rents left to the agents in future periods by deducting such rents from their consumption in earlier periods grows with the distance between the period at which the rent is provided and the one at which it is (partially) recouped (equivalently, with the number of periods the rent is rolled backwards). We

also notice that such effect is absent in economies in which the agents' types change only once (e.g., Garrett and Pavan (2015), and Makris and Pavan (2018)). In such economies, the variance of the agents' types always increases between period 1 (where it is zero) and period 2 (where it is positive), and the only relevant impulse responses are the one between period 2 and period 1 (that is, there are no intermediate impulse responses). As a result, in such simplified economies, the transaction costs from rolling rents backwards cannot be disentangled from the costs of shielding the agents from risk.

While the exposition in the paper favors macro/DPF applications, the results apply more broadly to many dynamic mechanism design problems, including the design of managerial compensation schemes (e.g., Garrett and Pavan (2015)), matching with unknown and time-varying preferences (e.g., Fershtman and Pavan (2017)), and the sale of experiences goods with habit formation (e.g., Bose and Makris (2016)).

Layout. The rest of the paper is organized as follows. Section 2 describes the model. Section 3 provides a characterization of the first-best allocations, that is, the allocations that would be sustained in the absence of the agents' private information. Section 4 develops the recursive approach that yields the second-best allocations, that is, the allocations sustained when the agent possesses private information. Section 5 uses the results in the previous two sections to arrive at the general formula for the wedges mentioned above. It then uses the formula to shed new lights on what drives the dynamics of distortions in various economies considered in the literature.

2 Environment

Hereafter, we refer to the party designing the contractual relationship as the principal (“she”) and to the informed party as the agent (“he”). The relationship between the two parties lasts for T periods, where T is finite. The agent receives new private information in each period. The principal offers a contract that must respect the agent's incentives (i.e., be incentive-compatible) and a certain redistribution constraint described in detail below.

In the new dynamic public finance literature, the agent's type represents the agent's productivity, and the allocations the profile of type-dependent consumption and earnings (alternatively, output). In a canonical buyer-seller model, the agent's type represents a taste parameter and the allocations the profile of type-dependent output and monetary transfers. In a managerial compensation setting, the agent's type may represent his ability to generate cash flows for the firm and an allocation represents a profile of type-dependent effort and compensation, with the latter specifying a transfer from the firm to the manager as a function of performance measures correlated with both the agent's type and the agent's effort.

To fix ideas, hereafter, we will focus on an economy that resembles those studied in the new dynamic public finance literature. It should be easy to see, though, that the results extend to many other dynamic mechanism design problems.

We start with some preliminary notation. Subscripts t denote time, with $t = 1, 2, \dots, T$, where T can either be finite or infinite. Superscripts t , instead, denote histories up to, and including, period t . Thus, for any variable a , $a^t \equiv \{a_1, \dots, a_t\}$, with a_t denoting the period- t value of a . Furthermore, for any t , any $j \geq 0$, $a_t^{t+j} \equiv \{a_t, \dots, a_{t+j}\}$, whereas $a_t^{t-j-1} \equiv \{\emptyset\}$. For any set A , A^0 denotes the empty set. We also use the convention that, when $l < k$, $\prod_{i=k}^l a_{i+1} = 1$ and $\sum_{i=k}^l a_{i+1} = 0$. Finally, we denote by $\mathbb{I}_A(a)$ the indicator function taking value 1 when $a \in A$ and 0 otherwise.

In each period t , the agent produces *output* $y_t \in Y_t = \mathbb{R}_+$ at a cost $\psi(y_t, \theta_t)$, with θ_t denoting the agent's period- t productivity/skill. The latter is the agent's private information and is learned by the agent at the beginning of period t . The function $\psi(y_t, \theta_t)$ is thrice differentiable, increasing, and convex in y_t . We then let

$$\psi_y(y_t, \theta_t) \equiv \partial\psi(y_t, \theta_t)/\partial y_t, \quad \psi_\theta(y_t, \theta_t) \equiv \partial\psi(y_t, \theta_t)/\partial\theta_t, \quad \text{and} \quad \psi_{y\theta}(y_t, \theta_t) \equiv \partial^2\psi(y_t, \theta_t)/\partial\theta_t\partial y_t.$$

We assume that $\psi_{y\theta} < 0$, which implies that higher types are more productive in the sense of having a lower marginal disutility of labor. We also assume that $\psi_{y\theta}$ is nonincreasing in y which guarantees a well-behaved solution to the principal's problem. The agent's productivity at any period $t \geq 2$ is a function of the agent's productivity in the previous period, θ_{t-1} , the agent's output in the previous period, y_{t-1} , and some shock ε_t . That is,

$$\theta_t = z_t(\theta_{t-1}, y_{t-1}, \varepsilon_t),$$

for some function $z_t(\cdot)$ equi-Lipschitz continuous and differentiable, increasing in both the first and third arguments. The dependence of z_t on past output may capture learning-by-doing, as in Makris and Pavan (2018), or other investments in human capital measurable in the agent's past productivity and in past output. While in many applications of interest it is natural to assume that z_t is increasing in y_{t-1} , the analysis below does not hinge on such assumption and accommodates for the possibility that z_t be decreasing in y_{t-1} , or non-monotone in y_{t-1} . A negative dependence of z_t on past output may reflect a substitutability between current productivity and past production, as in models of learning-or-doing, or, in a trade model, a substitutability between current and past consumption, capturing the idea that a buyer may gradually lose interest in a product he consumed intensively in the past.

The shock ε_t is drawn from the interval $E_t \equiv (\underline{\varepsilon}_t, \bar{\varepsilon}_t) \subseteq \mathbb{R}$ according to some cumulative distribution function G_t , absolutely continuous over the entire real line, and with density g_t strictly positive over E_t . Let $F_t(\theta_t | \theta_{t-1}, y_{t-1})$ denote the cumulative distribution function of the period- t productivity θ_t given the period- $(t-1)$ productivity θ_{t-1} and period- $(t-1)$ output, y_{t-1} , as implied by the combination of the distribution G_t with the function z_t . We assume that $F_t(\cdot | \theta_{t-1}, y_{t-1})$ is absolutely-continuous over the entire real line, with density $f_t(\theta_t | \theta_{t-1}, y_{t-1})$ strictly positive over a compact subset of $\Theta_t \equiv (\underline{\theta}_t, \bar{\theta}_t) \subseteq \mathbb{R}$. The set Θ_t defines the support of the marginal distribution of the period- t productivity. Note that the monotonicity of $z_t(\cdot)$ in θ_{t-1} implies that

$$\frac{\partial F_{t+1}(\theta_{t+1} | \theta_t, y_t)}{\partial \theta_t} \leq 0.$$

The above distributions describe the evolution of the agent's private information at all periods $t > 1$. The period-1 productivity, instead, is exogenous and drawn from the interval $\Theta_1 \equiv (\underline{\theta}_1, \bar{\theta}_1) \subseteq \mathbb{R}$ according to a cumulative distribution function F_1 , absolutely-continuous over the entire real line, with density f_1 strictly positive over Θ_1 . For future reference, we also let

$$\eta_1(\theta_1) \equiv \frac{f_1(\theta_1)}{1 - F_1(\theta_1)}$$

denote the hazard rate of the period-1 distribution.

Denote by $c_t \in \mathbb{R}$ the agent's period- t consumption. Let Θ^t denote the set of period- t productivity histories, with generic element $\theta^t \in \Theta^t$, and $\theta \equiv \theta^T$. Interpret $y^t \in Y^t$, $c^t \in \mathbb{R}^t$, $y = y^T$ and $c = c^T$ analogously. Hereafter, we will refer to θ^t as the agent's period- t type, and to θ as the agent's complete type.

The principal's *lifetime* utility is given by

$$U^P(\theta, y, c) \equiv \sum_{t=1}^T \delta^{t-1} (v^P(y_t) - c_t)$$

whereas the agent's *lifetime* utility is given by

$$U^A(\theta, y, c) \equiv \sum_{t=1}^T \delta^{t-1} (v^A(c_t) - \psi(y_t, \theta_t)),$$

with $v^i : \mathbb{R} \rightarrow \mathbb{R}$ increasing, weakly concave, and twice differentiable, $i = A, P$. The function U^i is the Bernoulli utility function player i uses to evaluate lotteries over (θ, y, c) . We denote by

$$U_\tau^A(\theta, y, c) \equiv \sum_{t=\tau}^T \delta^{t-\tau} (v^A(c_t) - \psi(y_t, \theta_t)) \quad \text{and} \quad U_\tau^P(\theta, y, c) \equiv \sum_{t=\tau}^T \delta^{t-\tau} (v^P(y_t) - c_t)$$

the two players' *continuation payoffs* in the restriction of the game that starts with period τ .

We are interested in environments in which output and consumption are strictly positive in each period. For this reason, we assume the following Inada conditions hold: (a) $\lim_{c \rightarrow 0} v^A(c) = \infty$ (when $v^A''(c) < 0$),² and (b) $\lim_{y_t \rightarrow 0} \{v^P(y_t) - \psi_y(y_t, \theta_t)\} > 0$ for all θ_t . To ease the notation, we will also drop the superscript A from the various functions referring to the agent's payoffs, unless there is risk of confusion.

Consistently with the rest of the new dynamic public finance literature, we assume the agent cannot privately save (that is, his savings can be controlled by the principal). Output, y_t , and consumption c_t , are contractible in all periods. The principal can commit to a contract (equivalently, a mechanism) specifying, for each period, consumption c_t and output y_t , as a function of messages sent by the agent in current and past periods. Without loss of optimality, we will restrict attention to (deterministic) direct revelation mechanisms that, in each period, ask the agent to report his new

²When v^A is linear, as in the micro literature, the distribution of consumption over time is indeterminate, in which case assuming consumption is positive in every period is without loss of optimality.

private information, θ_t , and which are *incentive compatible*, meaning that, at each history, it is in the agent’s interest to report truthfully in the continuation game starting with that history, irrespective of whether the agent reported truthfully in the past.³

An incentive-compatible contract can thus be described by the history-dependent allocations it induces. Hereafter, we will refer to the mapping from complete types θ to a pair of consumption and output streams (c, y) as an *allocation rule* and denote the latter by $\chi : \Theta \rightarrow \mathbb{R}^{2T}$. Given χ , we then let $\chi_t(\theta^t) = (y_t(\theta^t), c_t(\theta^t))$ be the period- t allocation under χ , $\chi^t(\theta^t) \equiv (\chi_1(\theta_1), \dots, \chi_t(\theta^t))$ the history of allocations up to, and including, period t , and $\chi(\theta) \equiv \chi^T(\theta^T)$ the complete allocation, when the agent’s complete type is θ . We then denote by $\lambda[\chi]|\theta^t$ the endogenous process over Θ that is obtained by combining the kernels F described above with the allocation rule χ starting from history θ^t , and by $\lambda[\chi]$ the ex-ante distribution over Θ , under the rule χ . Similarly, we will denote by $\lambda[\chi]|\theta^t, y_t$ the endogenous probability distribution over Θ that is obtained by combining the kernels F with the allocation rule χ , starting from history θ^t and period- t output y_t . Given this notation, let

$$V_1(\theta_1) \equiv \mathbb{E}^{\lambda[\chi]|\theta_1} [U^A(\tilde{\theta}, \chi(\tilde{\theta}))]$$

denote type θ_1 ’s expected lifetime utility, under the rule χ (hereafter, we denote by \sim random vectors).

In addition to the aforementioned incentive-compatibility constraints, the principal must guarantee that, for any $\theta_1 \in \Theta_1$, the following constraint holds

$$(1 - r)V_1(\theta_1) + r \int q(V_1(\theta'_1))dF_1(\theta'_1) \geq \kappa, \tag{1}$$

where κ is a constant, $r \in \{0, 1\}$, and the function $q(\cdot)$ is increasing and (weakly) concave and captures possible non-linear Pareto weights assigned by the principal to the agent’s expected lifetime utility. We will refer to the above constraint as the “*redistribution/participation constraint*”. Note that, when $r = 0$, this constraint is equal to the interim participation constraint typically assumed in the micro literature. In a taxation setting, the case $r = 0$, instead, corresponds to the problem of a government with *Rawlsian preferences for redistribution*.⁴ The case $r = 1$ and $q(V_1) = V_1$ for all V_1 , instead, corresponds to an ex-ante participation constraint. Equivalently, in a taxation problem, such a case corresponds to the problem of a planner with “Utilitarian” preferences for redistribution. More generally, a strictly concave function q captures the principal’s inequality aversion (see, for instance, Saez, 2001, Farhi and Werning, 2013, and Best and Kleven, 2013), with stronger concavity corresponding to higher inequality aversion.⁵

³The restriction to mechanism that are incentive compatible at all histories is without loss of optimality in this class of Markov environments; see Pavan, Segal, and Toikka (2014).

⁴This is because incentive compatibility requires $V_1(\theta_1)$ to be non-decreasing. Hence, the relevant period-1 type for which the constraint in (1) binds when $r = 0$ is the lowest one, implying that (1) is equivalent to the familiar constraint that $V_1(\theta_1) \geq 0$ assumed in the taxation literature with a Rawlsian objective.

⁵The results below also apply to a different version of the redistribution constraint in which the weights are for the different period-1 types, i.e., where the constraint takes the form $(1 - r)V_1(\theta_1) + r \int q(\theta'_1)V_1(\theta'_1)dF_1(\theta'_1) \geq \kappa$ with the weighting function $q(\cdot)$ normalized so that $\int q(\theta'_1)dF_1(\theta'_1) = 1$.

We are interested in comparing the second-best allocations for the environment described above to their first-best counterparts.

Definition 1. The rule χ identifies the *first-best allocations* for the environment described above if and only if χ maximizes the principal's ex-ante expected payoff $\mathbb{E}^{\lambda[\chi]}[U^P(\tilde{\theta}, \chi(\tilde{\theta}))]$ over all rules satisfying the participation/redistribution constraint (1). The rule χ identifies the *second-best allocations* for the environment described above if and only if χ maximizes the principal's ex-ante expected payoff $\mathbb{E}^{\lambda[\chi]}[U^P(\tilde{\theta}, \chi(\tilde{\theta}))]$ over all rules that are incentive compatible for the agent and satisfy the participation/redistribution constraint (1).

3 First-Best

In the absence of private information, the principal's optimal allocation rule is obtained by maximizing $\mathbb{E}^{\lambda[\chi]}[U^P(\tilde{\theta}, \chi(\tilde{\theta}))]$ subject to the redistribution constraint (1). This constraint clearly binds at the optimum. The first-best allocation rule can then be described recursively as follows. For any $t < T$, any θ^t , any $y_t(\theta^t)$, let

$$\Pi_{t+1}(\theta^t) \equiv \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t))$$

denote the expected period- $(t+1)$ agent's continuation payoff, given the period- t productivity history θ^t and the period- t output $y_t(\theta^t)$. If $T < +\infty$, then let $\Pi_{T+1}(\theta^T) \equiv 0$. Equipped with this notation, note that the agent's period- t continuation payoff under the rule χ at history θ^t can be written as

$$V_t(\theta^t) \equiv \mathbb{E}^{\lambda[\chi]|\theta^t}[U_t^A(\tilde{\theta}, \chi(\tilde{\theta}))] = v(c_t(\theta^t)) - \psi(y_t(\theta^t), \theta_t) + \delta \Pi_{t+1}(\theta^t).$$

Next, let $C \equiv v^{-1}$. The first-best allocation rule can then be expressed recursively by letting

$$\begin{aligned} Q_t^{FB}(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1})) \equiv \\ \max_{y_t(\theta^{t-1}, \cdot), V_t(\theta^{t-1}, \cdot), \Pi_{t+1}(\theta^{t-1}, \cdot)} \int \{v^P(y_t(\theta^t)) - C(V_t(\theta^t) + \psi(y_t(\theta^t), \theta_t) - \delta \Pi_{t+1}(\theta^t)) \\ + \delta Q_{t+1}^{FB}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t))\} dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \end{aligned}$$

subject to

$$\Pi_t(\theta^{t-1}) = \int V_t(\theta^t) dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})), \text{ for } t > 1 \quad (2)$$

and

$$\kappa = (1-r)V_1(\theta_1) + r \int q(V_1(\theta'_1)) dF_1(\theta'_1), \text{ all } \theta_1 \in \Theta_1, \quad (3)$$

with

$$\Pi_{T+1}(\theta) = 0, \text{ for all } \theta \in \Theta, \text{ if } T \text{ is finite.}$$

Note that $Q_\tau^{FB}(\theta^{\tau-1}, y_{\tau-1}(\theta^{\tau-1}), \Pi_\tau(\theta^{\tau-1}))$ is the period- τ value function of the principal's problem, given the period- τ "state variables" $(\theta^{\tau-1}, y_{\tau-1}(\theta^{\tau-1}), \Pi_\tau(\theta^{\tau-1}))$. To understand the formalization,

note that, in each period $t = 1, \dots, T$, given the state $(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}))$, the choice of the period- t output schedule, $y_t(\theta^{t-1}, \cdot)$, along with the choice of the agent's period- t continuation payoff, $V_t(\theta^{t-1}, \cdot)$, and the agent's expected future "promised utility" $\Pi_{t+1}(\theta^{t-1}, \cdot)$, determine the agent's period- t consumption schedule $c_t(\theta^{t-1}, \cdot)$. For all $\theta^t = (\theta^{t-1}, \theta_t)$, the latter is simply given by

$$c_t(\theta^t) = C(V_t(\theta^t) + \psi(y_t(\theta^t), \theta_t) - \delta \Pi_{t+1}(\theta^t)).$$

The first-best allocation, expressed in recursive form, is then obtained by choosing policies $y_t(\theta^{t-1}, \cdot)$, $V_t(\theta^{t-1}, \cdot)$, and $\Pi_{t+1}(\theta^{t-1}, \cdot)$, for each period t , that jointly maximize the principal's expected continuation payoff subject to the consistency (or, equivalently, promise keeping) constraint (2) that the average of the agent's period- t continuation utility be equal to the level promised in the previous period. Importantly, the expectation over θ_t is computed given the period- $(t-1)$ productivity, θ_{t-1} , and the period- $(t-1)$ output, $y_{t-1}(\theta^{t-1})$. In period one, the promise-keeping constraint is replaced by the (binding) redistribution constraint (3).

Next, for any rule χ and any truthful history $(\theta^t, \theta^{t-1}, \chi^{t-1}(\theta^{t-1}))$, define the principal's period- t continuation payoff under the rule χ by

$$V_t^P(\theta^t) \equiv \mathbb{E}^{\lambda[\chi]|\theta^t} [U_t^P(\tilde{\theta}, \chi(\tilde{\theta}))].$$

Finally, let

$$LD_t^{FB;\chi}(\theta^t) \equiv \delta \frac{\partial}{\partial y_t} \int \left\{ V_{t+1}^P(\theta^{t+1}) + \frac{V_{t+1}(\theta^{t+1})}{v'(c_t(\theta^t))} \right\} dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \quad (4)$$

denote the effect of a marginal variation in the period- t output on the expected sum of the principal's and of the agent's period- $(t+1)$ continuation payoffs, with the latter weighted by the inverse marginal utility of the period- t consumption $v'(c_t(\theta^t))$ — if T is finite, let $V_{T+1}(\theta^{T+1}) \equiv V_{T+1}^P(\theta^{T+1}) \equiv 0$. Importantly, note that the marginal variation captured by the function $LD_t^{FB;\chi}(\theta^t)$ is computed holding fixed the mapping from future productivity histories into allocations, as specified by the rule χ . The following result summarizes all the key properties of the first-best allocations:

Proposition 1. *Suppose the rule $\chi = (y, c)$ identifies the first-best allocations. Then the following conditions hold at all interior points with $\lambda[\chi]$ -probability one:*

$$v^P'(y_t(\theta^t)) + LD_t^{FB;\chi}(\theta^t) = \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} \quad \text{all } t = 1, \dots, T, \quad (5)$$

$$v'(c_t(\theta^t)) = v'(c_{t+1}(\theta^t, \theta_{t+1})), \quad \text{any } t = 1, \dots, T-1 \quad (6)$$

and

$$rq'(V_1(\theta_1))v'(c_1(\theta_1)) = rq'(V_1(\theta'_1))v'(c_1(\theta'_1)).$$

The first condition describes the first-best output schedule. At each period, the principal equalizes the marginal benefit of asking the agent for an extra unit of output (taking into account its effect

on future payoffs stemming from the endogeneity of the process) with its marginal cost. The latter in turn reflects the fact that the principal needs to increase the agent’s utility of consumption by an amount equal to the agent’s marginal disutility of higher output. The monetary cost of compensating the agent for the extra disutility is obtained by dividing the marginal disutility of output by the marginal utility of consumption. Naturally, a high degree of risk aversion (equivalently, a fast declining v') increases the cost to the principal of compensating the agent for the extra disutility of output. When productivity evolves exogenously (equivalently, in the last period T , if the latter is finite), condition (5) then reduces to the familiar optimality condition $v^{P'}(y_t(\theta^t)) = \psi_y(y_t(\theta^t), \theta_t)/v'(c_t(\theta^t))$. To appreciate the effects due to the endogeneity of the process, consider the case where $F_{t+1}(\theta_{t+1}|\theta_t, y_t)$ is non-increasing in y_t , meaning that higher period- t output shifts the distribution of θ_{t+1} in a first-order-stochastic-dominance way. Further assume that the sum of the principal’s and of the agent’s future continuation payoffs (adjusted by the agent’s marginal utility of consumption) is increasing in θ_{t+1} . Then, $LD_t^{FB;x}(\theta^t) \geq 0$. In this case, the endogeneity of the agent’s productivity thus induces the principal to ask for a higher output in period t compared to the level she would ask if productivity was exogenous. This is because higher output at present implies higher productivity in future periods (albeit, in a stochastic sense), which in turn brings higher discounted expected net surplus.

The second and third conditions in turn describe the optimal choice of consumption. When the agent is risk neutral (meaning his utility is linear in consumption) and the principal has an “utilitarian” objective, i.e., $r = 1$ and $q(V_1) = V_1$, all V_1 , the dynamics of consumption is indeterminate. The reason is that the agent does not have preferences for consumption smoothing. In the absence of any inequality aversion on the principal’s side, the distribution of continuation utility across types is then indeterminate. When, instead, the agent is risk averse, optimality requires the equalization of the agent’s marginal utility of consumption over any two consecutive type histories θ^t and (θ^t, θ_{t+1}) . Furthermore, away from the Rawlsian case (i.e., when $r = 0$), optimality also requires the equalization of the “marginal weights” $q'(V_1(\theta_1))v'(c_1(\theta_1))$ the principal assigns to the agent’s period-1 marginal consumption. In the Rawlsian case, instead, the principal equalizes the expected lifetime utility of any period-1 type to the participation threshold κ .

4 The Second Best

We now turn to the case where θ is the agent’s private information.

4.1 Incentive Compatibility

Let $I_t^\tau(\theta^\tau, y^{\tau-1})$ denote the period- τ impulse response of θ_τ to θ_t , as defined in Pavan, Segal, and Toikka (2014). The impulse response incorporates all the ways (direct and indirect) through which a marginal change in θ_t affects θ_τ , $\tau \geq t$, fixing the shocks ε^τ that, along with the decisions $y^{\tau-1}$,

are responsible for the type history θ^τ . Specifically, for all $t \geq 1$, all (θ^t, y^{t-1}) , $I_t^t(\theta^t, y^{t-1}) = 1$, and

$$I_t^{t+1}(\theta^{t+1}, y^t) = \frac{\partial z_{t+1}(\theta_t, y_t, \epsilon_{t+1})}{\partial \theta_t} \Big|_{\epsilon_{t+1} = e_{t+1}(\theta_t, \theta_{t+1}, y_t)} = \frac{\frac{\partial}{\partial \theta_t} [1 - F_{t+1}(\theta_{t+1} | \theta_t, y_t)]}{f_{t+1}(\theta_{t+1} | \theta_t, y_t)}$$

where $e_{t+1}(\theta_{t+1}, \theta_t, y_t)$ is defined implicitly by

$$z_{t+1}(\theta_t, y_t, e_{t+1}(\theta_{t+1}, \theta_t, y_t)) = \theta_{t+1}.$$

Using the above definition, the impulse response function over non-consecutive periods can then be defined inductively, for any $\tau > t$, any θ^τ , any $y^{\tau-1}$, as follows:

$$I_t^\tau(\theta^\tau, y^{\tau-1}) = \prod_{i=0}^{\tau-t-1} I_{t+i}^{t+1+i}(\theta^{t+1+i}, y^{t+i}). \quad (7)$$

For future reference, also note the following two key properties of these impulse response functions:

(a) for any $\tau \geq t$, θ^t , and χ ,

$$\mathbb{E}^{\lambda|\chi| \theta^t} [I_t^\tau(\tilde{\theta}^\tau, y^{\tau-1}(\tilde{\theta}^{\tau-1}))] = \frac{\partial}{\partial \theta_t} \mathbb{E}^{\lambda|\chi| \theta^t} [\tilde{\theta}_\tau | \theta^t, y^{t-1}(\theta^{t-1})]$$

and (b) for any differentiable and equi-Lipschitz continuous function $H(\theta_{t+1})$, any (θ_t, y^t) ,

$$\frac{\partial}{\partial \theta_t} \int H(\theta_{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t) = \int \frac{\partial H(\theta_{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y^t) dF_{t+1}(\theta_{t+1} | \theta_t, y_t). \quad (8)$$

The left-hand-side of (8) is the marginal variation of the expectation of the function $H(\theta_{t+1})$ due to a variation in the distribution of θ_{t+1} triggered by a marginal variation in θ_t . The right-hand side of (8) is simply the expectation of the product of the derivative of the H function with the impulse response of θ_{t+1} to θ_t , holding (θ_t, y_t) fixed.

Recall that $V_t(\theta^t)$ denotes the agent's continuation payoff at the truthful history $(\theta^t, \theta^{t-1}, \chi^{t-1}(\theta^{t-1}))$, under the rule χ . This is the payoff that the agent expects from period t onwards under a truthful strategy; note that, because χ is deterministic, truthful histories can be described entirely in terms of the realized type history θ^t . Therefore, hereafter, whenever there is no risk of confusion, we will be referring to a truthful history under the rule χ simply by θ^t . Then, let

$$D_t^X(\theta^{t-1}, \theta_t) = -\mathbb{E}^{\lambda|\chi| \theta^t} \left[\sum_{\tau=t}^T \delta^{\tau-t} I_t^\tau(\tilde{\theta}^\tau, y^{\tau-1}(\tilde{\theta}^{\tau-1})) \psi_\theta(y_\tau(\tilde{\theta}^\tau), \tilde{\theta}_\tau) \right]$$

denote the net present value of all future marginal variations in the disutility of labor, due to a marginal variation in the period- t productivity. Define $D_t^{\chi \circ \hat{\theta}_t}(\theta^{t-1}, \theta_t)$ in an analogous way for the allocation rule $\chi \circ \hat{\theta}_t$ that is obtained from χ by mapping any period- t message into the message $\hat{\theta}_t$ and then determining allocations according to χ as if the period- t message was $\hat{\theta}_t$. Theorems 1 and 3 in Pavan, Segal, and Toikka (2014) imply that the allocation rule χ is incentive compatible *if and only if*, for all t , all $(\theta^{t-1}, \theta_t), (\theta^{t-1}, \hat{\theta}_t) \in \Theta^t$, (a) the agent's equilibrium continuation payoff $V_t(\theta^{t-1}, \cdot)$ under χ is Lipschitz continuous over Θ_t with derivative given for almost all $\theta_t \in \Theta_t$ by

$$\frac{\partial V_t(\theta^t)}{\partial \theta_t} = D_t^X(\theta^{t-1}, \theta_t), \quad (9)$$

and (b)

$$\int_{\hat{\theta}_t}^{\theta_t} \left[D_t^x(\theta^{t-1}, r) - D_t^{x \circ \hat{\theta}_t}(\theta^{t-1}, r) \right] dr \geq 0. \quad (10)$$

Clearly, the marginal variation of V_t due to the marginal change in θ_t is related to the information rent that the principal must leave to the agent in the continuation game that starts with period t to induce truthful reporting. Note that the assumption that each z_s is increasing in θ_{s-1} implies that $I_t^r > 0$, and hence that $V_t(\theta^{t-1}, \theta_t)$ is non-decreasing in θ_t . Also note that Condition (9) is an envelope condition that relates the marginal variation in the net present value of the agent's payments to the marginal variation in the net present value of the non-monetary allocations, y_t^T . Condition (10), in turn, is a dynamic monotonicity condition requiring that the marginal variation in the net present value of the non-monetary allocations be sufficiently monotone in the period- t report.

4.2 Second-Best Allocations

As anticipated in the Introduction, our goal is to arrive at a general formula describing the evolution of the distortions in the second-best allocations for all such dynamic problems in which the First-Order/Myersonian Approach is valid. Such approach considers a *relaxed program* in which the various integral-monotonicity conditions of (10) are dropped and checked ex-post. Hereafter, we thus drop the integral monotonicity conditions of (10) and show that, when such conditions do not bind, the second-best allocations can be derived by expressing the principal's problem in a convenient recursive form that accounts for the endogeneity of the agent's private information.

As in the previous subsection, write the agent's (on-path) continuation payoff

$$V_t(\theta^t) \equiv v(c_t(\theta^t)) - \psi(y_t(\theta^t), \theta_t) + \delta \Pi_{t+1}(\theta^t)$$

in the continuation game that starts with the period- t history θ^t as the sum of the period- t flow payoff $v(c_t(\theta^t)) - \psi(y_t(\theta^t), \theta_t)$ and the discounted expected continuation payoff $\delta \Pi_{t+1}(\theta^t)$, where, for any $t < T$,

$$\Pi_{t+1}(\theta^t) \equiv \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)),$$

whereas for $t = T$ (when the latter is finite) $\Pi_{T+1}(\theta) \equiv 0$, all θ .

Next, for any $t < T$, let

$$Z_{t+1}(\theta^t) \equiv -\mathbb{E}^{\lambda^{|\chi|} | \theta^t} \left[\sum_{\tau=t+1}^T \delta^{\tau-t-1} I_t^r(\tilde{\theta}^\tau, y^{\tau-1}(\tilde{\theta}^{\tau-1})) \psi_\theta(y_\tau(\tilde{\theta}^\tau), \tilde{\theta}_\tau) \right] \quad (11)$$

with $Z_{T+1}(\theta) \equiv 0$ if $T < +\infty$. Using this notation, the various local incentive-compatibility constraints corresponding to Condition (9) above can also be conveniently written in recursive form as

$$\frac{\partial V_t(\theta^t)}{\partial \theta_t} = -\psi_\theta(y_t(\theta^t), \theta_t) + \delta Z_{t+1}(\theta^t).$$

Furthermore, using Condition (9) for period $t + 1$, along with the law of iterated expectations, and the law of motion of the impulse responses described in (7), we have that

$$Z_{t+1}(\theta^t) = \int [-\psi_\theta(y_{t+1}(\theta^{t+1}), \theta_{t+1}) + \delta Z_{t+2}(\theta^{t+1})] I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)).$$

The principal's (relaxed program) can then be conveniently rewritten in recursive form as follows

$$Q_t(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}), Z_t(\theta^{t-1})) \equiv \max_{y_t(\theta^{t-1}, \cdot), V_t(\theta^{t-1}, \cdot), \Pi_{t+1}(\theta^{t-1}, \cdot), Z_{t+1}(\theta^{t-1}, \cdot)}$$

$$\int \hat{Q}_t((\theta^{t-1}, \theta_t), y_t(\theta^{t-1}, \theta_t), V_t(\theta^{t-1}, \theta_t), \Pi_{t+1}(\theta^{t-1}, \theta_t), Z_{t+1}(\theta^{t-1}, \theta_t)) dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))$$

subject to⁶

$$\Pi_{T+1}(\theta) = Z_{T+1}(\theta) = 0, \text{ all } \theta \in \Theta, \quad (12)$$

$$\kappa = (1 - r)V_1(\underline{\theta}_1) + r \int q(V_1(\theta'_1)) dF_1(\theta'_1), \quad (13)$$

$$\frac{\partial V_t(\theta^{t-1}, \theta_t)}{\partial \theta_t} = -\psi_\theta(y_t(\theta^t), \theta_t) + \delta Z_{t+1}(\theta^t) \text{ all } t, \text{ all } \theta^t \in \Theta^t, \quad (14)$$

$$\Pi_t(\theta^{t-1}) = \int V_t(\theta^t) dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \text{ all } t > 1, \text{ all } \theta^{t-1} \in \Theta^{t-1}, \quad (15)$$

and

$$Z_t(\theta^{t-1}) = \int [-\psi_\theta(y_t(\theta^t), \theta_t) + \delta Z_{t+1}(\theta^t)] I_{t-1}^t((\theta_{t-1}, \theta_t), y_{t-1}(\theta^{t-1})) dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \\ \text{all } t > 1, \text{ all } \theta^{t-1} \in \Theta^{t-1}, \quad (16)$$

where

$$\hat{Q}_t(\theta^t, y_t(\theta^t), V_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t)) \equiv \\ v^P(y_t(\theta^t)) - C(V_t(\theta^t) + \psi(y_t(\theta^t), \theta_t) - \delta \Pi_{t+1}(\theta^t)) \\ + \delta Q_{t+1}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t)).$$

In words, the principal's problem consists in choosing, for each period $t \geq 1$, history of past types θ^{t-1} , period- $(t - 1)$ output $y_{t-1}(\theta^{t-1})$, promised expected utility $\Pi_t(\theta^{t-1})$, and promised marginal continuation utility $Z_t(\theta^{t-1})$, a period- t output schedule $y_t(\theta^{t-1}, \cdot)$, a period- t continuation utility $V_t(\theta^{t-1}, \cdot)$ (including the agent's period- t flow payoff), a period- t promised expected future utility $\Pi_{t+1}(\theta^{t-1}, \cdot)$, and a period- t promised expected future marginal utility $Z_{t+1}(\theta^{t-1}, \cdot)$. All these period- t functions must be selected jointly to maximize the combination of the principal's flow and future

⁶Note that, in writing the redistribution/participation constraint $(1 - r)V_1(\theta_1) + r \int q(V_1(\theta'_1)) dF_1(\theta'_1) \geq \kappa$, all θ_1 , we used the fact that, at the optimum, such constraint necessarily binds, along with the fact that, when $r = 0$, $V_1(\theta_1) \geq \kappa$ all θ_1 if and only if $V_1(\underline{\theta}_1) \geq \kappa$.

expected payoff, taking into account that the principal will face a similar optimization problem in future periods, as is always the case in dynamic programming. In addition, the period- t schedules must respect the promise-keeping constraints that the average period- t continuation utility and marginal continuation utility be equal to what was promised in the previous periods, i.e., $\Pi_t(\theta^{t-1})$ and $Z_t(\theta^{t-1})$, respectively. Clearly, such promise-keeping constraints apply only to period $t > 1$. In period $t = 1$, instead, the principal is only constrained by the redistribution/participation constraint (13), which plays the role of an initial condition for the dynamics of the various policies.⁷

Note two key differences with respect to the recursive problem identifying the first-best optimal policies. The first one is the presence of the various constraints (14) on the derivatives of the agent's continuation payoff. As explained above, such constraints are necessary conditions for incentive compatibility. The second is the presence of the promise-keeping constraints (16) requiring that the agent's expected marginal continuation utility

$$\frac{\partial}{\partial \theta_{t-1}} \mathbb{E}[V_t(\theta^{t-1}, \tilde{\theta}_t) | \theta_{t-1}, y_{t-1}(\theta^{t-1})]$$

from period t (included) onwards with respect to the period- $(t-1)$ type θ_{t-1} be equal to what was promised in the previous period, $Z_t(\theta^{t-1})$.⁸ Note that, in writing (16), we used property (8) of the impulse response functions, along with Condition (14) to write

$$\begin{aligned} \frac{\partial}{\partial \theta_{t-1}} \mathbb{E}[V_t(\theta^{t-1}, \tilde{\theta}_t) | \theta_{t-1}, y_{t-1}(\theta^{t-1})] &= \mathbb{E}\left[\frac{\partial V(\theta^{t-1}, \tilde{\theta}_t)}{\partial \theta_t} I_{t-1}^t(\tilde{\theta}^t, y_{t-1}(\theta^{t-1})) | \theta^{t-1}, y_{t-1}(\theta^{t-1})\right] \\ &= \int [-\psi_\theta(y_t(\theta^t), \theta_t) + \delta Z_{t+1}(\theta^t)] I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1})) dF_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})). \end{aligned} \quad (17)$$

The above problem can thus be seen as a collection of interdependent optimal control problems. To state our next proposition, which summarizes the solution to the above problem, we need to introduce some further definitions. First, for any $t \geq 1$, any $\theta^t \in \Theta^t$, let $\lambda_t(\theta^t)$ denote the shadow cost of increasing the agent's continuation utility at history θ^t .⁹ Next, let

$$LD_t^\chi(\theta^t) \equiv \delta \frac{\partial}{\partial y_t} Q_{t+1}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))$$

⁷The above principal problem can be read in two ways: (a) it describes optimality conditions that the policy χ must satisfy on-path, i.e., at histories $(\theta^{t-1}, \chi^{t-1}(\theta^{t-1}))$ that are consistent with the implementation of the decisions specified by the policy χ for periods $s \leq t-1$; (b) it also describes the recursion that can be used to reduce the dimensionality of the dynamic programming problem —note that the solution to the problem depends on the “state” $(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}), Z_t(\theta^{t-1}))$ only through the four numbers $(\theta_{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}), Z_t(\theta^{t-1}))$.

⁸Note that the derivative $\frac{\partial}{\partial \theta_{t-1}} \mathbb{E}[V_t(\theta^{t-1}, \tilde{\theta}_t) | \theta_{t-1}, y_{t-1}(\theta^{t-1})]$ is with respect to the agent's true period- $(t-1)$ type θ_{t-1} , holding fixed the history of output and consumption decisions $y^{t-1}(\theta^{t-1})$ and $c^{t-1}(\theta^{t-1})$.

⁹Formally, as we show in the Appendix, $\lambda_t(\theta^t)$ coincides with the multiplier $\xi_{t+1}(\theta^t)$ associated with the promised marginal utility constraint (16) at period $t+1$. To see why this is the case, note that, by virtue of (17), (16) is a constraint on the extra utility $\frac{\partial}{\partial \theta_t} \mathbb{E}[V_{t+1}(\tilde{\theta}^{t+1}) | \theta^t, y_t(\theta^t)]$ the principal must leave, for incentives reasons, to all period- t types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t \geq \theta_t$, when she increases the continuation utility $\mathbb{E}[V_{t+1}(\tilde{\theta}^{t+1}) | \theta^t, y_t(\theta^t)]$ of type $\theta^t = (\theta^{t-1}, \theta_t)$. The multiplier $\xi_{t+1}(\theta^t)$ to such constraint thus captures the shadow cost of increasing type θ^t 's continuation payoff.

denote the marginal effect of higher period- t output on the future value of the above optimization problem (equivalently, on the principal's continuation payoff, taking into account all future constraints), when the values of Π_{t+1} , Z_{t+1} and y_t at history θ^t are the ones determined by the output and consumption schedules specified by the policy χ . We then have the following result:

Proposition 2. *Suppose the rule $\chi = (y, c)$ identifies the second-best allocations. The following optimality conditions must then hold at all interior points with $\lambda[\chi]$ -probability one:*

$$\frac{1}{v'(c_t(\theta_t))} = \int \frac{1}{v'(c_{t+1}(\theta^t, \theta_{t+1}))} dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^{t-1})), \text{ for any } t < T, \quad (18)$$

and

$$v^{P'}(y_t(\theta^t)) + LD_t^\chi(\theta^t) = \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} - \psi_{y\theta}(y_t(\theta^t), \theta_t) \cdot \lambda_t(\theta^t), \text{ for all } t, \quad (19)$$

where

$$\begin{aligned} \lambda_t(\theta^t) = & \frac{1-F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} \frac{dF_t(\theta'_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{1-F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \\ & - \frac{1-F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \\ & + I_{t-1}^t(\theta^t, y^{t-1}(\theta^{t-1})) \lambda_{t-1}(\theta^{t-1}), \text{ for all } t > 1, \end{aligned} \quad (20)$$

and

$$\lambda_1(\theta_1) = \frac{\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1)}{f_1(\theta_1)} - r \frac{\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1)}{f_1(\theta_1)} \pi, \quad (21)$$

with

$$\pi \equiv \frac{\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1)}{\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1)}.$$

Together, the conditions in Proposition 2 provide a complete characterization of the second-best allocations. Condition (18) is the familiar Rogerson-inverse-Euler condition. It appears in various works in the new dynamic public finance literature (see, among others, Albanesi and Sleet (2006), Kapicka (2013), Farhi and Werning (2013), and Golosov et al. (2016)). It also appears in various papers in the managerial compensation literature (see, e.g., Garrett and Pavan (2015), and the references therein). Such a condition describes the optimal intertemporal allocation of consumption, for a given output policy. This condition is absent in the micro literature, for, in this literature, the agent's payoff is linear in consumption, in which case the intertemporal distribution of payments/consumption is indeterminate.

Condition (19), in turn, describes the optimal output schedule. The left-hand side is the marginal benefit to the principal of increasing output at history θ^t , taking into account the effect that higher period- t output has on the principal's continuation surplus, and accounting for the cost of future incentives, as captured by the term $LD_t^\chi(\theta^t)$. The right-hand side is the marginal cost to the principal. It combines the direct monetary cost of compensating the agent for the marginal increase

in the disutility of generating higher output, as captured by the term $\psi_y(y_t(\theta^t), \theta_t)/v'(c_t(\theta^t))$, with the cost of raising the information rent (equivalently, the continuation payoff) of all agents whose period- t productivity exceeds θ_t . This cost combines the familiar term $-\psi_{y\theta}(y_t(\theta^t), \theta_t)$ that captures the disutility of effort that can be saved by those agents with higher period- t productivity when mimicking type θ_t , with the shadow cost $\lambda_t(\theta^t)$ of increasing the period- t continuation payoff of all period- t types whose period- t productivity exceeds θ_t (more on this shadow cost below). As usual, both the marginal benefit and the marginal cost are per worker of period- t productivity θ_t . As a result, the shadow cost is obtained by dividing the aforementioned costs by the density $f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))$.

The last condition in the proposition, Condition (20), characterizes the dynamics of such shadow cost, with initial condition given by (21). To understand the initial condition (21), note that, for incentive reasons, when the principal increases the expected lifetime utility of type θ_1 by one unit, she also needs to increase the expected lifetime utility of all higher period-1 types by the same amount.¹⁰ The first term in (21) is the direct cost to the principal, in consumption terms, of providing such extra utility, taking into account the heterogeneity in the marginal utility of consumption of all types above θ_1 . This is the only relevant term in the micro literature, where the redistribution/participation constraint takes the familiar Rawlsian form $v(\theta_1) \geq \kappa$ all θ_1 , which is formally equivalent to $r = 0$ in the redistribution/participation constraint (13)). When, instead, the principal's preferences for redistribution are less extreme than in the Rawlsian case (formally, when $r = 1$ in (13), in which case the redistribution/participation constraint takes the form $\int_{\underline{\theta}_1}^{\bar{\theta}_1} q(V(\theta'_1))dF_1(\theta'_1) \geq \kappa$, for some non-increasing function $q(\cdot)$ describing the non-linear Pareto weights used by the principal to evaluate the various types' lifetime expected utilities), increasing the lifetime expected utility of all agents whose period-1 productivity exceeds θ_1 comes with the benefit of relaxing the redistribution/participation constraint (13). This benefit is captured by the second term in (21). In particular, the term

$$\int_{\theta_1}^{\bar{\theta}_1} q'(V(\theta'_1))dF_1(\theta'_1)$$

is the marginal value of increasing by one util the expected lifetime utility of all period-1 types above θ_1 . The term π , on the other hand, is the shadow value of relaxing the redistribution/participation constraint (13). To see this, note that increasing the value of the left-hand side of the redistribution constraint (13) by one unit, while ensuring incentive compatibility, can be achieved by increasing the lifetime utility of each period-1 type by an amount equal to $1/\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta'_1))dF_1(\theta'_1)$. Aggregating across agents, while accounting for the heterogeneity in the different types' marginal utility of consumption, we have that the total effect of such relaxation on the resources that the principal can appropriate (in consumption units) is equal to π . As mentioned above, such benefit is absent in the Rawlsian case (i.e., when $r = 0$ in (13)), for, in this case, the principal does not value increasing the expected lifetime utility of any period-1 type above $\underline{\theta}_1$.

¹⁰To see this, use (14).

Next, consider the law of motion (20) describing the dynamics of the shadow costs λ_t of increasing the agent's continuation payoffs at any period after the first one. The first term in the right-hand side of (20) is the direct marginal cost, in consumption terms, of increasing the period- t expected continuation utility of all agents whose period- t productivity exceeds θ_t (fixing the period- $(t-1)$ history θ^{t-1}). As usual, such cost is computed taking into account the heterogeneity in the types' marginal utility of consumption. To understand the second term in the right-hand side of (20), note that, when the principal increases by one util the continuation payoff of all period- t types $(\theta^{t-1}, \tilde{\theta}_t)$ with $\tilde{\theta}_t > \theta_t$, then all agents at history θ^{t-1} expect an increase in their lifetime utility by $\delta[1 - F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))]$. The principal can then reduce the compensation she provides in period $t-1$ to such agents by

$$\frac{\delta [1 - F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{v'(c_{t-1}(\theta^{t-1}))}$$

while maintaining unchanged the continuation utility of all such agents. Next, use the Rogerson-inverse-Euler Condition (18) to observe that

$$\frac{\delta [1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{v'(c_{t-1}(\theta^{t-1}))} = \delta \int_{\underline{\theta}_t}^{\tilde{\theta}_t} \frac{1 - F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})),$$

which implies that, in terms of period- t consumption, the amount the principal can recoup (normalized by the density $f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))$) by reducing the compensation paid to the agent at history θ^{t-1} is equal to the second term in the the right-hand side of (20).

The last term in the right-hand side of (20), for $t > 1$, is the extra cost to the principal of increasing by one util the continuation utility of all period- $(t-1)$ types $(\theta^{t-2}, \theta'_{t-1})$ with $\theta'_{t-1} > \theta_{t-1}$. To see this, observe that, when the principal increases the continuation utility of all period- t types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t > \theta_t$, by one util, to preserve incentives at period $t-1$, she then needs to increase the continuation utility of all period- $(t-1)$ types $(\theta^{t-2}, \theta'_{t-1})$, with $\theta'_{t-1} > \theta_{t-1}$, by

$$\delta \frac{\partial [1 - F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{\partial \theta_{t-1}}.$$

This is true despite the adjustment in the period- $(t-1)$ compensation of type $\theta^{t-1} = (\theta^{t-2}, \theta_{t-1})$ described above.¹¹ Such increase is necessary to discourage these types from mimicking type $(\theta^{t-2}, \theta_{t-1})$ and originates from the fact that such types attach a higher probability to having a period- t productivity above θ_t than type $(\theta^{t-2}, \theta_{t-1})$ and hence enjoy the extra rent promised to the period- t types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t > \theta_t$. Using the fact that

$$\frac{\frac{\partial [1 - F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{\partial \theta_{t-1}}}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} = I_{t-1}^t(\theta^t, y^{t-1}(\theta^{t-1})),$$

¹¹Importantly, note that if the principal did not reduce the period- $(t-1)$ consumption of type θ^{t-1} so as to hold the latter type's continuation utility constant, she would then have to increase the continuation utility of all period- $(t-1)$ types $(\theta^{t-2}, \theta'_{t-1})$ with $\theta'_{t-1} > \theta_{t-1}$ by

$$\delta \partial [1 - F_t(\theta_t|\theta^{t-1}, y_{t-1}(\theta^{t-1}))]/\partial \theta_{t-1} + \delta [1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))].$$

we then have that the shadow cost of increasing the continuation utility of all period- $(t - 1)$ types $(\theta^{t-2}, \theta'_{t-1})$, with $\theta'_{t-1} > \theta_{t-1}$ (again, in terms of period- t consumption, and normalized by the density $f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))$) is equal to the last term in the right-hand side of (20).

Note that the first two terms in the right-hand side of (20) vanish when the agent's utility is linear in consumption, as in the micro literature. This is because, under transferable utility, the increase in the expected future payoff of type θ^{t-1} is perfectly offset by the decrease in the compensation paid to this type in period $t - 1$. Combining (20) with (21), it is then easy to see that, with transferable utility, the shadow cost of increasing the period- t continuation payoff of all period- t types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t \geq \theta_t$, in terms of period-1 consumption, is equal to the impulse response

$$I_1^t(\theta^t, y^{t-1}(\theta^{t-1})) = \prod_{i=1}^{t-1} I_i^{i+1}(\theta^{i+1}, y^i)$$

of θ_t to θ_1 , as noticed in the micro literature (see, e.g., Pavan, Segal, and Toikka (2014)). In this case, the dynamics of distortions under second-best allocations are entirely driven by the dynamics of the impulse responses I_1^t of period- t types to period-1 types. In contrast, with non-transferable utility, the shadow cost of increasing the period- t continuation utility of all period- t types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t \geq \theta_t$, in terms of period-1 consumption, depends also on the impulse responses I_s^t , $t > s$, of period- t types to intermediate types. This property has important implications for the dynamics of distortions (see the discussion after Theorem 1).

5 Wedges

Equipped with the results in the previous sections, we now show how the various forces responsible for the dynamics of distortions under second-best allocations can be summarized in a concise formula for the wedges.

Definition 2. The period- t “wedge” at history θ^t , under the rule χ , is given by

$$W_t(\theta^t) \equiv v^{P'}(y_t(\theta^t)) + LD_t^{FB;\chi}(\theta^t) - \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))}.$$

To understand the definition, recall that (first-best) efficiency requires that the marginal cost to the principal of asking for higher period- t output (the term $\psi_y(y_t(\theta^t), \theta_t)/v'(c_t(\theta^t))$ in the formula for the wedges) be equalized to its marginal benefit, where, with endogenous private information, the latter takes into account also the effect of higher period- t output on the principal's and the agent's joint future surplus, as captured by the term $LD_t^{FB;\chi}(\theta^t)$ introduced above. The period- t wedge is thus the discrepancy between the marginal benefit and the marginal cost of higher period- t output at the proposed allocation. Importantly, such discrepancy is computed holding fixed the policies that determine future allocations, so as to highlight the part of the inefficiency that pertains to the period- t allocations.

In a taxation problem, wedges are also related to non-linear marginal income tax rates. To see this, note that, in a taxation problem, $v^P(y) = y$, in which case the principal's (dual) problem consists in maximizing the net present value of tax revenues subject to the agents' incentive compatibility constraints and the redistribution constraint (1). As in most of the macro literature, let the agents' return on saving (after deducting any possible linear capital tax rate) be equal to $\tilde{r} = 1/\delta - 1$. Then let $\mathcal{T}_t(y^t)$ denote the total period- t tax bill charged to any agent whose period- t history of past and current incomes is equal to y^t and denote by

$$R_t(\theta^t) \equiv \mathbb{E}^{\lambda|\chi|} \left[\sum_{s=t}^T \delta^{s-t} \mathcal{T}_s(y^s(\tilde{\theta}^s)) \right]$$

the period- t expected net present value of current and future tax bills, for any worker of period- t productivity history equal to θ^t , given the tax code $\mathcal{T} \equiv (\mathcal{T}_t(\cdot))$. As we show in the Appendix, the period- t wedge is related to the current and future marginal tax rates by the following relationship¹²

$$\begin{aligned} W_t(\theta^t) &= \frac{\partial \mathcal{T}_t(y^t(\theta^t))}{\partial y_t} + \delta \frac{\partial}{\partial y_t} \int R_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &+ \delta \mathbb{E}^{\lambda|\chi|} \left[\sum_{s=t+1}^T \delta^{s-t-1} \frac{\partial \mathcal{T}_s(y^s(\tilde{\theta}^s))}{\partial y_t} \frac{v'(c_s(\tilde{\theta}^s))}{v'(c_t(\theta^t))} \right] \end{aligned} \quad (22)$$

with the last two terms equal to zero when $t = T$. In the formulas for $W_t(\theta^t)$ in (22), the income and consumption policies $\chi = (y, c)$ are the ones induced, in equilibrium, by the tax code \mathcal{T} , when agents make their consumption and labor decisions optimally. Contrary to static settings, in a dynamic environment, wedges are thus related to both current and future marginal tax rates taking into account the effects of labor supply at present on the evolution of the agents' productivity. Importantly, note that the dependence of future taxes on current incomes is both direct (the third term in (22)), and indirect, via the effect of a variation in current income on the distribution of future productivity (the second term in (22)).¹³

Following the tradition in the macro and public finance literature, hereafter we will consider the following transformation of the wedges

$$\hat{W}_t(\theta^t) \equiv \frac{W_t(\theta^t)}{\frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))}}$$

¹²The derivative in the second term in the right-hand side of (22) is with respect to the measure F_{t+1} , holding the function $R_{t+1}(\theta^{t+1})$ constant.

¹³The formula in (22) can be interpreted as a "labor" wedge. In addition, there is also a "savings" wedge which is given by

$$W_t^S(\theta^t) \equiv 1 - \frac{v'(c_t(\theta^t))}{(1 + \tilde{r}) \delta \int v'(c_{t+1}(\theta^{t+1})) dF_{t+1,y}(\theta_{t+1} | \theta_t, y_t(\theta^t))}$$

where \tilde{r} is the rate of return on savings net of any possible linear capital tax. Such intertemporal wedges have been studied extensively in the received new dynamic public finance literature (see, e.g., Albanesi and Armenter (2012) and the references therein). The forces shaping such wedges are fairly well understood. In this paper, we thus abstain from discussing them further and, instead, focus on the labor wedges.

which captures the wedge relative to marginal cost of higher period- t output (in consumption terms), at the proposed allocations. The advantage of expressing the inefficiencies in terms of relative wedges comes from the fact that the latter are absolute (i.e., percentage) numbers, and hence unit free. We will refer to $\hat{W}_t(\theta^t)$ as the *relative wedge*.

Fix the rule χ , and let

$$h_t(\theta^t, y^t) \equiv -\frac{I_1^t(\theta^t, y^{t-1})}{\eta_1(\theta_1)} \psi_\theta(y_t, \theta_t)$$

denote the period- t ‘‘handicap.’’ Note that, while the wedges measure marginal distortions in the allocations, the ex-ante net present value of the handicaps measures the expected total surplus the principal must leave to the agent, over and above the level $V_1(\underline{\theta}_1)$ left to the lowest period-1 type, to induce truthful revelation of the agent’s private information:

$$\mathbb{E} \left[V_1(\tilde{\theta}_1) \right] = \mathbb{E} \left[\sum_t \delta^{t-1} h_t(\tilde{\theta}^t, y^t(\tilde{\theta}^t)) \right] + V_1(\underline{\theta}_1).$$

Then let

$$\begin{aligned} \hat{W}_t^{RRN}(\theta^t) &\equiv \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1})) |\psi_{y\theta}(y_t(\theta^t), \theta_t)|}{\eta_1(\theta_1) \psi_y(y_t(\theta^t), \theta_t)}, \\ \Omega_t(\theta^t) &\equiv \frac{\frac{\partial}{\partial y_t} \mathbb{E}^{\lambda[\chi]|\theta^t, y_t(\theta^t)} [\sum_{\tau=t+1}^T \delta^{\tau-t} h_\tau(\tilde{\theta}^\tau, y^\tau(\tilde{\theta}^\tau))]}{\psi_y(y_t(\theta^t), \theta_t)}, \\ TC_t(\theta^t) &\equiv \frac{1-F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \left[\begin{array}{l} \int_{\bar{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} \frac{dF_t(\theta'_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}{1-F_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \\ - \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \end{array} \right], \\ RA_t(\theta^t) &\equiv v'(c_t(\theta_t)) \left[\eta_1(\theta_1) \sum_{\tau=2}^t \frac{TC_\tau(\theta^\tau)}{I_1^\tau(\theta^\tau, y^{\tau-1}(\theta^{\tau-1}))} + \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} \frac{dF_1(\theta'_1)}{1-F_1(\theta_1)} \right], \end{aligned} \quad (23)$$

and

$$D_t(\theta^t) \equiv r v'(c_t(\theta_t)) \int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1) \left(\frac{\int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) \frac{dF_1(\theta'_1)}{1-F_1(\theta_1)}}{\int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1)} \right).$$

The next theorem is the paper’s main result:

Theorem 1. *Suppose the rule $\chi = (y, c)$ identifies the second-best allocations. Then at any period $t \geq 1$, with $\lambda[\chi]$ -probability one, the relative wedge is given by*

$$\hat{W}_t(\theta^t) = [RA_t(\theta^t) - D_t(\theta^t)] \left[\hat{W}_t^{RRN}(\theta^t) + \Omega_t(\theta^t) \right]. \quad (24)$$

The theorem provides a convenient representation of wedge dynamics under second-best allocations. In particular, it illustrates how the agent’s risk-aversion, the persistence and endogeneity of the agent’s private information, and the principal’s preferences and constraints for redistribution interact over time in shaping the allocations under optimal contracts.

The term \hat{W}_t^{RRN} in (24) is the period- t (relative) wedge when the agent is risk neutral (that is, when $v(c) = c$, all c) and the principal's preferences for redistribution are Rawlsian (equivalently, when the redistribution constraint (1) takes the familiar form $V(\theta_1) \geq \kappa$ all θ_1 , as typically assumed in the micro literature).

The term RA_t , in turn, is a correction due to the agent's risk-aversion. This correction accounts for all the extra costs to the principal of moving compensation across time and across (endogenous) productivity histories, over and above the costs coming from the need to leave surplus to the agent to incentivize truthful information revelation. Clearly such correction is absent when the agent is risk neutral, in which case $RA_t(\theta^t) = 1$, all θ^t . To see this more clearly, recall from the discussion following Proposition 2 that, with non-transferable utility, the period- t shadow cost to the principal of providing one extra unit of continuation utility to all types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t > \theta_t$ is equal to $\lambda_t(\theta^t)$. Now combine (20) with the initial condition (21) to observe that

$$\lambda_t(\theta^t) = \sum_{\tau=2}^t I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1})) TC_\tau(\theta^\tau) + I_1^t(\theta^t, y^{t-1}(\theta^{t-1})) \lambda_1(\theta_1).$$

The functions $TC_t(\theta^t)$ can be thought of as the net *transaction costs* the principal has to incur to increase by one unit the continuation utility of all period- τ types $(\theta^{\tau-1}, \theta'_\tau)$ with $\theta'_\tau > \theta_\tau$, while adjusting the compensation of type $\theta^{\tau-1}$ so as to hold this type's continuation utility constant. In other words, these are the costs to the principal of backward shifting (i.e., of rolling backward) the information rents given to the agents in periods other than the first one.

Recall that, when the principal increases by one unit the continuation utility of all types $(\theta^{t-1}, \theta'_t)$ with $\theta'_t > \theta_t$, she then also needs to increase by

$$\delta \frac{\partial [1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{\partial \theta_{t-1}}$$

the continuation utility of all period- $(t-1)$ types $\theta^{t-1} = (\theta^{t-2}, \theta'_{t-1})$ with $\theta'_{t-1} > \theta_{t-1}$. Such an increase is necessary to preserve such types' incentives to report truthfully. Because of such an increase, the principal can, however, reduce the compensation of type θ^{t-2} to hold the latter type's continuation utility constant. The net cost to the principal of increasing the continuation utility of all period- $(t-1)$ types $(\theta^{t-2}, \theta'_{t-1})$, with $\theta'_{t-1} > \theta_{t-1}$, normalized by the density $f_{t-1}(\theta_{t-1} | \theta_{t-2}, y_{t-2}(\theta^{t-2}))$ is thus equal to $TC_{t-1}(\theta^{t-1}) I_{t-1}^t(\theta^t, y^{t-1}(\theta^{t-1}))$. Continuing backwards all the way to period one and summing over all intermediate periods, we thus have that the *total* shadow cost of providing one extra unit of continuation utility to all types $(\theta^{t-1}, \theta'_t)$, with $\theta'_t > \theta_t$, in terms of period-1 consumption, is equal to $\lambda_t(\theta^t)$.

Importantly, note that, with non-transferable utility, the shadow costs $\lambda_t(\theta^t)$ depend not only on the impulse response of period- t types to period-1 types, but also on the impulse responses of period- t types to all intermediate period- s types, $1 < s < t$. This property, which has been ignored in the micro literature by focusing on the transferable utility case, has important implications for the dynamics of wedges under second-best allocations.

Next, consider the term D_t in the formula for the wedges. This term is a second correction that controls for the principal's preferences for redistribution when the latter differ from the ones in the Rawlsian benchmark (equivalently, for the possibility that the redistribution/participation constraint (1) takes the integral form $\int_{\underline{\theta}_1}^{\bar{\theta}_1} q(V_1(\theta'_1)) dF_1(\theta'_1) \geq \kappa$, as typically assumed in the macro and/or (dynamic) public finance literature, instead of the interim form $V(\theta_1) \geq \kappa$, all θ_1 , as typically assumed in the micro literature). While, with risk-averse agents, the term RA_t tends to amplify the other two terms in the wedge formula, the term D_t tends to dampen them. The reason is that, when the principal values the utility of period-1 types above $\underline{\theta}_1$, the costs to the principal of increasing the agent's expected lifetime utility for incentive reasons is diminished by the positive effect that such an increase has on the participation/redistribution constraint (1).

Finally, consider the term Ω_t in the wedge formula. This term summarizes all the effects due to the endogeneity of the agent's private information. Naturally, this term is equal to zero when the type process is exogenous. As anticipated above, when impulse responses $I_1^t(\theta^t, y^{t-1})$ are nondecreasing in past output, y_{t-1} , and the kernels $F_t(\theta_t|\theta_{t-1}, y_{t-1})$ are nonincreasing in y_{t-1} , the term Ω_t is typically positive, thus contributing to higher wedges. In models of taxation with learning-by-doing, a positive dependence of the impulse responses on past output may capture the idea that learning-by-doing has stronger effects for highly productive workers than for less productive ones, a property that seems plausible, albeit whose empirical support has not been overwhelming (see, for example, the discussion in Stantcheva, 2017). Similarly, in a trade model with habit formation, the property of impulse responses increasing in past trades may reflect the idea that habit formation is stronger for individuals with high willingness to consume. In all these cases, the endogeneity of the agent's private information calls for additional distortions in the second-best allocations. By further distorting downwards the allocations y_t , the principal economizes on the costs of future incentives by shifting the period- $(t+1)$ type distribution towards levels that command smaller continuation payoffs (equivalently, smaller rents).

As anticipated above, the theorem favors a useful reinterpretation of the key forces responsible for the dynamics of distortions in the special cases considered in the literature. Start with the case of a planner with Rawlsian preferences for redistribution facing a risk-neutral agent (recall that this specification also corresponds to the case of transferable utility and interim participation constraints considered in the micro literature). In this case, $RA_t(\theta^t) = 1$ all t , all θ^t , and $D_t(\theta^t) = 0$, all θ^t . Consider first the case where the agent's private information is exogenous, which is the benchmark in most of the existing literature. In this case, the formula for the wedges reduces to

$$\hat{W}_t(\theta^t) = \hat{W}_t^{RRN}(\theta^t) = \frac{I_1^t(\theta^t) |\psi_{y\theta}(y_t(\theta^t), \theta_t)|}{\eta_1(\theta_1) \psi_y(y_t(\theta^t), \theta_t)}.$$

When $t = 1$, the above formula coincides with the one

$$\hat{W}_1^{RRN}(\theta_1) = \frac{1}{\eta_1(\theta_1)\theta_1} \epsilon_{\theta}^{\psi_y}(y_1(\theta_1), \theta_1)$$

familiar from the early public finance literature (see, e.g., Mirrlees (1971), Diamond (1998), Saez (2001)), with

$$\epsilon_{\theta}^{\psi_y}(y_1(\theta_1), \theta_1) \equiv \frac{|\psi_{y\theta}(y_1(\theta_1), \theta_1)|\theta_1}{\psi_y(y_1(\theta_1), \theta_1)}$$

denoting the elasticity of the marginal disutility of labor with respect to the agent's productivity. For any period $t > 1$, instead, the formula for $\hat{W}_t^{RRN}(\theta^t)$ is adjusted by taking into account the intertemporal informational linkage between type θ^t and type θ_1 , as captured by the impulse response $I_1^t(\theta^t)$ of θ_t to θ_1 . A higher impulse response of θ^t to θ_1 implies a larger increase in the information rent that the principal must leave to all agents whose period-1 productivity exceeds θ_1 when she asks for higher output at history θ^t . Hence, the higher the impulse, the higher the distortion. The dynamics of the wedges in this case are then driven entirely by the (exogenous) dynamics of the impulse responses, as discussed at length in the micro literature (see, e.g., Pavan, Segal, and Toikka, (2014)).

Next, consider the case in which the planner's aversion to inequality is less than Rawlsian (i.e. $r = 1$), but continue to assume that the agent's utility is linear in consumption and that private information is exogenous. In this case, the correction in the wedges due to the principal's lower aversion to inequality is constant over time and equal to

$$1 - D_1(\theta_1) := 1 - \frac{\int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) \frac{dF_1(\theta'_1)}{1 - F_1(\theta_1)}}{\int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1)}.$$

The dynamics of the wedges then continue to be driven by the dynamics of the impulse responses, with the only difference relative to the benchmark with Rawlsian preferences for redistribution coming from the scaling of the wedges $\hat{W}_t^{RRN}(\theta^t)$ by the correction term $1 - D_1(\theta_1)$ which is non-negative, less than one, and increasing in θ_1 . Importantly, while such correction has no effect on the qualitative properties of the dynamics of distortions, it may have non-negligible effects on the progressivity of the period-1 wedges.

Next, consider the case in which the agent is strictly risk averse (i.e., his utility over consumption v is strictly concave), and continue to assume that the process governing the evolution of the agent's private information is exogenous. A special case of interest to the new dynamic public finance literature is that of a planner with utilitarian preferences for redistribution (see, e.g., Farhi and Werning (2013), Kapicka (2013) and Stantcheva (2016)). In this case,

$$RA_t(\theta^t) = v'(c_t(\theta_t)) \left[\eta_1(\theta_1) \sum_{\tau=2}^t \frac{TC_{\tau}(\theta^{\tau})}{I_1^{\tau}(\theta^{\tau}, y^{\tau-1}(\theta^{\tau-1}))} + \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} \frac{dF_1(\theta'_1)}{1 - F_1(\theta_1)} \right],$$

and

$$D_t(\theta^t) \equiv v'(c_t(\theta_t)) \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1).$$

Relative to the Rawlsian-risk-neutral benchmark, the net correction to the wedges is then equal to

$$RA_t(\theta^t) - D_t(\theta^t) = v'(c_t(\theta_t)) \left[\int_{\theta_1}^{\bar{\theta}_1} \frac{\eta_1(\theta_1) \sum_{\tau=2}^t \frac{TC_\tau(\theta^\tau)}{I_1^\tau(\theta^\tau, y^{\tau-1}(\theta^{\tau-1}))} + \frac{1}{v'(c_1(\theta'_1))} \frac{dF_1(\theta'_1)}{1-F_1(\theta_1)} - \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1) \right].$$

This literature also typically assumes, in the calibrations, that the process governing the evolution of the agents' private information is a random walk, which amounts to assuming that $I_1^\tau = 1$, all τ . The key finding in this literature is that, in such economies, distortions tend to increase over the lifecycle. The literature has attributed this finding to the fact that the risk the agents are exposed to, captured by the variance of θ_t , as perceived at time 1, increases with time. To shield the agents from risk, the principal distorts downwards the agents' labor supply, which permits him to reduce the volatility in the agents' compensation necessary to incentivize them to reveal their private information. Because such risk increases over time, so do the distortions.

Our formula permits us to qualify that such intuition is incomplete. The dynamics of the wedges in such economies are also driven by the dynamics of the transaction costs necessary to roll backwards the agents' information rents (the terms $TC_\tau(\theta^\tau)$ identified above). Such costs can grow over time also when the variance of the agents' types remains constant, or even declines with t . Such costs grow with the number of periods a rent is rolled backwards, as explained above. As a result, to contain such costs, the planner may optimally increase the distortions over time, even if the risk the agents are exposed to remains constant, or declines, with t .

Another special case, of interest for example to the managerial compensation literature, is that of a principal with Rawlsian preferences for redistribution (recall that the latter property is equivalent to the imposition of interim participation constraints) facing a risk-averse agent. In this case,

$$RA_t(\theta^t) = v'(c_t(\theta_t)) \left[\eta_1(\theta_1) \sum_{\tau=2}^t \frac{TC_\tau(\theta^\tau)}{I_1^\tau(\theta^\tau, y^{\tau-1}(\theta^{\tau-1}))} + \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} \frac{dF_1(\theta'_1)}{1-F_1(\theta_1)} \right],$$

while $D_t(\theta^t) = 0$. The relative wedges are then equal to

$$\begin{aligned} \hat{W}_t(\theta^t) &= \frac{|\psi_{y\theta}(y_t(\theta^t), \theta_t)|}{\psi_y(y_t(\theta^t), \theta_t)} v'(c_t(\theta_t)) \sum_{\tau=2}^t I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1})) TC_\tau(\theta^\tau) \\ &\quad + \frac{|\psi_{y\theta}(y_t(\theta^t), \theta_t)|}{\psi_y(y_t(\theta^t), \theta_t)} v'(c_t(\theta_t)) \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_1(\theta_1)} \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} \frac{dF_1(\theta'_1)}{1-F_1(\theta_1)}. \end{aligned}$$

As the above formula shows, in this case, the dynamics of the wedges are also driven by the dynamics of the impulse responses of period- t types to all intermediate types (i.e., by I_τ^t , $\tau = 1, \dots, t-1$). To the best of our knowledge, this channel too has not been noticed in previous work and plays a major role for the dynamics of distortions in economies with risk-averse agents. As anticipated in the Introduction, such intermediate impulse responses contribute to the total cost the principal must incur to roll backwards the agents' information rents. In future work, it would be interesting to study in more detail how these intermediate impulse responses contribute to the dynamics of the distortions in various economies of interest (see also the numerical analysis in the next section).

Lastly, consider models with endogenous private information. In such economies, the terms $\hat{W}_t^{RRN}(\theta^t)$ and $RA_t(\theta^t) - D_t(\theta^t)$ in the formulas for the wedges retain the interpretation discussed above. However, naturally, the conditional distributions and the impulse response functions in the formulas for $\hat{W}_t^{RRN}(\theta^t)$ and $RA_t(\theta^t) - D_t(\theta^t)$ should now be interpreted as the ones corresponding to the endogenous process induced by the rule χ . The novel effects originating in the endogeneity of the agent's private information are the ones captured by the interactions of the terms $\Omega_t(\theta^t)$ with the correction terms $RA_t(\theta^t) - D_t(\theta^t)$.

Note that the expectational term in the numerator of $\Omega_t(\theta^t)$ is the expected discounted sum of all future handicaps, $h_s(\theta^s)$, $s > t$, with the latter capturing the consumption losses to the principal due to the need to leave information rents to the agent in future periods to incentivize truthful reporting. As anticipated above, in general, there are two channels by which the endogeneity of the process affects the expectation of future handicaps. The first one is by shifting the distribution of future types, holding fixed all future handicaps. In the case of taxation with learning-by-doing, the future handicaps are typically increasing in the agents' future types. When this is the case, by shifting the distribution of future productivity towards levels for which the handicaps are higher, learning-by-doing contributes positively to the expectation of future handicaps. Other things equal, this channel thus contributes to larger distortions in labor supply in the form of higher wedges (see Makris and Pavan (2018)). We expect similar effects in trade models with habit formation, where higher past consumption increases future willingness to pay. On the contrary, when past consumption reduces the interest in future purchases, as in certain trade models featuring intertemporal substitution, the endogeneity of the agent's private information may contribute to a reduction in the wedges, as one can see from the decomposition in Theorem 1.

The second channel by which the endogeneity of the type process affects the wedges is through the effect of current output on the impulse responses of future types to current ones, and thereby through its direct effect on future handicaps, for a given distribution of future types. As discussed above, in general, future handicaps may be either increasing or decreasing in current output, depending on whether the impulse responses of future types to the current ones are increasing or decreasing in current output. For example, in the case of taxation under learning-by-doing, impulse responses are increasing in current output when current skills and current output are complements in the determination of future skills. In this case, this second channel adds to the first one and learning-by-doing contributes unambiguously to higher wedges. When, instead, current skills and current output are substitutes in the determination of future skills so that higher output at present reduces the impulse responses of future types to current ones, this second channel contributes to a lower negative effect of learning-by-doing on expected future losses and hence to lower wedges.

As discussed in Makris and Pavan (2018), the endogeneity of the agent's private information may have non-negligible effects not only on the level of the wedges but also on their dynamics and progressivity. In the next section, we investigate how such endogeneity interacts with the agents' risk aversion and the principal's preferences for redistribution in shaping the dynamics of

distortions. In particular, we investigate whether such effects are more pronounced at the beginning of the relationship, when the change in the agents' types has impact over a larger number of periods, or later in the relationship when there are fewer periods ahead in which the benefits of the change can be exploited, but also more urgency to act in case the realized type is not the desired one.

6 Numerical Analysis

TBA

7 Appendix

Proof of Proposition 1. For any $t > 1$, let $\pi_t(\theta^{t-1})$ be the multiplier of the period- t promise keeping constraint (2). Let also $q'_t(V_t(\theta_t)) \equiv 1$ for $t > 1$ and $q'_t(V_t(\theta_t)) \equiv q'(V_1(\theta_1))$ for $t = 1$.

Start with the case of $r = 1$, and let π_1 be the multiplier of the redistribution/participation constraint (3), which is an integral constraint. At the solution to the planner's problem, the following necessary conditions with respect to $y_t(\theta^t)$, $V_t(\theta^t)$ and $\Pi_{t+1}(\theta^t)$ must hold with $\lambda[\chi]$ -probability one:

$$\begin{aligned} v^{P'}(y_t(\theta^t)) - \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} + \delta \frac{\partial}{\partial y_t} Q_{t+1}^{FB}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t)) &= 0 \text{ for any } t = 1, \dots, T \\ \frac{1}{v'(c_t(\theta^t))} + \pi_t(\theta^{t-1}) q'_t(V_t(\theta_t)) &= 0 \text{ for any } t = 1, \dots, T \\ \frac{1}{v'(c_t(\theta^t))} + \frac{\partial}{\partial \Pi_{t+1}} Q_{t+1}^{FB}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t)) &= 0 \text{ for any } t < T \end{aligned}$$

with $\pi_1(\theta^0) \equiv \pi_1$.

Next, use the envelope theorem to establish that

$$\delta \frac{\partial}{\partial y_T} Q_{T+1}^{FB}(\theta, y_T(\theta), \Pi_{T+1}(\theta)) = 0$$

while, for $t < T$,

$$\delta \frac{\partial}{\partial y_t} Q_{t+1}^{FB}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t)) = \delta \frac{\partial}{\partial y_t} \int \{V_{t+1}^P(\theta^{t+1}) - \pi_{t+1}(\theta^t) V_{t+1}(\theta^{t+1})\} dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t))$$

and

$$\frac{\partial}{\partial \Pi_{t+1}} Q_{t+1}^{FB}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t)) = \pi_{t+1}(\theta^t).$$

Combining the above optimality conditions and using the definition of $LD_t^{FB;\chi}(\theta^t)$ gives the result.

Now consider the case where $r = 0$. Recall that, in this case, the redistribution/participation constraint (3) takes the form $V_1(\theta_1) \geq \kappa$ all $\theta_1 \in \Theta_1$. Let $\hat{\pi}_1(\theta_1)$ be the (Kuhn-Tucker) multiplier of the redistribution constraint (3). At the solution to the planner's problem, the same necessary

conditions as for the case $r = 1$ must hold, except for the second one for $t = 1$ which must be replaced by

$$\frac{1}{v'(c_1(\theta_1))} + \hat{\pi}_1(\theta_1) = 0.$$

Combining the various necessary conditions yields the result. Q.E.D.

Proof of Proposition 2. In any period t , the principal is facing an optimal control problem with integral constraints, for any given "state" $(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}), Z_t(\theta^{t-1}))$.

For any $t > 1$, let $\pi_t(\theta^{t-1})$ and $\xi_t(\theta^{t-1})$ be the multipliers of the two "promise-keeping" integral constraints (15) and (16) associated with the levels of expected utility, $\Pi_t(\theta^{t-1})$, and marginal expected utility, $Z_t(\theta^{t-1})$, promised in period $t - 1$. Also note that the redistribution/participation constraint (13) can be conveniently rewritten as

$$\Pi_1 = (1 - r)V_1(\underline{\theta}_1) + r \int q(V_1(\theta_1))dF_1(\theta_1) \quad (25)$$

by letting $\Pi_1 \equiv \kappa$. Then let π_1 be the multiplier associated with constraint (25) and let $\xi_1(\theta^0) \equiv 0$. With some abuse of notation, also let $q'_1(V; r) \equiv r q'(V)$ and, for any $t > 1$, $q'_t(V; r) = 1$, all V , all $r \in \{0, 1\}$. Finally, for any $t \geq 1$, any $\theta^t \in \Theta^t$, let $\mu_t(\theta^t)$ denote the co-state variable associated with the law of motion of the state variable $V_t(\theta^t)$, as specified by Condition (14) in the principal's problem.

We start with the case $r = 1$, and then move to the case $r = 0$.

Case $r = 1$.

In this case, the redistribution constraint (25) becomes

$$\int q(V_1(\theta_1))dF_1(\theta_1) = \Pi_1. \quad (26)$$

The solution to the principal's problem must then satisfy the following optimality conditions with $\lambda[\chi]$ -probability one:

$$\left\{ v^{Pr}(y_t(\theta^t)) - \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} + LD_t^X(\theta^t) \right\} f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) - \psi_{y\theta}(y_t(\theta^t), \theta_t) [\mu_t(\theta^t) - \xi_t(\theta^{t-1}) I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1})) f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))] = 0, \text{ for all } t, \quad (27)$$

$$\frac{\partial \mu_t(\theta^t)}{\partial \theta_t} = f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \times \left\{ \frac{1}{v'(c_t(\theta^t))} + \pi_t(\theta^{t-1}) q'_t(V_t(\theta^t); r) \right\}, \text{ for all } t, \quad (28)$$

$$\mu_t(\theta^{t-1}, \underline{\theta}_t) = 0, \text{ for all } t, \quad (29)$$

$$\mu_t(\theta^{t-1}, \bar{\theta}_t) = 0, \text{ for all } t, \quad (30)$$

$$\frac{\delta}{v'(c_t(\theta^t))} + \delta \frac{\partial Q_{t+1}}{\partial \Pi_{t+1}} = 0, \text{ for all } 1 \leq t < T, \quad (31)$$

$$\begin{aligned} & \delta [\mu_t(\theta^t) - \xi_t(\theta^{t-1})I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1}))f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))] + \\ & + \delta \frac{\partial Q_{t+1}}{\partial Z_{t+1}} f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) = 0, \text{ for all } 1 \leq t < T, \end{aligned} \quad (32)$$

where $\partial Q_{t+1}/\partial \Pi_{t+1}$ and $\partial Q_{t+1}/\partial Z_{t+1}$ are shortcuts for

$$\frac{\partial}{\partial \Pi_{t+1}} Q_{t+1}(\theta^t, y_t(\theta^t), V_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))$$

and

$$\frac{\partial}{\partial Z_{t+1}} Q_{t+1}(\theta^t, y_t(\theta^t), V_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))$$

respectively, and where $\pi_1(\theta^0) \equiv \pi_1$.

Next, use the envelope theorem to observe that $\frac{\partial Q_{t+1}}{\partial \Pi_{t+1}} = \pi_{t+1}$ and $\frac{\partial Q_{t+1}}{\partial Z_{t+1}} = \xi_{t+1}$. It follows that Conditions (31) and (32) can be rewritten as

$$\frac{1}{v'(c_t(\theta^t))} + \pi_{t+1}(\theta^t) = 0, \text{ for all } 1 \leq t < T, \quad (33)$$

and

$$\begin{aligned} & [\mu_t(\theta^t) - \xi_t(\theta^{t-1})I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1}))f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))] + \\ & + \xi_{t+1}(\theta^t)f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) = 0, \text{ for all } 1 \leq t < T. \end{aligned} \quad (34)$$

The solution to the principal's problem is then given by Conditions (27), (28), (29), (30), (33), and (34) above, along with the laws of motions for the agent's continuation payoff as given by (14), the constraints (15), (16), and (25), the consumption identity

$$c_t(\theta^t) = C(V_t(\theta^t) + \psi(y_t(\theta^t), \theta_t) - \delta \Pi_{t+1}(\theta^t))$$

and the fact that, for $t = T$,

$$LD_T^X(\theta) = \delta \frac{\partial}{\partial y_T} Q_{T+1}(\theta, y_T(\theta), \Pi_{T+1}(\theta), Z_{T+1}(\theta)) = 0$$

while, for $t < T$,

$$\begin{aligned} LD_t^X(\theta^t) &= \delta \frac{\partial}{\partial y_t} Q_{t+1}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t)) \\ &= \delta \frac{\partial}{\partial y_t} \int \hat{Q}_{t+1}(\theta^{t+1}, y_{t+1}(\theta^{t+1}), V_{t+1}(\theta^{t+1}), \Pi_{t+2}(\theta^{t+1}), Z_{t+2}(\theta^{t+1})) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &\quad - \pi_{t+1}(\theta^t) \delta \frac{\partial}{\partial y_t} \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &\quad - \xi_{t+1}(\theta^t) \delta \frac{\partial}{\partial y_t} \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)). \end{aligned} \quad (35)$$

Note that in writing (35) we used the fact that

$$\begin{aligned} & \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &= \int [-\psi_\theta(y_{t+1}(\theta^{t+1}), \theta_{t+1}) + \delta Z_{t+2}(\theta^{t+1})] I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)). \end{aligned}$$

We proceed to show that the above optimality conditions imply those in the proposition. Let $\lambda_t(\theta^t) \equiv -\mu_t(\theta^t)/f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) + \xi_t(\theta^{t-1})I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1}))$ for all $t \geq 1$. It is then easy to see that Condition (27) is equivalent to Condition (19) in the proposition.

Next, combine Conditions (33) and (34) with the law of motion for the co-state variable (28) for $t > 1$, and use the boundary conditions (29) and (30) for $t > 1$ to obtain the following alternative representation of the law of motion of the co-state variables:

$$\begin{aligned} \mu_t(\theta^t) = \mu_t(\theta^{t-1}, \bar{\theta}_t) - \int_{\theta_t}^{\bar{\theta}_t} \frac{\partial \mu_t(\theta^{t-1}, \theta'_t)}{\partial \theta_t} d\theta'_t = & - \int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \\ & + \frac{[1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{v'(c_{t-1}(\theta^{t-1}))}. \end{aligned} \quad (36)$$

When evaluated at $\theta_t = \underline{\theta}_t$, the above expression yields

$$\begin{aligned} 0 = \mu_t(\theta^{t-1}, \underline{\theta}_t) = & - \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \\ & + \frac{1}{v'(c_{t-1}(\theta^{t-1}))}. \end{aligned} \quad (37)$$

It follows that the expression in (37) is equivalent to

$$\frac{1}{v'(c_{t-1}(\theta^{t-1}))} = \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \quad (38)$$

which is the Rogerson-inverse-Euler condition (18) in the proposition.

Combining (36) with (38), we also have that

$$\begin{aligned} \frac{-\mu_t(\theta^t)}{f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} = & \frac{1}{f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} \int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) \\ & - \frac{\int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1})) [1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))]}{f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} \end{aligned} \quad (39)$$

Using the definition of $\lambda_t(\theta^t)$ above, we then have that, for any $t > 1$, Condition (39) is equivalent to

$$\begin{aligned} \lambda_t(\theta^t) = & \frac{\int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} - \frac{\int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1 - F_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{v'(c_t(\theta^{t-1}, \theta'_t))} dF_t(\theta'_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} \\ & + I_{t-1}^t(\theta^t, y^{t-1}(\theta^{t-1}))\xi_t(\theta^{t-1}). \end{aligned} \quad (40)$$

Then note that, by definition of $\lambda_t(\theta^t)$ and (34), for all $t > 1$, $\lambda_{t-1}(\theta^{t-1}) = \xi_t(\theta^{t-1})$. This implies that Condition (40) is equivalent to Condition (20) in the proposition.

Next, consider $t = 1$. Applying the law of motion of the co-state variable (28) to $t = 1$, and using the boundary conditions (29) and (30) for $t = 1$, we then have that

$$\frac{-\mu_1(\theta_1)}{f_1(\theta_1)} = \frac{1}{f_1(\theta_1)} \left[\int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1) + \pi_1 \int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1) \right]. \quad (41)$$

When evaluated at $\theta_1 = \underline{\theta}_1$, the above expression yields the following formula (recall that $\mu_1(\underline{\theta}_1) = 0$)

$$0 = \int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1) + \pi_1 \int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1),$$

from which we obtain that

$$\pi_1 = \frac{-\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1)}{\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta'_1)) dF_1(\theta'_1)} < 0.$$

Letting $\pi = -\pi_1$, and recalling that, by definition, $\lambda_1(\theta_1) \equiv -\mu_1(\theta_1)/f_1(\theta_1)$, we obtain Condition (21) in the proposition. This completes the proof for the case of $r = 1$.

Case $r = 0$.

The derivation of Conditions (18), (19), and (20) follows from the same steps as in the case $r = 1$. The only differences pertain to the derivation of Condition (21). First, observe that, when $r = 0$, the redistribution constraint (26) in the optimal control problem described above for the case $r = 1$ is absent and replaced by the constraint $V_1(\underline{\theta}_1) = \kappa$. As a result, the multiplier π_1 in the law of motion (28) for $t = 1$ must be set equal to zero. Second, the only relevant period-1 transversality condition is now $\mu_1(\bar{\theta}_1) = 0$. Setting $\pi_1 = 0$ in (28) for $t = 1$, and using the transversality condition $\mu_1(\bar{\theta}_1) = 0$, we thus have that

$$0 = \mu_1(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{\partial \mu_1(\theta'_1)}{\partial \theta_1} d\theta'_1 = \mu_1(\underline{\theta}_1) + \int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1)$$

which implies

$$\mu_1(\underline{\theta}_1) = -\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1). \quad (42)$$

Combining (42) with (28) for $t = 1$, we obtain that

$$\lambda_1(\theta_1) \equiv \frac{-\mu_1(\theta_1)}{f_1(\theta_1)} = \frac{\int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta'_1))} dF_1(\theta'_1)}{f_1(\theta_1)} \quad (43)$$

which is equivalent to Condition (21) in the proposition when specialized to the case $r = 0$. Q.E.D.

Proof of Theorem 1. From Proposition 2, recall that, under the second-best allocation policies, the following optimality condition for output must hold with $\lambda[\chi]$ -probability one:

$$v^{P'}(y_t(\theta^t)) + LD_t^\chi(\theta^t) = \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} - \psi_{y\theta}(y_t(\theta^t), \theta_t) \cdot \lambda_t(\theta^t), \text{ for all } t, \quad (44)$$

where

$$LD_t^\chi(\theta^t) \equiv \delta \frac{\partial}{\partial y_t} Q_{t+1}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))$$

and where $Q_{t+1}(\theta^t, y_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))$ denotes the principal's expected continuation payoff from period $t + 1$ (included) onward, given the period- t history θ^t , the period- t output choice $y_t(\theta^t)$,

the agent's promised expected continuation utility $\Pi_{t+1}(\theta^t)$, and the promised expected marginal continuation utility $Z_{t+1}(\theta^t)$, as defined in the main text.

Next, use (35) to observe that

$$\begin{aligned} LD_t^X(\theta^t) &= \delta \frac{\partial}{\partial y_t} \int [v^P(y_{t+1}(\theta^{t+1})) - c_{t+1}(\theta^{t+1}) \\ &+ \delta Q_{t+2}(\theta^{t+1}, y_{t+1}(\theta^{t+1}), \Pi_{t+2}(\theta^{t+1}), Z_{t+2}(\theta^{t+1}))] dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &- \pi_{t+1}(\theta^t) \delta \frac{\partial}{\partial y_t} \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &- \xi_{t+1}(\theta^t) \delta \frac{\partial}{\partial y_t} \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)), \end{aligned}$$

where recall that $\pi_{t+1}(\theta^t)$ and $\xi_{t+1}(\theta^t)$ are, respectively, the period- $(t+1)$ multiplier of the promise-keeping expected utility constraint (15) and of the promise-keeping expected marginal utility constraint (16).

Next use the result from the proof of Proposition that $\pi_{t+1}(\theta^t) = -1/v'(c_t(\theta^t))$ and $\xi_{t+1}(\theta^t) = \lambda_t(\theta^t)$ to rewrite the above expression for $LD_t^X(\theta^t)$ as

$$\begin{aligned} LD_t^X(\theta^t) &= \delta \frac{\partial}{\partial y_t} \int [v^P(y_{t+1}(\theta^{t+1})) - c_{t+1}(\theta^{t+1}) \\ &+ \delta Q_{t+2}(\theta^{t+1}, y_{t+1}(\theta^{t+1}), \Pi_{t+2}(\theta^{t+1}), Z_{t+2}(\theta^{t+1}))] dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &+ \frac{\delta}{v'(c_t(\theta^t))} \frac{\partial}{\partial y_t} \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \\ &- \delta \lambda_t(\theta^t) \frac{\partial}{\partial y_t} \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}((\theta_t, \theta_{t+1}), y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)). \end{aligned}$$

Now observe that

$$v^P(y_{t+1}(\theta^{t+1})) - c_{t+1}(\theta^{t+1}) + \delta Q_{t+2}(\theta^{t+1}, y_{t+1}(\theta^{t+1}), \Pi_{t+2}(\theta^{t+1}), Z_{t+2}(\theta^{t+1}))$$

is the principal's continuation payoff from period $t+1$ (included) onwards, under the second-best allocations, and thus coincides with $V_{t+1}^P(\theta^{t+1})$, as defined in the main text. Using the definition of $LD_t^{FB;X}(\theta^t)$, we thus have that

$$\begin{aligned} LD_t^X(\theta^t) &= LD_t^{FB;X}(\theta^t) \\ &- \delta \lambda_t(\theta^t) \frac{\partial}{\partial y_t} \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}((\theta_t, \theta_{t+1}), y_t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)). \end{aligned} \quad (45)$$

Using Condition (9) to express $\partial V_{t+1}(\theta^{t+1})/\partial \theta_{t+1}$ as a function of future decisions, along with the property that, for any (τ, s) , $\tau > s$, $I_1^\tau = I_1^s I_s^\tau$, and the definition of period- t handicaps in the main text, we finally have that

$$\begin{aligned} LD_t^X(\theta^t) &= LD_t^{FB;X}(\theta^t) - \frac{\lambda_t(\theta^t) \eta(\theta_1)}{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))} \\ &\cdot \delta \frac{\partial}{\partial y_t} \mathbb{E}^{\lambda[X] | \theta^t, y_t(\theta^t)} \left[\sum_{\tau=t+1}^T \delta^{\tau-t-1} h_\tau(\tilde{\theta}^\tau, y^\tau(\tilde{\theta}^\tau)) \right]. \end{aligned} \quad (46)$$

Replacing (46) into (19), and using the definition of the relative wedges in the main text and of the functions W_t^{RRN} , RA_t , D , and Ω_t yields, after some algebra, the formula for the relative wedges in the theorem. Q.E.D.

Proof of (22). Consider first the worker's problem when he faces a tax schedule \mathcal{T} and a (net of any linear capital tax) rate of return to savings $r^\# = 1/\delta - 1$. At any period t , given current productivity θ_t , history of past earnings, y^{t-1} , and net savings from the previous period, S_t , the associated value function of the worker is given by:

$$\widehat{V}_t(S_t, \theta_t, y^{t-1}) = \max_{c_t, y_t} \left\{ v(c_t) - \psi(y_t, \theta_t) + \delta \int \widehat{V}_{t+1}(S_{t+1}, \theta_{t+1}, (y^{t-1}, y_t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t) \right\}$$

where $\widehat{V}_{T+1}(S_{T+1}, \theta_{T+1}, y^T) \equiv 0$, and where the current net savings transferred to the next period are given by the following law-of-motion

$$S_{t+1} = y_t - \mathcal{T}_t(y_t) - c_t + (1 + r^\#)S_t$$

with $S_{T+1} = 0$ and where we assumed (without loss of generality) that $S_1 = 0$. The FOCs of this problem (after suppressing, for notational simplicity, the dependence of the optimal solution on the "state" (S_t, θ_t, y^{t-1})) are as follows.¹⁴ The FOC with respect to c_t , $t < T$, is

$$v'(c_t) = \delta \int \frac{\partial \widehat{V}_{t+1}(S_{t+1}, \theta_{t+1}, y^t)}{\partial S_{t+1}} dF_{t+1}(\theta_{t+1} | \theta_t, y_t). \quad (47)$$

In turn, using (47), we have that the FOC with respect to y_t is

$$\left[1 - \frac{\partial \mathcal{T}_t}{\partial y_t} \right] v'(c_t) = \psi_y(y_t, \theta_t) - \delta \frac{\partial}{\partial y_t} \int \widehat{V}_{t+1}(S_{t+1}, \theta_{t+1}, y^t) dF_{t+1}(\theta_{t+1} | \theta_t, y_t)$$

where the second term in the right-hand side of the above condition is zero when $t = T$.

Now take the second-best rule χ and, given any tax code \mathcal{T} , define $S(\theta)$ iteratively by¹⁵

$$S_{t+1}(\theta^t) = y_t(\theta^t) - \mathcal{T}_t(y_t(\theta^t)) - c_t(\theta^t) + (1 + r^\#)S_t(\theta^{t-1})$$

with $S_1(\theta^0) \equiv 0$. Forwarding the above law of motion and combining the result with the definition of V^P , we then have that, when $v^P(y) = y$ and $(1 + r^\#)\delta = 1$,

$$V_{t+1}^P(\theta^{t+1}) = -S_{t+1}(\theta^t)/\delta + R_{t+1}(\theta^{t+1}) \quad (48)$$

¹⁴Here, we discuss the problem of the worker when T is finite, but the proof can easily be modified to allow for an infinite T . In such a case, one could show, by following standard techniques, that a well-defined time-invariant value function exists that satisfies the Bellman equation counterpart of the one that defines \widehat{V}_i above. Moreover, the transversality condition $\lim_{t \rightarrow \infty} S_t$ would replace S_{T+1} above.

¹⁵Note that, given an allocation χ and a tax code \mathcal{T} , we can think of the path $\mathcal{T}_t(y^t)$ as the path of the principal's appropriation, in each period t , of the agent's period- t earnings, given the history of earnings y^t , and of $S_{t+1}(\theta^t)$ as the transfer of resources in the next period.

where

$$R_{t+1}(\theta^{t+1}) \equiv \mathbb{E}^{\lambda[x]|\theta^{t+1}} \left[\sum_{s=t+1}^T \delta^{s-t-1} \mathcal{T}_s(y^s(\tilde{\theta}^s)) \right]$$

is the NPV of tax revenues from period $t + 1$ onwards. Moreover, for any t and any θ^t ,

$$V_t(\theta^t) = \widehat{V}_t(S_t(\theta^{t-1}), \theta_t, y^{t-1}(\theta^{t-1})). \quad (49)$$

Furthermore, after using the envelope theorem for the worker's problem and the FOC above for c_{t+1} , we have that, for any $s \geq t$,

$$\begin{aligned} & \frac{\partial \widehat{V}_{s+1}(S_{s+1}(\theta^s), \theta_{s+1}, y^s(\theta^s))}{\partial y_t} = \\ & - \frac{\partial \mathcal{T}_s(y^{s+1}(\theta^{s+1}))}{\partial y_t} v'(c_{s+1}(\theta^{s+1})) + \delta \int \frac{\partial \widehat{V}_{s+2}(S_{s+2}(\theta^{s+1}), \theta_{s+2}, y^{s+1}(\theta^{s+1}))}{\partial y_t} dF_{s+2}(\theta_{s+2} | \theta_{s+1}, y_{s+1}(\theta^{s+1})). \end{aligned}$$

By iterating forward the last condition over s , we obtain that

$$\frac{\partial \widehat{V}_{t+1}(S_{t+1}(\theta^t), \theta_{t+1}, y^t(\theta^t))}{\partial y_t} = -\mathbb{E}^{\lambda[x]|\theta^{t+1}} \left[\sum_{s=t+1}^T \delta^{s-t-1} \frac{\partial}{\partial y_t} \mathcal{T}_s(y^s(\tilde{\theta}^s)) v'(c_s(\tilde{\theta}^s)) \right]. \quad (50)$$

Combining the definition of wedges with (4), (48), (49) and (50), and using the FOC above with respect to y_t , we have that

$$\frac{\partial}{\partial y_t} \int S_{t+1}(\theta^t) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) = 0.$$

We conclude that any given wedge schedule (and thereby any given allocation rule) can be implemented via a system of non-linear marginal income tax schedule that satisfies:

$$\begin{aligned} W_t(\theta^t) &= \frac{\partial \mathcal{T}_t(y^t(\theta^t))}{\partial y_t} + \\ & \delta \left\{ \frac{\partial}{\partial y_t} \int V_{t+1}^P(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) - \frac{1}{v'(c_t(\theta^t))} \int \frac{\partial}{\partial y_t} \widehat{V}_{t+1}(S_{t+1}(\theta^t), \theta_{t+1}, y^t(\theta^t)) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) \right\} = \\ & \frac{\partial \mathcal{T}_t(y^t(\theta^t))}{\partial y_t} + \delta \frac{\partial}{\partial y_t} \int R_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} | \theta_t, y_t(\theta^t)) + \\ & \delta \mathbb{E}^{\lambda[x]|\theta^t} \left[\sum_{s=t+1}^T \delta^{s-t-1} \frac{\partial \mathcal{T}_s(y^s(\tilde{\theta}^s))}{\partial y_t} \frac{v'(c_s(\tilde{\theta}^s))}{v'(c_t(\theta^t))} \right] \end{aligned}$$

with the last two terms equal to zero when $t = T$. Q.E.D

References

- [1] Albanesi, S. and R. Armenter (2012): "Intertemporal Distortions in the Second Best," *Review of Economic Studies*, 79(4), 1271–1307.

- [2] Albanesi, S. and C. Sleet (2006): “Dynamic Optimal Taxation with Private Information,” *Review of Economic Studies*, 73(1): 1-30.
- [3] Best, M.C., and H.J. Kleven, (2013): “Optimal Income Taxation with Career Effects of Work Effort,” mimeo, LSE.
- [4] Bose, S. and M. Makris, (2016): “Rational Learning of Stochastic Intertemporal Complementarity,” mimeo University of Southampton.
- [5] Diamond, P.A., (1998): “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates,” *American Economic Review*, 88(1), 83-95.
- [6] Farhi, E., and I. Werning, (2013): “Insurance and Taxation over the Life Cycle,” *Review of Economic Studies*, 80, 596-635.
- [7] Fershtman, D., and A. Pavan (2017): “Matching Auctions,” mimeo, Northwestern University.
- [8] Garrett, D., and A. Pavan, (2015): “Dynamic Managerial Compensation: a Variational Approach,” *Journal of Economic Theory* 159, 775–818.
- [9] Golosov, M., M. Troshkin, and A. Tsyvinski, (2016): “Redistribution and Social Insurance,” *American Economic Review*, 106(2): 359-86.
- [10] Golosov, M., A. Tsyvinski, and I. Werning, (2006): “New Dynamic Public Finance: A User’s Guide”, NBER Macroeconomic Annual 2006, MIT Press.
- [11] Gorry, A. and E. Oberfield, (2012): “Optimal Taxation Over the Life Cycle”, *Review of Economic Dynamics*, 15(4):551–572.
- [12] Kapicka, M., (2006): “Optimal Income Taxation with Human Capital Accumulation and Limited Record Keeping”, *Review of Economic Dynamics*, 9, 612–639.
- [13] Kapicka, M., (2013): “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach,” *Review of Economic Studies*, 80(3), 1027-1054.
- [14] Kapicka, M., (2015a): “Optimal Mirrleesian Taxation in a Ben-Porath Economy,” *American Economic Journal: Macroeconomics*, Vol. 7, no. 2, 219-48.
- [15] Kapicka, M., (2015b): “Pareto Efficient Income Taxation with Learning by Doing”, mimeo, CERGE-EI.
- [16] Kapicka, M. and J. Neira (2016): “Optimal Taxation with Risky Human Capital,” mimeo, CERGE-EI.
- [17] Kocherlakota, N. R., (2010): *The New Dynamic Public Finance*. Princeton University Press, USA.

- [18] Krause, A., (2009): “Optimal Nonlinear Income Taxation with Learning-by-doing”, *Journal of Public Economics*, 93(9-10), 1098-1110.
- [19] Makris, M. and A. Pavan (2018): “Taxation under Learning-By-Doing,” mimeo, Northwestern University.
- [20] Mirrlees, J., (1971): “An Exploration in the Theory of Optimal Income Taxation,” *Review of Economic Studies*, 38, 175-208.
- [21] Pavan, A., S. Segal, and J. Toikka, (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82(2), 601-653.
- [22] Pavan, A., (2017): “Dynamic Mechanism Design: Robustness and Endogenous Types,” *Advances in Economics and Econometrics*, Proceeding of the 2015 World Congress of the Econometric Society, forthcoming.
- [23] Rogerson, W.P., (1985): “Repeated Moral Hazard,” *Econometrica*, 53(1), 69-76.
- [24] Saez, E., (2001): “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, 68, 205-229.
- [25] Stantcheva, S., (2015): “Learning and (or) Doing: Human Capital Investments and Optimal Taxation” NBER Working Paper 21381.
- [26] Stantcheva, S., (2017): “Optimal Taxation and Human Capital Policies over the Life Cycle,” *Journal of Political Economy*, 125 (6), 1931–1990.