Incentives for Endogenous Types: Taxation under Learning-by-Doing*

Miltiadis Makris
University of Southampton

Alessandro Pavan
Northwestern University

June 2016

Abstract

We study the effects of learning-by-doing on the dynamics of optimal allocations. We show that learning-by-doing may increase expected future welfare losses stemming from the agents' private information. When sufficiently strong, this novel channel may lead to an increase in wedges and to taxes that decrease over time (even in the presence of an insurance motive). Furthermore, learning-by-doing can lead to marginal distortions that are increasing in income, thus contributing to the progressivity of the optimal tax code. We also provide a novel decomposition of the wedge, accommodating for agents' risk aversion, inequality aversion on the principal's side, and endogenous types. This also allows us to clarify various results in the dynamic mechanism design and taxation literature and underline the relative importance of various channels in shaping incentives under optimal contracts.

*For comments and useful suggestions, we thank Marek Kapicka, Paul Klein, Alex Mennuni, Ludovic Renou, Arnau Valladares-Esteban, and seminar participants at the 2014 European Meetings of the Econometric Society, the 2014 Latin America and Caribbean Meetings of the Econometric Society, the 2015 World Congress of the Econometric Society, the 2015 Shanghai Macroeconomics Workshop, CERGE-EI, the Chicago FED Research Department, Northwestern University, and the University of Western Ontario. Laura Doval provided excellent research assistance. The usual disclaimer applies.

1 Introduction

Learning-by-doing refers to a situation where spending more time working on the particular task at hand at present has a positive effect on future productivity/skills/wages (used interchangeably). Thus, one can think of it as human capital investment which is a side-product of the production process. Learning-by-doing is believed to be an important source of productivity growth (Arrow, 1962, Lucas, 1988). A vast literature in labor economics documents the effect of labor experience on wages.[1] Dustmann and Meghir (2005), for instance, find that in the first two years of employment, wages grow at an average rate of 7% for skilled workers, and at a rate of 10% in the first year and 8% in the second for non-skilled workers. Insofar variations in earnings reflect – to some extent – variations in skills/productivities, these findings suggest that labor experience has a significant impact on human capital accumulation.

In this paper, we study the effects of learning-by-doing (hereafter, LBD) on the dynamics of optimal contracts, when the agents' private information (their productivity) evolves endogenously over time (i.e., depends stochastically on past contracted output). In particular, we focus on the effects of LBD on the dynamics of wedges (that is, distortions relative to the complete-information benchmark) introduced in the optimal contract to limit information rents.

In the presence of LBD, agents have incentives to produce more. This implies that, all other things equal, output is higher under LBD than in its absence. However, all other things are not equal. When the agent's private information is endogenous, LBD also affects the expectation of future welfare losses for the principal due to asymmetric information (equivalently, of future information rents) through two novel channels. The first one is by changing the distribution of future types, for given future losses. The second one is through its effect on the impulse responses of future types to current ones, and thereby through its direct effect on future losses for given distribution of future types.

In most cases of interest, future welfare losses of the principal are increasing in future types. Therefore, through the first channel, LBD can contribute to higher expected future losses by shifting the distribution of future types towards productivity levels for which future information rents are higher. This channel thus contributes to higher wedges at present to mitigate the magnified cost of future rents brought in by LBD.

This effect can be amplified or weakened by the second channel. Holding fixed the distribution of future types, future losses may be either increasing or decreasing in current output, depending on whether the impulse responses of future types to current ones are increasing or decreasing in current output. In turn, this depends on the interaction between current skills and current output in the

[1] For instance, Willis (1986) and Altug and Miller (1998) show that past work experience has a significant effect on current wage earnings. Moreover, Jacobson et al. (1993) show that displaced workers suffer sizable wage losses, while Topel (1991) shows how job tenure affects wage profiles.

determination of future productivity. When current skills and current output are complements in the determination of future skills, impulse responses are increasing in current output. In this case the second channel adds to the first one in contributing to higher wedges, whereas the opposite is true when current output and current skills are substitutes.

We show that, when sufficiently strong, these novel effects may lead to marginal distortions (wedges) decreasing over time. This is true even in the presence of an insurance motive, which is known to contribute, in the absence of LBD, to distortions increasing over time (see, e.g., Farhi and Werning, 2013, and the discussion therein). The intuition is straightforward: the impact of LBD on future skills is most pronounced when there are many periods ahead—at the extreme, the effects of LBD on wedges are absent on the last day at work.

We also show that, in contrast to common intuition, LBD can lead to an increase in the level of the wedges. In particular, we show that this is the case when the impulse responses of future types to current ones are independent of current output, which is true for example under the Cobb-Douglas skill-transition technology typically assumed in the literature (i.e., when future types are given by a multiplicative power function of past types, past output – or labor supply – and shocks).

Finally, we show that LBD may lead to an increase in wedges that is more pronounced at the top of the current-period type distribution than at the bottom. When such effect is strong enough, it can lead to progressivity in wedges (equivalently, to marginal tax rates increasing in current income), under the same conditions that would lead to regressivity in the absence of LBD. This is because the benefit of distorting downwards labor supply in the present period so as to economize on the future costs of incentives is stronger for higher types, given that these are the types that expect, on average, larger rents in future periods.

Another contribution of the paper is in fully characterizing the optimal dynamic contract, accommodating for non-transferable utility (i.e., agents' risk aversion), arbitrary forms of inequality aversion on the principal's side, and endogenous types. The derived formula for the wedges generalizes formulae that appear in the received literature on optimal income taxation. This enables us to clarify various results in the dynamic mechanism design and taxation literature and isolate the relative importance of various channels in shaping the dynamics of distortions under optimal contracts.

Related literature.

The closest work is Krause (2009), Stantcheva (2014), Best and Kleven (2013), Kapicka (2013b) and Kapicka and Neira (2014). Krause (2009) and Best and Kleven (2013) study LBD in a two-period model. In addition, the former paper studies an economy with only two productivity levels, whereas the latter paper restricts attention to taxes that depend only on age and current income, thus abstracting from the characterization of fully optimal tax codes. Stantcheva (2014), Kapicka (2013b), and Kapicka and Neira (2014) study the effects on optimal non-linear taxes of *direct* invest-

ments in human capital.[2] In all these works, past (labor/investment) decisions affect productivity in a deterministic manner, with productivity being positively related also to the agent's private information.[3] The key difference between our paper and this recent body of work is that we consider an economy in which LBD affects the worker's *future private information* (identified in our context with his stochastic talent/productivity), whereas the evolution of the worker's private information is exogenous in the aforementioned body of work. To the best of our knowledge, ours is the first paper to study the effects of LBD on optimal taxes when the worker's private information is drawn from a continuous distribution that depends endogenously of past output (or labor supply). Another key difference between our model and the models that assume direct investment in human capital (e.g., Stantcheva, 2014, Kapicka, 2013b, and Kapicka and Neira, 2014) is the following. In the latter class of models, the instruments the principal can use to manipulate the wedges (e.g., training) are independent of effort/labor supply. In our model, instead, the instruments are related. In particular, any manipulation of the supply of labor in the current period has implications for the evolution of the agent's productivity in future periods. Such differences have important implications for the properties of optimal taxes, as we show shortly.

Our paper is also related to the dynamic mechanism design literature. See, for instance, Pavan, Segal, and Toikka (2014), and Bergemann and Pavan (2015) for a discussion of recent developments of this literature[4], as well as the surveys by Golosov et al. (2006) and Kocherlakota (2010) for the application of this literature to optimal taxation. The key contribution of the present paper relative to the earlier dynamic taxation literature (aka new dynamic public finance) is in the endogeneity of the type process. As mentioned above, such endogeneity may lead to taxes that decline, on average, over time, under the same preference/technology specifications that are conductive to taxes that are increasing over time when the type process is exogenous (see, e.g., Farhi and Werning (2013), Golosov et al. (2011) and Kapicka (2013a)). The key contribution relative to Pavan, Segal, and Toikka (2014) is in relaxing the assumption that utility is quasilinear, thus introducing a preference for consumption-smoothing and insurance.[5] A second contribution is in deriving testable implications for the endogeneity of the type process on the dynamics of incentives —the analysis in Pavan, Segal, and Toikka (2014) accommodates such a possibility but does not investigate its implications for the

[2]Human capital is contractible in Kapicka and Neira (2014) and non-contractible in Kapicka (2013b). Stantcheva (2014) considers both the case of contractible and non-contractible human capital, but focuses primarily on the former case.

[3]With private information being allowed to be perfectly correlated across time in some of these papers.

[4]That paper provides a general treatment of incentive compatibility in dynamic settings. It extends previous work by Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Battaglini (2005), Eso and Szentes (2007), and Kapicka (2013), among others, by allowing for more general payoffs and stochastic processes and by identifying the role of impulse responses as the key driving force for the dynamics of optimal contracts.

[5]See also Garrett and Pavan (2015) for a model of dynamic contracting with non-transferable utility, but exogenous types.

dynamics of distortions under optimal contracts.

The organization of the paper is the following. Section 2 presents a very simple two-period economy to illustrate the key ideas. This simple economy is, however, flexible enough to illustrates how learning-by-doing can change the predictions about the dynamics and the progressivity of the wedges identified in the literature. Section 3 extends the analysis to a fairly general class of economies with an arbitrary number of periods. It uses a recursive approach to arrive at a decomposition of the wedges into three driving forces that, jointly, are responsible for the dynamics and the progressivity of the wedges. Section 4 relates the results to the dynamic mechanism design and to the taxation literature. Section 5 concludes. All proofs are relegated to an Appendix in Section 5.

2 Simple two-period taxation economy

To fix ideas, consider an optimal nonlinear income taxation environment.

Agents, productivity, and information. The economy is populated by a unit-mass continuum of agents of different productivities, each living for $T = 2$ periods. Each agent's productivity is independent of all other agents' productivities and evolves endogenously over time.

In each period $t = 1, 2$, each agent produces *income* $y_t \in Y_t = \mathbb{R}_+$ at a cost $\psi(y_t, \theta_t)$, where θ_t denotes the agent's *productivity* (equivalently, her skill) and is the agent's private information. The function $\psi(y_t, \theta_t)$ is assumed to be thrice differentiable, increasing, and convex in y_t. Consistently with most of the literature[6], we assume here that ψ takes the iso-elastic form

$$\psi(y_t, \theta_t) = \frac{1}{1+\phi}\left(\frac{y_t}{\theta_t}\right)^{1+\phi}$$

and then denote by

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

its partial and cross derivatives. Hence, under this specification, ϕ is the inverse Frish elasticity. Note that for positive income, $\psi_\theta < 0$ and $\psi_{y\theta} < 0$.

Productivity in period $t = 1$ is exogenous and drawn from an absolutely-continuous distribution F_1 with density f_1 strictly positive over $\Theta_1 = (\underline{\theta}_1, \bar{\theta}_1])$. Productivity in period $t = 2$, instead, is given by the function

$$\theta_2 = z_2(\theta_1, y_1, \varepsilon_2) = \theta_1 y_1^{\zeta} \varepsilon_2 \tag{1}$$

with ε_2 drawn from some distribution G with support $E \subset \mathbb{R}_+$, independently across agents and independently from θ_1. We assume that $\zeta \leq \phi/(1+\phi)$ in order to ensure a well-defined solution for the first-period output schedule. Then let

$$\Theta_2 = \{\theta_2 : \theta_2 = z_2(\theta_1, y_1, \varepsilon_2),\ (\theta_1, y_1, \varepsilon_2) \in \Theta_1 \times \mathbb{R}_+ \times E\}$$

[6]See, among others, Kapicka (2013), Farhi and Werning (2013), and Best and Kleven (2013).

denote the set of possible period-2 productivities, and $\Theta = \Theta_1 \times \Theta_2$. The dependence of period-2 productivity on period-1 income is what captures learning-by-doing (LBD). Note that, with income being an increasing function of effort and productivity, the above representation is flexible enough to encompass both the case where LBD is generated through past "effort/labor supply" as well as past "income" directly. Note, further, that, under this specification, the intensity of LBD is conveniently parametrized by the uni-dimensional parameter $\zeta \geq 0$; the case of no LBD corresponds to $\zeta = 0$, while higher ζ capture stronger LBD effects. Also note, under the technology in (1), for any $\theta \equiv (\theta_1, \theta_2) \in \Theta$ and any y_1, the impulse response of θ_2 to θ_1 (that is, the marginal effect of a variation in θ_1 on θ_2, holding fixed the shock $\varepsilon_2 = \theta_2/\theta_1 y_1^\zeta$ that, given θ_1 and y_1, is responsible for θ_2) is independent of y_1 and is given by

$$I_1^2(\theta, y_1) \equiv \left. \frac{\partial z_2(\theta_1, y_1, \varepsilon_2)}{\partial \theta_1} \right|_{\varepsilon_2 = \theta_2/\theta_1 y_1^\zeta} = \frac{\theta_2}{\theta_1} \tag{2}$$

The advantage of this specification is that it permits us to highlight the channel by which LBD affects wedges that is most specific to our analysis; namely, the effect of LBD on the distribution of future productivity.

Also note that, given the above specification, for any $(\theta_1, y_1) \in \Theta_1 \times \mathbb{R}_+$, θ_2 is distributed according to

$$F_2(\theta_2|\theta_1, y_1) = G\left(\frac{\theta_2}{\theta_1 y_1^\zeta}\right)$$

Preferences. Next, denote by $c_t \in \mathbb{R}_+$ the agent's period-t consumption and let δ be the common discount factor.[7] The *lifetime* utility of each agent is given by

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

Government's problem. The government's problem consists in designing an intertemporal tax code that maximizes an aggregator of the agents' intertemporal payoffs subject to the requirement that the fiscal deficit be no greater than an exogenous level G. As is typical in this literature, we will solve the dual problem associated to this primal, in which the Government maximizes intertemporal tax revenue subject to the constraint that the aggregator of the agents' intertemporal payoffs be greater than a given threshold.

Formally, this dual problem can be stated as follows. Let $\chi : \Theta \to \mathbb{R}^{2T}$, denote an *allocation rule*, with $\chi(\theta) = (y_t(\theta^t), c_t(\theta^t))$, $t = 1, 2$ where $\theta^1 \equiv \theta_1$ and $\theta^2 \equiv \theta$. Note that such rule describes the lifetime profile of income-consumption pairs for each agent with period-1 productivity θ_1 and with period-2 productivity θ_2. Then denote by $\lambda[\chi]$ the endogenous probability distribution over Θ

[7] In turn, δ is also equal to the inverse of the gross interest rate (See for instance Best and Kleven (2013), Kapicka (2013), Farhi and Werning (2013), Golosov et. al. (2014) and Stantcheva (2016)).

that is obtained by combining the period-1 exogenous distribution F_1 with the endogenous period-2 distribution F_2 that one obtains when $y_1 = y_1(\theta_1)$. Further, let $\lambda[\chi]|\theta_1$ denote the endogenous distribution over Θ that obtains under the rule χ, when the agent's initial productivity is θ_1. Finally, let

$$V_1(\theta_1) \equiv \mathbb{E}^{\lambda[\chi]|\theta_1}[U(\theta, \chi(\theta))] = \mathbb{E}^{\lambda[\chi]|\theta_1}\left[\sum_{t=1}^{T} \delta^{t-1}\left(c_t(\theta^t) - \psi(y_t(\theta^t), \theta_t)\right)\right]$$

denote the expected lifetime utility of each agent of initial productivity θ_1, under the allocation rule χ. Importantly, note that the dependence on χ is both through the policies $c_t(\cdot)$ and $y_t(\cdot)$, and through the dependence of the period-2 distribution F_2 on period-1 income $y_1(\theta_1)$, because of LBD.

In the simple economy here, the Government's problem then consists in maximizing expected tax revenue

$$R = \mathbb{E}^{\lambda[\chi]}\left[\sum_{t=1}^{T} \delta^{t-1}\left(y_t(\theta^t) - c_t(\theta^t)\right)\right]$$

subject to the constraint that

$$\min_{\theta_1} V_1(\theta_1) \geq \kappa$$

and the constraint that the rule χ is incentive compatible (that is, each agent finds it optimal to generate income over the life-cycle as specified by the policy $\chi(\cdot)$).

Remarks. Implicit in the above formulations are the following assumptions. (i) Consumption is a deterministic function of income, with the two linked by the tax code according to $c_t = y_t - \tau_t(y^t)$ where $y^1 = y_1$ and $y^2 = y \equiv (y_1, y_2)$. Note that the above specification accommodates for the possibility that the total taxes the agent pays in period 2 depend on her income in both periods y. In equilibrium, an agent with lifetime productivity history $\theta = (\theta_1, \theta_2)$ thus consumes $c_1(\theta_1) = y_1(\theta_1) - \tau_1(y_1(\theta_1))$, and $c_2(\theta) = y_2(\theta) - \tau_2(y_1(\theta_1), y_2(\theta))$. (ii) The Government commits in advance to the intertemporal tax code τ. (iii) The Government is constrained by the agent's intertemporal payoff be at least equal to κ. That is, the Government's primal problem consists in maximizing the expected lifetime utility of those agents whose initial productivity is the lowest (i.e., $\theta_1 = \underline{\theta_1}$) subject to an exogenous fiscal budget constraint. In other words, the Government's preferences for redistribution are consistent with a Rawlsian welfare objective. Because incentive-compatibility requires that the indirect utility $V_1(\theta_1)$ be non-decreasing, the latter assumption is equivalent to the assumption that the principal puts equal weight to the utility of each agent. (iv) The agent here is risk neutral when it comes to evaluating lotteries over intertemporal consumption.

We will relax most of these assumptions in the analysis of the general model of Section 3.

2.1 First Best

Suppose each agent's productivity is verifiable (that is, agents do not possess private information). In this simple economy, the Government would then require that each agent produces output/income

according to the policies

$$\psi_y(y_1(\theta_1), \theta_1) = 1 + \delta LD_1^{FB;\chi}(\theta),$$

and

$$\psi_y(y_2(\theta), \theta_2) = 1,$$

where

$$LD_1^{FB;\chi}(\theta) \equiv \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[y_2(\theta) - \psi(y_2(\theta), \theta_2) \right].$$

and $\lambda[\chi]|\theta_1, y_1$ denotes the endogenous distribution over Θ that obtains under the rule χ, when the agent's initial productivity is θ_1 and output is y_1. In both periods, the optimal choice of income is thus obtained by equalizing each agent's marginal disutility of labor with the marginal benefit of higher output. In the second period, the latter simply coincides with the extra resources that are made available when the agent works harder. In the first period, instead, the benefit of asking for higher output also takes into account the effect that the latter has on the distribution of period-2 productivity. Because the period-2 policies are set optimally, usual envelope arguments then imply that, in a first-best world, the extra benefit of asking for higher period-1 output originating in LBD is given by the function $LD_1^{FB;\chi}(\theta)$. Importantly, note that this function is computed holding fixed the period-2 policy $y_2(\cdot)$ as specified by the allocation rule χ. The expectation in the formula for $LD_1^{FB;\chi}(\theta)$ is with respect to the endogenous distribution over Θ under the rule χ, starting from period-1 productivity θ_1 and period-1 income $y_1 = y_1(\theta)$.

Given the optimal income choices, the optimal consumption policies are then given by any combination of $c_1(\cdot)$ and $c_2(\cdot)$ that satisfy

$$c_1(\theta_1) + \delta \mathbb{E}^{\lambda[\chi]|\theta_1} \left[c_2(\theta) \right] = \kappa + \psi(y_1(\theta_1), \theta_1) + \delta \mathbb{E}^{\lambda[\chi]|\theta_1} \left[\psi(y_2(\theta), \theta_2) \right].$$

The first-best allocations can be sustained, for example, through productivity-specific non-linear taxes of the form

$$\tau_1(y_1(\theta_1), \theta_1) = y_1(\theta 1) - \psi(y_1(\theta_1), \theta_1) - K \text{ and } \tau_2(y, \theta) = y_2 - \psi(y_2(\theta), \theta_2) - (\kappa - K)/\delta.$$

with K be a scalar. Because the agent is risk-neutral, the distribution of taxes over time, however, is irrelevant in this simple economy.

2.2 Second Best

When the agents' productivities are their own private information, the Government faces additional constraints to its ability to tax labor income. In particular, incentive compatibility requires that highly productive types receive informational rents necessary to dissuade them from mimicking less productive ones. Letting

$$V_2(\theta) \equiv c_2(\theta) - \psi(y_2(\theta), \theta_2)$$

denote the agent's period-2 continuation utility, we have that period-2 incentive-compatibility requires that each agent's continuation utility satisfy the familiar envelope formula

$$V_2(\theta_1, \theta_2) = V_2(\theta_1, \underline{\theta}_2) - \int_{\underline{\theta}_2}^{\theta_2} \psi_\theta(y_2(\theta_1, s), s) ds$$

along with the requirement that the income policy $y_2(\theta_1, \cdot)$ be non-decreasing in θ_2.

Period-1 incentive compatibility in turn requires that the expected lifetime utility of each agent satisfy an analogous envelope formula given by

$$V_1(\theta_1) = V_1(\underline{\theta}_1) - \int_{\underline{\theta}_1}^{\theta_1} \left\{ \psi_\theta(y_1(s), s) ds + \delta \mathbb{E}^{\lambda[\chi]|s} \left[I_1^2(\theta, y_1(\theta_1)) \psi_\theta(y_2(\theta), \theta_2) \right] \right\} ds \tag{3}$$

along with the requirement that, for any $\theta_1, \hat{\theta}_1 \in \Theta_1$, the following integral-monotonicity condition holds

$$\begin{aligned} & \int_{\hat{\theta}_1}^{\theta_1} \left\{ \psi_\theta(y_1(s), s) + \delta \mathbb{E}^{\lambda[\chi]|s} \left[I_1^2(\theta, y_1(\theta_1)) \psi_\theta(y_2(\theta), \theta_2) \right] \right\} \\ \leq & \int_{\hat{\theta}_1}^{\theta_1} \left\{ \psi_\theta(y_1(\hat{\theta}_1), s) + \delta \mathbb{E}^{\lambda[\chi]|s, y_1(\hat{\theta}_1)} \left[I_1^2(\theta, y_1(\hat{\theta}_1)) \psi_\theta(y_2(\theta), \theta_2) \right] \right\} ds \end{aligned} \tag{4}$$

Hereafter we follow the same first-order approach as in most of the dynamic mechanism design literature by ignoring the integral monotonicity constraint in (4) and verify that holds ex-post.

Using the fact that, when the agents are risk neutral, the tax revenues are equal to

$$R = \mathbb{E}^{\lambda[\chi]} \left[\sum_{t=1}^{T} \delta^{t-1} \left(y_t(\theta^t) - \psi(y_t(\theta^t), \theta_t) \right) - V_1(\theta_1) \right]$$

along with (3), we then have that, under asymmetric information, the government's objective can be expressed as

$$R = \mathbb{E}^{\lambda[\chi]} \left[\sum_{t=1}^{T} \delta^{t-1} \left(y_t(\theta^t) - \psi(y_t(\theta^t), \theta_t) + \frac{I_1^t(\theta, y_1(\theta_1))}{\eta_1(\theta_1)} \psi_\theta(y_t(\theta^t), \theta_t) \right) - V_1(\underline{\theta}_1) \right], \tag{5}$$

where $I_1^1(\theta, y_1(\theta_1)) \equiv 1$, all θ, and where

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

denotes the hazard rate of the period-1 distribution F_1. The Government's second-best policies thus maximize (5) subject to the constraint that $V_1^\chi(\underline{\theta}_1) \geq \kappa$. The Government's problem in a second-best world is thus similar to its problem in a first-best world, except for the fact that tax revenues are lower because of the informational rents that must be left to the agents to induce them to reveal their private information. Such rents are given by the "handicaps"

$$h_1(\theta_1, y_1(\theta_1)) \equiv -\frac{1}{\eta_1(\theta_1)} \psi_\theta(y_1(\theta_1), \theta_1) \text{ and } h_2(\theta, y(\theta)) \equiv -\frac{I_1^2(\theta, y_1(\theta_1))}{\eta_2(\theta_1)} \psi_\theta(y_2(\theta), \theta_2)$$

in the tax revenue function R, where $y(\theta) \equiv (y_1(\theta_1), y_2(\theta))$.

The second-best income policies are then chosen to trade off the marginal effects of higher output on current and future surplus, as in a first-best world, with the marginal effects that higher output has on current and future handicaps. Differentiating R with respect to $y_1(\theta_1)$ and $y_2(\theta)$ we then have that the optimal policies must satisfy the optimality conditions

$$\psi_y(y_2(\theta), \theta_2) - \frac{I_1^2(\theta, y_1(\theta_1))}{\eta_1(\theta_1)} \psi_{\theta y}(y_2(\theta), \theta_2) = 1 \tag{6}$$

and

$$\psi_y(y_1(\theta_1), \theta_1) - \frac{1}{\eta_1(\theta_1)} \psi_{\theta y}(y_1(\theta_1), \theta_1) - \delta \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\frac{I_1^2(\theta, y_1(\theta_1))}{\eta_1(\theta_1)} \psi_\theta(y_2(\theta), \theta_2) \right] = 1 + LD_1^{FB;\chi}(\theta), \tag{7}$$

where recall that

$$LD_1^{FB;\chi}(\theta) \equiv \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[y_2(\theta) - \psi(y_2(\theta), \theta_2) \right]$$

are the first-best effects of changing the endogenous distribution of θ_2. The above conditions pin down the optimal output schedules. The left-hand side in each of these conditions is the marginal cost of asking the agent for higher output in period t, whereas the right-hand side is the marginal benefit. Consider first (6). The marginal cost of asking for higher output in period-2 from an agent of productivity history θ has two parts. The first one is the marginal adjustment in the agent's consumption necessary to compensate him for the extra disutility of labor, $\psi_y(y_2(\theta), \theta_2)$. This part is standard and is the same as in the first-best benchmark. The interesting part is the second one. To understand this part, note that, under asymmetric information, when the principal asks for higher output from an agent with productivity $\theta = (\theta_1, \theta_2)$, she then needs to increase the consumption of all period-1 types whose period-1 productivity is higher than θ_1. The cost of such adjustment is higher the higher the inverse period-1 hazard rate $1/\eta_1(\theta_1)$ and the higher the intertemporal informational linkage between θ_1 and θ_2 as captured by the impulse response $I_1^2(\theta, y_1(\theta_1))$. The cost of such informational rent, as perceived from the perspective of period one is given by the period-2 handicap $h_2(\theta, y(\theta))$. Because the latter is increasing in y_2, at the optimum, the labor supply of an agent of productivity history θ is distorted downwards relative to its first-best level.

Next, consider the optimal choice of period-1 output, as determined by (7). The benefits of asking for higher y_1 naturally take into account the effect of changing the distribution of period-2 productivity coming from LBD. These benefits are determined by the same function $LD_1^{FB;\chi}(\theta)$ introduced above in the first-best case. However, the value of this function is now different because the period-2 policies y_2 are now distorted relative to the first-best, as explained above. As a result, the benefits of LBD are now reduced relative to their first-best counterparts. Furthermore, the costs of asking for higher period-1 output to an agent of period-1 productivity θ_1 are now augmented by the effect that such higher period-1 output has on the information rents that must be left to all

period-1 types above θ_1. The marginal costs of increasing such rents are given by the second and the third terms in the left-hand side in (7). The second term is the familiar one and coincide with the corresponding term in the optimality condition for period-2 output, except for the fact that the impulse response of θ_1 to itself, which is equal to $I_1^1(\theta, y_1(\theta_1)) = 1$, is typically higher than the impulse response $I_1^2(\theta, y_1(\theta 1))$ of θ_2 to θ_1, due to imperfect and declining persistence in productivity. The interesting novel effects due to LBD are captured by the third term in the left-hand side in (7). These are the novel effects that originate from the endogeneity of the agents' private information in future periods. In the presence of LBD, asking for a higher period-1 output has two effects on the expected period-2 handicap $h_2(\theta, y(\theta))$. The first effect is through the variation in the impulse responses. This effect, however, is absent under the technology specification of (1). The second effect is via the change in the distribution of θ_2, holding fixed the period-2 handicap function $h_2(\theta, y(\theta))$.

The novel effects due to LBD have important implications for the level and the progressivity of the labor wedges, as we show next. Let

$$W_1(\theta_1) \equiv 1 - \frac{\psi_y(y_1(\theta_1), \theta_1)}{1 + LD_1^{FB;\chi}(\theta)} \text{ and } W_2(\theta) \equiv 1 - \psi_y(y_2(\theta), \theta_2).$$

Recall that efficiency requires that the marginal disutility of extra period-t output be equalized to the marginal benefit, where, in the first period the latter takes into account also the effect of higher period-1 output on the the distribution of future total surplus, as captured by the term $LD_1^{FB;\chi}(\theta)$. The period-t wedge W_t is then defined to be the discrepancy between the ratio of marginal cost and marginal benefit of higher period-t output at the efficient allocation and the corresponding ratio at the proposed second-best allocation. Importantly, in period-1, such discrepancy is computed holding fixed the period-2 policies, so as to highlight the part of the inefficiency that pertains to the period-1 decisions. It is easy to see that the the wedges W_t are directly related to the period-t marginal tax rates; that is, the sensitivity of current taxes to current income, holding fixed past incomes and all future tax schedules (with the latter allowed to depend on the entire history of reported incomes).

As is customary in the taxation literature, instead of studying the behavior of W_t, we will consider the monotone transformation

$$\hat{W}_t(\theta^t) \equiv \frac{W_t(\theta^t)}{1 - W_t(\theta^t)}$$

which captures the wedge relative to the ratio of the agent's marginal disutility and total marginal benefit. We will be referring to $\hat{W}_t$ as the *relative wedge*.[8] Using the optimality conditions above, we then have that

$$\hat{W}_t(\theta^t) = \hat{W}_t^{RN}(\theta^t) + \Omega_t(\theta^t) \tag{8}$$

where

$$\hat{W}_t^{RN}(\theta^t) \equiv -\frac{I_1^t(\theta, y_1(\theta_1))\psi_{\theta y}(y_t(\theta^t), \theta_t)}{\eta_1(\theta_1)\psi_y(y_t(\theta), \theta_t)}$$

[8] For example, when $t = 1$, $\hat{W}_1(\theta_1) = W_1(\theta_1) / \left[\frac{\psi_y(y_1(\theta_1),\theta_1)}{1+LD_1^{FB;\chi}(\theta_1)}\right]$.

are the relative period-t wedges when the agent is risk neutral (as assumed here) and there are no LBD effects, and where the term

$$\Omega_1(\theta_1) = -\delta \frac{\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)}\left[\frac{I_1^2(\theta, y_1(\theta_1))}{\eta_1(\theta_1)}\psi_\theta(y_2(\theta), \theta_2)\right]}{\psi_y(y_1(\theta_1), \theta_1)}$$

summarizes the novel effects due to LBD (clearly such effects are absent in the second period and hence $\Omega_2(\theta) = 0$ all θ). As we will show in the next section, the decomposition of the wedges in (8) extends to richer economies, modulo an amplification term that combines the agents' risk aversion with preferences for redistribution other than the Rawlsian ones considered here. One can also verify that, under the technology and preference structure of this simple economy, the above formulas reduce to

$$\hat{W}_t^{RN}(\theta^t) = \frac{1+\phi}{\theta_1 \eta_1(\theta_1)},$$

and

$$\Omega_1(\theta_1) = \frac{\delta}{\psi_y(y_1(\theta_1), \theta_1)} \hat{W}_1^{RN}(\theta_1) \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)}\left[\psi(y_2(\theta), \theta_2)\right] \tag{9}$$

Note that, because of the iso-elastic disutility of labor and the Rawlsian preferences of the planner, the first-period wedges in the absence of LBD are independent of the allocation itself. Moreover, under the specification in (1), the impulse responses are invariant in the period-1 output. As a result, for any productivity history $\theta = (\theta_1, \theta_2)$, the period-2 wedges are invariant in the intensity of the LBD effects and coincide with the period-1 wedges in the absence of LBD. These properties facilitate the discussion of the effects of LBD but are not essential to the results below.

The wedges $\hat{W}_t^{RN}(\theta^t)$ are the well-known period-t wedges under transferable utility, iso-elastic disutility of labor, and Rawlsian objective, in the absence of LBD (see e.g., Diamond (1998)). It is easy to see that $\hat{W}_t^{RN}(\theta^t)$ are nonincreasing in productivity if, and only if, $\theta_1 \eta_1(\theta_1)$ is nondecreasing in θ_1, as typically assumed in the taxation literature. In the absence of LBD, the theory thus predicts marginal taxes that are constant over time and nonincreasing in earnings (when earnings are nondecreasing in productivity). We now show that both these predictions can be overturned if the planner designs optimal wedges under LBD.

The effect of LBD is captured by the novel term $\Omega_t(\theta^t)$. We immediately have that the introduction of LBD in this simple model has an effect only on first-period wedges. In more detail, it will lead to higher first-period wedges if and only if $\Omega_1(\theta_1) > 0$. In addition, the first-period wedge will be less regressive under LBD when $\Omega_1(\theta_1)$ is increasing. The term $\Omega_1(\theta_1)$ measures the (discounted) effect of LBD on the expected future welfare loss due to the information problem faced by the planner. In particular, it is the change in the expected period-2 handicap (measured in units of the first-period marginal disutility from labor) after a marginal increase in first-period output. Therefore, if it is positive (rest. negative), then the principal wants to set the period-1 wedges above (rest. below) the level she would have selected in the absence of LBD to depress labor and thereby reduce her

expected welfare loss due to asymmetric information. In fact, if higher first-period types expect, on average, larger rents in future, then the effect of LBD on the period-2 handicap (i.e. $\Omega_1(\theta_1)$) is increasing in the first-period type. Consequently, in this case, the planner wants also to implement a less "regressive" wedge than the one in the absence of LBD.

Both these are in fact possible under conditions that have been deployed in the existing literature. Suppose, for instance, Pareto shocks, which can approximate relatively well the upper tail of the skill distribution (see, for instance, Diamond (1998) and Kapicka (2013)). In this case, the wedges in the absence of LBD are constant across time *and* skills. However, $\Omega_1(\theta_1)$ is strictly positive and increasing, leading thus to decreasing wedges over time and a first-period wedge which is increasing (referred below as "progressive"). This is illustrated in Figure 1 for the case of the Pareto distribution used in the computations of Kapicka (2013), and income levels evaluated at their optimal level (i.e. under rule χ). As the figure shows, stronger LBD effects (here parametrized by a higher level of the parameter ζ) are responsible for higher period-1 wedges and for more progressivity at all income percentiles, but in particular at high percentiles.[9] This result is in fact robust to allowing for more plausible distributions of skill-shocks as we show in our next proposition. For instance, Figure 2 illustrates the first-period wedge for the case of a skills distribution with Pareto-tail parameter $\lambda = 5$, which echoes one the distribution highlighted in Diamond (1998).

To state our next proposition we need to introduce some definitions and notation.

Definition 1 *The period-1 wedge is more progressive over the interval (θ_1', θ_1'') in the presence of LBD than in its absence if and only if $\hat{W}_1(\cdot)$ is strictly steeper than $\hat{W}_1^{RN}(\cdot)$ over $(\theta_1', \theta_1'') \subset \Theta_1$. The period-1 wedge under LBD, $\hat{W}_1(\cdot)$, is more progressive than the period-1 wedge in the absence of LBD, $\hat{W}_1^{RN}(\cdot)$, if and only if $\hat{W}_1(\cdot)$ is weakly steeper than $\hat{W}_1^{RN}(\cdot)$ over the entire support Θ_1 of the period-1 distribution, and strictly steeper over a subset $(\theta_1', \theta_1'') \subset \Theta_1$.*

Using (8) for $t = 1$, we have that the period-1 wedge is more progressive over the interval (θ_1', θ_1'') in the presence of LBD than in its absence if and only if the function $\Omega_1(\theta_1)$ is strictly increasing over (θ_1', θ_1''). The proposition below identifies necessary and sufficient conditions for this to be the case over the entire support Θ_1.

Proposition 1 *Consider the two-period economy described above. (i) For all $\theta_1 \in \Theta_1$, $\hat{W}_1(\theta_1) > \hat{W}_1^{RN}(\theta_1)$; (ii) For all $\theta = (\theta_1, \theta_2)$, $\hat{W}_1(\theta_1) > \hat{W}_2(\theta)$; (iii) There exists a function $\Gamma : \Theta_1 \to \mathbb{R}$ such that the period-1 wedge $\hat{W}_1(\theta_1)$ is more progressive than the period-1 wedge $\hat{W}_1^{RN}(\theta_1)$ over any*

[9] Recall that the non-normalized wedge (or effective marginal tax rate) is given by $W_1(\theta_1) = \frac{\hat{W}_1(\theta_1)}{1+\hat{W}_1(\theta_1)}$. The parameter $eta = 0$ in the figure indicates that the result is for the case of a risk-neutral agent. Also note that the figure assumes $\phi = 2$, i.e., a Frisch elasticity of 0.5, as in Farhi and Werning (2013), Kapicka (2013), and Stantcheva (2014). Finally, the parameter $\rho = 1$ in the caption of the figure indicates that we have assumed here (as in the entire analysis in this section) a skill-persistence parameter of 1, as in Farhi and Werning (2013), Kapicka (2013), and Stantcheva (2014).

arbitrary interval $(\theta_1', \theta_1'') \subset \Theta_1$ *if and only if*

$$\Gamma(\theta_1) \geq 0 \ \forall \theta_1 \in (\theta_1', \theta_1''). \tag{10}$$

(iv) Suppose F_1 *is Pareto (in which case,* $\theta_1 \eta_1(\theta_1) = \lambda$ *for all* θ_1*). (a) In the absence of LBD,* $\hat{W}_1^{RN}(\theta_1) = (1+\phi)/\lambda$*, for all* θ_1*. (b) In the presence of LBD,* $\hat{W}_1(\theta_1)$ *is strictly increasing in* θ_1 *over the entire support* $\Theta_1 = \mathbb{R}_+$*. (v) Assume* $\hat{W}_1^{RN}(\theta_1)$ *is nonincreasing. The solution to the relaxed program also solves the full program.*

Hence, LBD contributes to higher period-1 wedges across all productivity levels. Furthermore, in the presence of LBD, period-1 wedges are higher than period-2 wedges across all histories $\theta = (\theta_1, \theta_2)$. Finally, LBD contributes to a higher progressivity of the period-1 wedges over any arbitrary interval of period-1 productivities $(\theta_1', \theta_1'') \subset \Theta_1$ if and only if Condition (10) is satisfied, which is the case, for any $(\theta_1', \theta_1'') \subset \Theta_1$, for example when θ_1 is drawn from a Pareto distribution.

More specifically, as we show in the proof of the proposition in the Appendix, the first-period wedge is given by

$$\hat{W}_1(\theta_1) = \hat{W}_1^{RN}(\theta_1) \left\{ 1 + \frac{\frac{1}{\phi} \theta_1^{\frac{1+\phi}{\phi}} y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{1}{\delta\zeta\bar{\varepsilon}(\phi)} \left[1 + \hat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}} + \theta_1^{\frac{1+\phi}{\phi}} y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}} \right\} \tag{11}$$

with $y_1(\theta_1)$ uniquely defined by

$$\left[1 + \hat{W}_1^{RN}(\theta_1)\right]^{-1} \theta_1^{1+\phi} + \delta\zeta\bar{\varepsilon}(\phi)\theta_1^{\frac{(1+\phi)^2}{\phi}} \left[1 + \hat{W}_1^{RN}(\theta_1)\right]^{-\frac{1+\phi}{\phi}} y_1^{\frac{\zeta(1+\phi)-\phi}{\phi}} - y_1^{\phi} = 0$$

with $\bar{\varepsilon}(\phi) \equiv \mathbb{E}\left[\varepsilon_2^{\frac{1+\phi}{\phi}}\right]$.

To understand the results in the proposition, recall from the discussion above that, under the specification in (1), LBD affects the expected net present value (NPV) of future welfare losses (equivalently, expected future handicaps) only through its effect on the distribution of period-2 productivity, for given period-2 handicaps (which, under the assumed specification, are given by $h_2(\theta, y(\theta)) = \hat{W}_1^{RN}(\theta_1)\psi(y_2(\theta), \theta_2)$ and are invariant in y_1 and increasing in θ_2). As we show in the proof of the proposition in the Appendix, the marginal effect of higher period-1 output on the expectation of the period-2 handicaps at history θ_1 and at output level $y_1(\theta_1)$ is proportional to the term

$$\theta_1^{\frac{1+\phi}{\phi}} y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}$$

that appears in the formula of the period-1 wedges in (11). As anticipated earlier in this subsection, when the period-2 handicap is increasing in the period-2 productivity, as is the case here, LBD, by shifting the distribution of period-2 productivity towards higher levels, contributes to a higher expectation of the period-2 handicap (equivalently, to higher expected future welfare losses). As

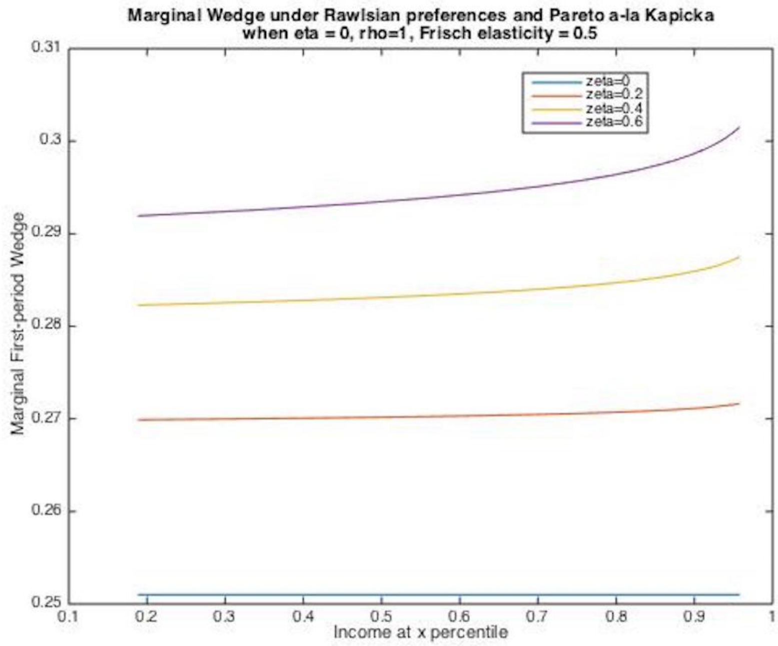

Figure 1: The Risk-Neutral Rawlsian Pareto case

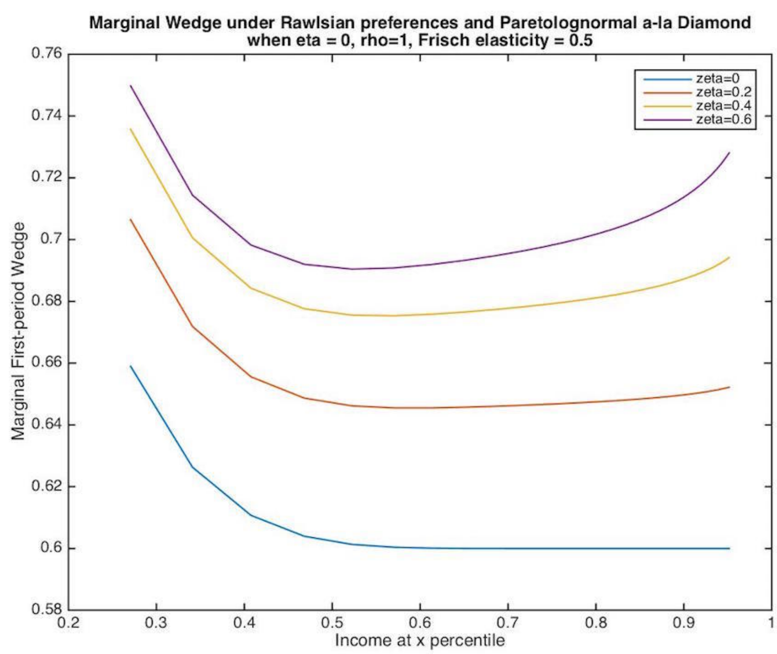

Figure 2: The Risk-Neutral Rawlsian Paretolognormal case

a result, at the optimum, the principal raises the period-1 wedges above the level she would have selected in the absence of LBD. This explains why, all other things equal, period-1 wedges are higher in an economy with LBD than in one without it.

Next observe that the effects of LBD on expected future welfare losses are typically most pronounced the more periods are ahead. In fact, these effects are completely absent in the last period at which the agents work. Hence, the positive effect of LBD on wedges should be expected to decline over time. All other things equal, this in turn implies that LBD can also revert the dynamics of the wedges. In the simple economy under examination here, wedges are constant over time in the absence of LBD, whereas they are decreasing over time in the presence of LBD.

Lastly, consider the effects of LBD on the progressivity of the wedges. By inspecting the formula for the period-1 wedges in (11), one can immediately see that LBD always contributes to a higher progressivity of the period-1 wedges when the term in the curly brackets in (11) is non-decreasing in θ_1, which is always the case when θ_1 is drawn from a Pareto distribution, for in this case $\hat{W}_1^{RN}(\theta_1)$ is constant and $y_1(\theta_1)$ is increasing in θ_1. More generally, Condition (10) in the proposition provides necessary and sufficient conditions for LBD to contribute to a higher progressivity (alternatively, to a lower progressivity of the wedges), at each productivity level θ_1.

As anticipated earlier in this subsection, the reason why LBD can contribute to a higher progressivity is that the benefit of distorting downwards labor supply so as to economize in the future cost of incentives (equivalently, on expected future information rents) by shifting the distribution of future types towards lower productivity levels (for which the cost of incentives is lower) can be stronger for higher types than for lower ones. This is because, typically, higher types expect, on average, larger rents in future periods. This is always the case when θ_1 is drawn from a Pareto distribution. More generally, in the economy under examination here, this is the case for any open interval (θ_1', θ_1'') at which $\eta_1(\theta_1)\theta_1$ is constant, as, for example, in the upper tail of a Pareto-Lognormal distribution, as illustrated in Figure 2 above.[10]

3 General case

3.1 The Environment

We now show how the results in the previous section extend to a broader class of economies, as well as to contractual problems other than optimal taxation. For this purpose, hereafter we will refer to the party designing the contractual relationship as the principal and to the informed party as the

[10] A Pareto-Lognormal distribution ε has support $(0, \infty)$, density $\frac{\lambda}{\varepsilon^{\lambda+1}} \exp(\lambda M + \lambda^2 \frac{\sigma^2}{2})\Phi(\frac{log(\varepsilon)-M-\lambda\sigma^2}{\sigma})$, and cdf $\Phi(\frac{log(\varepsilon)-M}{\sigma}) - \frac{1}{\varepsilon^{\lambda}} \exp(\lambda M + \lambda^2 \frac{\sigma^2}{2})\Phi(\frac{log(\varepsilon)-M-\lambda\sigma^2}{\sigma})$, where $\Phi(\cdot)$ is the c.d.f. of the standard Normal distribution. Such a distribution is similar to the Lognormal for small values but has a Pareto right-tail. Its mean for $\lambda > 1$ is the product of means of $Log - N(M, \sigma)$ and $Pareto(1, \lambda)$ distributions, i.e., $\frac{\lambda}{\lambda-1} \exp(M + \frac{\sigma^2}{2})$.

agent.

The principal ("she") wishes to determine contractually the allocations that govern her relationship with the agent ("he"). The relationship lasts for T periods, where T could be finite or infinite. The agent possesses private information in each period which is pertinent to the relationship. The principal offers a contract that has to respect the incentives of the agent (i.e., be incentive-compatible) and a certain redistribution constraint (described in detail below).

The agent's type could represent the agent's productivity, and the allocations the profile of type-dependent consumption and earnings. This profile then implies a non-linear earnings/income tax schedule for an economy with a unit-mass of agents of different productivities and types i.i.d. across individuals (but correlated over time for each individual). Alternatively, the agent's type could be a taste parameter and the allocation the profile of type-dependent consumption and efficiency units of labor. Such problems have been studied in the dynamic public finance (and insurance provision) literature. Finally, the allocation could be the profile of type-dependent consumption and output, with the profile implying the agent's compensation for producing the contracted output (which is valued by the principal). Such problems have been studied in the Managerial Compensation literature.[11] To fix ideas, hereafter, we will use the terminology akin to the dynamic public finance literature (DPF).

We start with some preliminary notation. Subscript t denotes period t, while superscript t denotes an history up to and including period t; that is, for a variable a, $a^t = \{a_1, .., a_t\}$ where $t = 1, 2, ..., T$. Furthermore, for any $j \geq 0$, $a_t^{t+j} = \{a_t, ..., a_{t+j}\}$, whereas $a_t^{t-j-1} = \{\varnothing\}$. For any set A, we let $A^0 = \{\varnothing\}$. Moreover, we 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 that takes value 1 when $a \in A$ and 0 otherwise.

In each period t, the agent produces *income/output* $y_t \in Y_t = \mathbb{R}_+$ at a cost $\psi(y_t, \theta_t)$, where θ_t (the agent's *productivity/skill*) is the agent's private information. 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 $\psi_{y\theta} < 0$, i.e. a more productive agent faces a lower marginal cost of production. As is common in such problems, we also assume that $\psi_{\theta y}(y, \theta)$ is nonincreasing in y in order to ensure a well-behaved second-best optimization problem for the principal. Productivity in period $t \geq 2$ is a function of the agent's productivity in the previous period, θ_{t-1}, the agent's income in the previous period, y_{t-1}, and some shock ε_t :

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

We will interpret the effect of past income on current productivity as learning-by-doing (LBD). Note

[11] The allocation could also be the profile of type-dependent price/transfer to a firm along with the specification of the output supplied by the firm. Such problems have been studied in the Procurement/Regulation literature.

that the above representation is general enough to encompass both the case where LBD is generated through past "effort/labor" as well as the case in which it is generated by past "produced output".

The productivity shock ε_t is distributed over $E_t \equiv (\underline{\varepsilon}_t, \overline{\varepsilon}_t) \subseteq \mathbb{R}$ according to some distribution G_t, with associated density g_t. We assume that the function $z_t(\cdot)$ is equi-Lipschitz continuous and differentiable, non-decreasing in its second argument, and increasing in its first and third arguments. Let $F_t(\theta_t|\theta_{t-1}, y_{t-1})$ denote the c.d.f. implied by G_t and the function z_t. We assume the period-1 productivity is exogenous and drawn from the absolutely-continuous distribution F_1 with density f_1. We the denote by $\Theta_t \equiv (\underline{\theta}_t, \overline{\theta}_t) \subseteq \mathbb{R}$ the support of the marginal distribution of the period$-t$ productivity. Note that our assumptions on $z_t(\cdot)$ imply the following First Order Stochastic Dominance properties:

$$F_{t+1,y}(\theta_{t+1}|\theta_t, y_t) \equiv \frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_t)}{\partial y_t} \leq 0 \text{ and } F_{t+1,\theta}(\theta_{t+1}|\theta_t, y_t) \equiv \frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_t)}{\partial \theta_t} \leq 0$$

We also introduce the following notation:

$$dF_{t+1,y}(\theta_{t+1}|\theta_t, y_t) \equiv \frac{\partial^2 F_{t+1}(\theta_{t+1}|\theta_t, y_t)}{\partial \theta_{t+1} \partial y_t} d\theta_{t+1} \text{ and } dF_{t+1,\theta}(\theta_{t+1}|\theta_t, y_t) \equiv \frac{\partial^2 F_{t+1}(\theta_{t+1}|\theta_t, y_t)}{\partial \theta_{t+1} \partial \theta_t} d\theta_{t+1}$$

Denote by $c_t \in C_t$ the agent's period-t consumption. Let $\Theta^t \equiv \prod_{\tau=1}^{t} \Theta_\tau$ denote the set of *feasible* period-t skill histories, with a generic element θ^t, and $\theta \equiv \theta^T$. Similarly for $y^t \in Y^t$ and $c^t \in \mathbb{R}^t$.

The *lifetime* utility of the principal is given by

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

whereas the *lifetime* utility of the agent is given by

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

with $v^i : \mathbb{R} \to \mathbb{R}$ increasing, weakly concave, and twice differentiable, $i = A, P$. U^i is the Bernoulli utility function of player i, which the player uses to evaluate lotteries over (θ, y, c). Player i is assumed to maximize the expectation of U^i. We denote by $U^i_\tau(\theta, y, c)$ the restriction of the Bernoulli utility from period $\tau \geq 2$ onwards (i.e. the player's continuation utility). That is, we let

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

and define $U^P_\tau(\theta, y, c)$ in a similar way.

We will confine our focus to environments in which output and consumption are strictly positive in each period. To this purpose we will assume the following Inada conditions: (a) $\lim_{c \to 0} v^{A\prime}(c) = \infty$

when $v^{A\prime\prime}(c) < 0$,[12] and (b) $\lim_{y_t \to 0} \left\{ v^{P\prime}(y_t) - \psi_y(y_t, \theta_t) \right\} > 0$ for all θ_t. To ease the notation, in what follows we will also drop the superscript A from the variables that refer to the agent's payoffs, unless there is risk of confusion.

We also 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, whereas the agent's skill θ_t is the agent's private information. We will refer to θ^t as to the agent's period-t type (history), while we will refer to θ as to the agent's type. The principal can commit to a contract that specifies for each period consumption c_t and output y_t, possibly as a function of messages sent in current and past periods. Without loss of optimality, we will restrict attention to (deterministic) direct revelation mechanisms that are incentive compatible, that is, that induce the agent to report truthfully in each period and for every history.[13]

An incentive-compatible contract gives rise to type-dependent allocations. We will refer to the mapping from types to (i) consumption, and (ii) income as an *allocation rule*. Given an allocation rule $\chi : \Theta \to \mathbb{R}^{2T}$, let, $\chi_t(\theta^t) = (y_t(\theta^t), c_t(\theta^t))$, $\chi^t(\theta^t) = (\chi_1(\theta_1), ..., \chi_t(\theta^t))$ and $\chi(\theta) = \chi^T(\theta^T)$. We then denote by $\lambda[\chi]|\theta^t$ the endogenous probability distribution over Θ that is obtained by combining the kernels F described above with the allocation rule χ starting from history θ^t, and denote by $\lambda[\chi]$ the ex-ante distribution under the rule χ. Similarly, denote by $\lambda[\chi]|\theta^t, y_t$ the endogenous probability distribution over Θ that is obtained by combining the kernels F described above with the allocation rule χ starting from history θ^t and period-t output y_t.

In addition, we will restrict attention to incentive-compatible contracts that satisfy

$$(1-r)V_1(\theta_1) + r \int q(V_1(\theta_1)) dF_1(\theta_1) \geq \kappa \tag{12}$$

where $r \in \{0, 1\}$ and where

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

is type θ_1's expected utility. The function $q(\cdot)$ is an increasing and (weakly) concave function that captures possible non-linear Pareto weights assigned by the principal to the agent's utility. With an abuse of terminology, we will refer to the above constrain as a "*redistribution constraint*". To understand this terminology, consider first the case of $r = 1$. In such an environment, in a managerial compensation or an optimal insurance problem, the case where $q(V) = V$ corresponds to an ex-ante participation constraint. In a taxation (dual) problem, it corresponds to the case of a planner with a "Utilitarian" objective. A strictly concave function q, instead, captures inequality aversion on the part of the planner (see, for instance, Farhi and Werning, 2013, and Best and Kleven, 2013), with stronger concavity corresponding to higher inequality aversion. The scenario of $r = 0$ is of

[12] As we will see, under risk neutrality, the allocation of an agent's consumption over time is indeterminate, and so assuming that consumption in every period is positive comes without loss of generality.

[13] The requirement of truth-telling and obedience for every history is without loss of generality in this Markov environment; see Pavan, Segal, and Toikka (2014).

special interest in a taxation problem, as it captures a planner with a "Rawlsian" objective. Because incentive compatibility requires that $V_1(\theta_1)$ be nondecreasing, in this case, de facto, the principal values only the payoff of the lowest period-1 type, in which case the above constraint is then equivalent to $V_1(\underline{\theta}_1) \geq \kappa$. This constraint is also the relevant (i.e., binding) interim participation constraint in managerial compensation problems in which the agent's participation must be guaranteed after the agent has learned his period-1 type. We will be refering to $q(V_1(\theta_1))$ as the principal's evaluation of the agent's period-1 expected payoff.

Hereafter, we will be interested in characterizing how LBD affects the properties of allocations sustained under *constrained-efficient* rules.

Definition 2 *A constrained-efficient allocation rule χ maximizes the principal's ex-ante expected payoff $\mathbb{E}^{\lambda[\chi]}[U^P(\theta, \chi(\theta))]$ among all allocation rules that are incentive compatible for the agent and satisfy the above redistribution constraint (12).*

3.2 The First-Best Benchmark

Under full information, the optimal allocation rule is obtained by maximizing $\mathbb{E}^{\lambda[\chi]}[U^P(\theta, \chi(\theta))]$ subject to the redistribution constraint (12). This constraint clearly binds at the optimum. Then, let $C \equiv v^{-1}$ and define the agent's period-t continuation payoff under the rule χ at any history θ^t recursively as follow

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

where

$$\Pi_{t+1}(\theta^t) \equiv \int V_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} \mid \theta_t, y_t(\theta^t)) \text{ if } t < T$$

and $\Pi_{t+1}(\theta^t) \equiv 0$ if $t = T < +\infty$, where we suppress the dependence of $V_\tau(\theta^\tau)$ on the rule χ to lighten the notation.

Note then that the solution to this problem can be conveniently described in recursive form as follows

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

subject to

$$\Pi_t(\theta^{t-1}) = \int V_t(\theta^t) dF_t(\theta_t|\theta_{t-1}, y_t(\theta^{t-1})), \text{ for } t > 1 \tag{13}$$

and

$$\kappa = (1-r)V_1(\theta_1) + r\int q(V_1(\theta_1)) dF_1(\theta_1) \tag{14}$$

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$-\tau$ value function of this recursive problem, given the period$-\tau$ 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 $t = 1, ..., T$, given the state, the choice of the current output schedule, $y_t(\theta^{t-1}, \cdot)$, along with the choice of the agent's current continuation payoff, $V_t(\theta^{t-1}, \cdot)$, and future "promised utility" $\Pi_{t+1}(\theta^{t-1}, \cdot)$, determine the current consumption schedule $c_t(\theta^{t-1}, \cdot)$. For all $\theta^t = (\theta^{t-1}, \theta_t)$, the latter is simply given by

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

The full-information benchmark, expressed in recursive form, is then obtained by choosing policies $(y_t(\theta^{t-1}, \cdot), V_t(\theta^{t-1}, \cdot), \Pi_{t+1}(\theta^{t-1}, \cdot))$ for each period t that maximize the principal's expected continuation payoff subject to the consistency (or, equivalently, promise keeping) constraint (13) that the mean of the agent's current period continuation payoff be equal to the utility level promised in the previous period. In addition, in period one, the above problem is augmented by the (binding) redistribution constraint (14). For $r = 0$ this constraint is

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

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

where we again suppress its dependence on χ to lighten the notation. Finally, let

$$LD_t^{FB;\chi}(\theta^t) \equiv \delta \int \left\{ V_{t+1}^P(\theta^{t+1}) + \frac{V_{t+1}(\theta^{t+1})}{v'(c_t(\theta^t))} \right\} dF_{t+1,y}(\theta_{t+1} \mid \theta_t, y_t(\theta^t)) \tag{15}$$

denote the effect of a marginal change in the period-t output on the expected sum of the principal's and 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, then let $V_{T+1}(\theta^{T+1}) = V_{T+1}^P(\theta^{T+1}) \equiv 0$. Importantly, note that the marginal effects captured by the function $LD_t^{FB;\chi}(\theta^t)$ are computed holding fixed the mapping from future productivity histories into allocations. We then have the following result.

Proposition 2 *The first-best allocation rule $\chi^* = (y^*, c^*)$ satisfies the following optimality conditions (at all interior points with $\lambda[\chi^*]$-probability one)*

$$v^{P'}(y_t^*(\theta^t)) + LD_t^{FB;\chi^*}(\theta^t) = \frac{\psi_y(y_t^*(\theta^t), \theta_t)}{v'(c_t^*(\theta^t))} \text{ all } t = 1, ..., T \tag{16}$$

$$v'(c_t^*(\theta^t)) = v'(c_{t+1}^*(\theta^t, \theta_{t+1})), \text{ any } t = 1, ..., T-1 \tag{17}$$

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 optimal output choice. The principal equalizes the marginal benefit of asking the agent to produce an extra unit of output (taking into account its effect on future payoffs from LBD which will be discussed shortly) 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 of labor is obtained by dividing the marginal disutility of labor 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 output. Under no LBD (or, equivalently, in the last period T, if the latter is finite), condition (16) 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 of LBD, recall that $F_{t+1,y}(\theta_{t+1}|\theta_t, y_t) \leq 0$, meaning that higher output at present shifts the future skills distribution in a first-order-stochastic-dominance way. In turn, this implies that, when the sum of the principal's and the agent's future payoffs (adjusted by the agent's marginal utility of consumption) is increasing in θ_{t+1}, then $LD_t^{FB;\chi^*}(\theta^t) \geq 0$. In such a case, LBD naturally induces the principal to ask for a higher output in period t compared to the level she would ask in the absence of LBD. This is because higher output implies higher ability in future periods, which in turn brings higher discounted expected net surplus.

The second and third conditions in turn describe the optimal choice of consumption. Under risk neutrality and an "utilitarian" objective, i.e., $v^{'} = q^{'} \equiv 1$, the dynamics of consumption is indeterminate. The reason is that under risk neutrality, the distribution of consumption over time is irrelevant; in the absence of inequality aversion, the allocation of utility across different types is also irrelevant. Under risk aversion, instead, optimality requires the equalization of the marginal utility of consumption between any two consecutive type histories θ^t and (θ^t, θ_{t+1}). Furthermore, with non-Rawlsian inequality aversion, optimality also requires equalizing across period-1 types the "weights" the principal assigns to the agent's period-1 marginal consumption, which are given by $q'(V_1^*(\theta_1))v^{'}(c_1^*(\theta_1))$. A Rawlsian planner, instead, would equalise the payoff of each first-period type with the threshold κ.

3.3 The Second Best

3.3.1 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, we have that $I_t^t(\theta^t, y^{t-1}) = 1$, and

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

where $e_{t+1}(\theta_{t+1}, \theta_t, y_t)$ is defined implicitly by[14]

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

Using the above definition, the impulse response function over non-consecutive periods can be defined inductively, for any $\tau > t$, as follows:

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

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,

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

and (b) for any differentiable and equi-Lipschitz continuous function $H(\theta^t, \cdot)$ of θ_{t+1}

$$\int H(\theta^{t+1})dF_{t+1,\theta}(\theta_{t+1}|\theta_t, y_t) = \int \frac{\partial H(\theta^{t+1})}{\partial\theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y^t)dF_{t+1}(\theta_{t+1}|\theta_t, y_t). \tag{18}$$

Next, recall that $V_t(\theta^t)$ denotes the agent's continuation payoff at the 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 and obedient strategy; note that, because χ is deterministic, truthful histories can be described entirely in terms of the 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. Finally, let

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

Likewise, 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 first mapping any period-t message into the message $\hat{\theta}_t$ and then determining allocations according to χ as if the period-t report was $\hat{\theta}_t$. Theorems 1 and 3 in Pavan et. al. (2014) establish that the allocation rule χ is incentive compatible *if and only if*, for all t, all θ^{t-1}, all $\theta_t, \hat{\theta}_t$, (a) the agent's equilibrium continuation payoff under χ is Lipschitz continuous with derivative given for almost all $\theta_t \in \Theta_t$ by (a)

$$\frac{\partial V_t(\theta^t)}{\partial\theta_t} = D_t^\chi(\theta^{t-1}, \theta_t) \tag{19}$$

and (b)

$$\int_{\hat{\theta}_t}^{\theta_t}\left[D_t^\chi(\theta^{t-1}, r) - D_t^{\chi\circ\hat{\theta}_t}(\theta^{t-1}, r)\right]dr \geq 0. \tag{20}$$

[14] Recall that here we are assuming that z_{t+1} is strictly increasing in ε_{t+1}. If this is not the case, the definition requires to take an average over all shocks that, along with θ_t and y_t lead to θ_{t+1}—see, Pavan, Segal, and Toikka (2014).

The marginal value $\partial V_t(\theta^t)/\partial\theta_t$ is related to the information rents that the principal must leave to the agent in the continuation to induce truthful and obedient behavior. Note that the assumption that each z_s is increasing in θ_{s-1} implies that $I_t^\tau > 0$, and hence $V_t(\theta^{t-1}, \theta_t)$ is increasing in θ_t.

Hereafter, we follow the so-called First-Order Approach (FOA) (also referred to, in the mechanism design literature, as the Myersonian approach) by considering a *relaxed program* in which the integral-monotonicity conditions (20) are dropped and checked ex-post. The discussion of the various channels through which LBD affects the dynamics of distortions in the subsequent sections should therefore be understood as being conditional on the First-Order Approach being valid.

3.3.2 Wedges

When the principal lacks information about the agent's type, she has an incentive to distort production in order to reduce the agent's information rents. These distortions give rise to the so-called (labor-) wedges, or marginal distortions, the dynamics of which is the focus of the paper.

Definition 3 *The period-t "wedge" or "marginal distortion" (equivalently, the "effective marginal income tax rate" in a taxation problem) at history θ^t, under the rule χ is given by*

$$W_t(\theta^t) \equiv 1 - \frac{\frac{\psi_y(y_t(\theta^t),\theta_t)}{v'(c_t(\theta^t))}}{\left[v^{P\prime}(y_t(\theta^t)) + LD_t^{FB;\chi}(\theta^t)\right]}.$$

To understand the formula, recall that efficiency requires that the marginal cost of extra period-t output (normalized by the agent's marginal utility of consumption) be equalized to the total weighted marginal benefit, where the latter takes into account both the flow marginal benefit to the principal, $v^{P\prime}$, as well as the effect of higher period-t output on the principal's and the agent's joint future surplus, as captured by the term $LD_t^{FB;\chi}(\theta^t)$ (see condition (16) above). The wedge is the discrepancy between the ratio of marginal cost and marginal benefit of higher period-t output at the efficient allocation and the corresponding ratio at the proposed allocation $(y_t(\theta^t), c_t(\theta^t))$. Importantly, such discrepancy is computed holding fixed the rule that determines future decisions so as to highlight the part of the inefficiency that pertains to the period-t decisions given past type θ^{t-1}. Note that, in a taxation problem, the wedge W_t is also equal to the period-t marginal tax rate; that is, the sensitivity of current taxes to current income, holding fixed past incomes and all future tax schedules (with the latter allowed to depend on the entire history of reported incomes).

As is usually the case in the taxation literature, instead of studying the behavior of $W_t(\theta^t)$, we will consider its monotone transformation

$$\hat{W}_t(\theta^t) \equiv \frac{W_t(\theta^t)}{1 - W_t(\theta^t)} = W_t(\theta^t) / \left[\frac{\frac{\psi_y(y_t(\theta^t),\theta_t)}{v'(c_t(\theta^t))}}{v^{P\prime}(y_t(\theta^t)) + LD_t^{FB;\chi}(\theta^t)}\right]$$

which captures the wedge relative to the ratio of the agent's adjusted marginal disutility and total weighted marginal benefit. We will be referring to $\hat{W}_t(\theta^t)$ as the *relative wedge*.

3.3.3 Wedges Dynamics under Optimal Contracts

To study the dynamics of the relative wedge under optimal contracts we need to analyze the principal's problem under asymmetric information. In general, this problem cannot be solved via simple pointwise maximization. As we show in the proof of the next proposition we can express this problem in a recursive formulation. This will, in turn, enable us to introduce a co-state variable that controls for the dynamics of the agent's continuation payoff, and then show how this permits us to arrive to optimality conditions that we then use to distill predictions about the dynamics of wedges under optimal contracts.

To state our next proposition, we need to introduce some definitions. First, for any $t > 1$, let $\mu_t(\theta^t)$ be the co-state variable associated with the local period-t envelope constraint, expressed in recursive form as (44). Next, let

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

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). Then let

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

denote the corresponding marginal value under the policy rule χ, where the values of $\Pi_{t+1}(\theta^t)$, $Z_{t+1}(\theta^t)$ and $y_t(\theta^t)$ are all determined using χ.

For any $t > 1$, also define the hazard rates of the period-t distributions by

$$\eta_t(\theta_t|\theta_{t-1}, y_{t-1}) \equiv \frac{f_t(\theta_t \mid \theta_{t-1}, y_{t-1})}{1 - F_t(\theta_t \mid \theta_{t-1}, y_{t-1})}$$

and, likewise, for period $t = 1$, let

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

To account for the fact that different types have different marginal utility of consumption, then let

$$p_t(\theta^{t-1}) \equiv \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta_t))} dF_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))$$

denote the period-t average inverse marginal utility of consumption, under the rule χ, and conditional on the type history θ^{t-1}. Similarly, define the average inverse marginal utility of consumption conditional on the agent's period-t type being above θ_t, and normalized by its unconditional mean value, by

$$\widehat{v}_t(\theta^t) \equiv \frac{\int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1},\tilde{\theta}_t))} \frac{dF_t(\tilde{\theta}_t|\theta_{t-1},y_{t-1}(\theta^{t-1}))}{1-F_t(\theta_t|\theta_{t-1},y_{t-1}(\theta^{t-1}))}}{p_t(\theta^{t-1})}.$$

It is nondecreasing in current productivity, when consumption is increasing in the current skills. Moreover, it equals to one when evaluated at the lowest productivity, and measures the amplification

of the cost of incentives for higher skill levels introduced by risk aversion. Analogously, define the average marginal evaluation by the principal of the agent's period-1 expected lifetime utility, relative to its unconditional mean value, under the rule χ, by

$$\widehat{q}_1(\theta_1, r) \equiv r \frac{\int_{\theta_1}^{\overline{\theta}_1} q'(V_1(\widetilde{\theta}_1)) \frac{dF_1(\tilde{\theta}_1)}{1-F_1(\theta_1)}}{\int_{\underline{\theta}_1}^{\overline{\theta}_1} q'(V_1(\tilde{\theta}_1)) dF_1(\tilde{\theta}_1)}.$$

This term that summarizes the effect of inequality-aversion here. It is constant under Rawlsian and Utilitarian preferences: equal to zero and one, respectively. For intermediate inequality-aversion, on the other hand, this term is nonincreasing in first-period productivity, when the agent's first-period value function is increasing in the current skills (which is true here). Further, it takes values between zero and one and is inversely related to the aversion to inequality on the part of the planner. Let also

$$\widehat{q}_t(\theta^t, r) \equiv 1 \text{ for any } t > 1$$

Further, denote

$$\widehat{\mu}_t(\theta^t) \equiv \frac{-\mu_t(\theta^t)}{\theta_t f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}$$

with $\widehat{\mu}_0(\theta^0) \equiv 0$, which represents the shadow cost of leaving information rents in period t to current type θ^t.

Finally, let

$$\epsilon_\theta^{\psi_y}(y_t, \theta_t) \equiv -\frac{\theta_t \psi_{y\theta}(y_t, \theta_t)}{\psi_y(y_t, \theta_t)}$$

denote the elasticity of the marginal disutility of labor with respect to the agent's productivity.

The following result summarizes all relevant optimality conditions:

Proposition 3 *The second-best allocation rule $\chi = (y, c)$ satisfies the following optimality conditions (at all interior points with $\lambda[\chi]$-probability one):*

$$\frac{1}{v'(c_t(\theta_t))} = \int_{\underline{\theta}_{t+1}}^{\overline{\theta}_{t+1}} \frac{1}{v'(c_{t+1}(\theta^t, \theta_{t+1}))} dF_{t+1}(\theta_{t+1} \mid \theta_t, y_t(\theta^{t-1})), \text{ for any } t < T \tag{21}$$

$$v^{P\prime}(y_t(\theta^t)) + LD_t^\chi(\theta^t) = \frac{\psi_y(y_t(\theta^t), \theta_t)}{v'(c_t(\theta^t))} \left[1 + \epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t) v'(c_t(\theta^t)) \widehat{\mu}_t(\theta^t)\right], \text{ for all } t \tag{22}$$

$$\widehat{\mu}_t(\theta^t) = \frac{p_t(\theta^{t-1})[\widehat{v}_t(\theta^t) - \widehat{q}_t(\theta^t, r)]}{\theta_t \eta_t(\theta_t | \theta_{t-1}, y_{t-1}(\theta^{t-1}))} + \left[\frac{\theta_{t-1}}{\theta_t} I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1}))\right] \widehat{\mu}_{t-1}(\theta^{t-1}), \text{ for all } t > 1. \tag{23}$$

Condition (21) is the familiar Rogerson-Euler condition. It appears also in the DPF literature and in the managerial compensation literature (see, e.g., Garrett and Pavan, 2015, for a validification of

this condition that does not use the first-order approach). It describes the optimal intertemporal allocation of consumption under any incentive-compatible contract; observe that it becomes redundant under risk neutrality.

Condition (22) describes the optimal output choice. It is the extension of (7) and (6) for the case of many periods, risk aversion on the part of the agent and a general labor disutility function.

Next, consider Condition (23). This describes the evolution of the shadow cost of information rents in the general model.[15] To understand this condition, let us start with the case of $t = 1$. Condition (23) for $t = 1$ pins down the weight the principal assigns to marginally increasing the expected lifetime utility of period-1 type θ_1. Using the aforementioned definitions, we have that, at the optimum, this weight is given by

$$-\mu_1(\theta_1) = \int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\tilde{\theta}_1))} dF_1(\tilde{\theta}_1) - r\pi_1 \int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\widetilde{\theta}_1))dF_1(\tilde{\theta}_1) \tag{24}$$

where

$$\pi_1 \equiv \frac{\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta_1))} dF_1(\theta_1)}{\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V(\theta_1))dF_1(\theta_1)} > 0$$

To see this, 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.[16] The first term in (24) 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 types above θ_1. However, there is also a benefit to a non-Rawlsian principal of providing such higher utility. This benefit stems from the effect that this has on the period-1 redistribution constraint. This benefit is captured by the second term in (24) for $r = 1$. In particular, the term

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

is the marginal value of increasing the expected lifetime utility of all period-1 types above θ_1, taking into account the (non-linear) Pareto weights the (non-Rawlsian) principal uses to evaluate the different utilities. The term π_1, on the other hand, is the shadow value of relaxing the redistribution constraint when $r = 1$, (43).[17] Combining the different terms and recalling the definition of $\hat{\mu}_t$, we thus have that, at the optimum, the initial value of the the co-state variable $\mu_1(\theta_1)$ must be given

[15] Observe that in the simple economy of the previous section this shadow cost is equal across periods and equal to $1/[\theta_1\eta_1(\theta_1)]$

[16] To see this, it suffices to use (44).

[17] To see this, note that reducing the value of the left-hand side of the redistribution constraint by one unit while ensuring incentive compatibility can be achieved by decreasing the lifetime utility of each period-1 type by an amount equal to the inverse of the denominator of π_1. Measuring this in terms of consumption and aggregating over all period-1 types leads directly to π_1.

by (24) with $r = 1$. Lastly note that, when $r = 0$, i.e. the (weaker version of the) redistribution constraint (43) becomes $V_1(\underline{\theta}_1) = \kappa$ (as in the managerial-compensation literature or in the taxation literature with a Rawlsian planner), then the principal does not attach any benefit in increasing the utility of any a period-1 type above $\underline{\theta}_1$. As a result, the second term in (24) vanishes.

Finally, consider Condition (23) for $t > 1$. This condition gives the optimal value of the period-t co-state variable, $\mu_t(\theta^t)$. As in the case of period 1, the co-state variable $\mu_t(\theta^t)$ is the shadow cost of increasing the expected lifetime utility of all period-t types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$, fixing the history θ^{t-1} (recall that, for incentive reasons, the principal needs to increase such types' continuation utility when she asks for a higher output to the period-t type (θ^{t-1}, θ_t)). The term

$$\frac{p_t(\theta^{t-1})\widehat{v}_t(\theta^t)}{\eta_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} = \frac{\int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1},\tilde{\theta}_t))} dF_t(\tilde{\theta}_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}$$

in the right-hand side of (23) for $t > 1$ is the direct marginal cost, in consumption terms, of increasing the compensation to all period-t types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$ (fixing θ^{t-1}), taking into account the heterogeneity in the marginal utility of consumption among such types (normalized by the conditional density, as usual). The second term

$$\frac{p_t(\theta^{t-1})}{\eta_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \tag{25}$$

in (23) for $t > 1$ takes into account the implications of such adjustment in terms of intertemporal consumption smoothing. To appreciate this term, note that, when the principal increases the continuation utility of all period-t types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$, she can then reduce the compensation she must provide in period $t-1$ to the period-$(t-1)$ type θ^{t-1}, while maintaining this latter type's value function unchanged. In particular, the amount the principal can deduct from the compensation to the period-$(t-1)$ type θ^{t-1} is equal to

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

Therefore, the total gain to the principal (normalized as usual by the conditional density of the period-t type) is equal to

$$\frac{1 - F_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))} v'(c_{t-1}(\theta^{t-1}))$$

which is equal to the term in (25) after using the definition of the period-t conditional hazard rate, η_t, and after using the Rogerson-Euler Condition

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

to express the inverse marginal utility of consumption in period $t-1$ in terms of the expected inverse marginal utility of consumption in period t.

The last term in the right-hand side of (23) for $t > 1$ is the extra cost of increasing the continuation utility of all period-t types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$ stemming from the fact that this increase also calls for an increase in the informational rents the principal must leave in previous periods because of the persistence of the agent's private information. To see this, use (44) and (46) jointly to observe that, when the principal increases the continuation payoff of all period-t types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$ by one, she then needs to increase the expected payoff of all period-$(t-1)$ types $(\theta^{t-2}, \tilde{\theta}_{t-1})$ with $\widetilde{\theta}_{t-1} \geq \theta_{t-1}$ by

$$\delta \int_{\theta_t}^{\bar{\theta}_t} \frac{\partial f_t(\widetilde{\theta}_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{\partial \theta_{t-1}} d\widetilde{\theta}_t = -F_{t,\theta}(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})).$$

Applying the same arguments above to period-$(t-1)$, we then have that the associated cost to the principal of this last adjustment is given by

$$-\frac{\mu_{t-1}(\theta^{t-1})}{f_{t-1}(\theta_{t-1} \mid \theta_{t-2}, y_{t-2}(\theta^{t-2}))} \left[-F_{t,\theta}(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))\right].$$

Normalizing, as usual, this cost by the conditional density of the period-t type and using the fact that

$$\frac{-F_{t,\theta}(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))} = I_{t-1}^t(\theta^t, y_{t-1}(\theta^{t-1}))$$

then yields the last term in (23) (recalling again the definition of $\hat{\mu}_t$).

Note that, (23) together with the definition $\widehat{\mu}_0(\theta^0) = 0$ provide a complete characterization of the dynamics of the period-t shadow cost of the informational rents (in terms of consumption) $\widehat{\mu}_t(\theta_t, \cdot)$. By substituting the shadow costs into the optimality conditions (22), and using also (21), one then obtains a complete description of the allocations sustained under optimal contracts. It is also worthwhile to notice that, when the principal has an Utilitarian objective (that is, when $r = 1$ and q' is identically equal to one) and the agent is risk neutral, then $\widehat{q}_1(\theta_1, 1) = 1$ and the solution to the above problem yields the same policies as in the first-best benchmark. This result thus extends to the environment with endogenous types under examination here a well-known result in optimal taxation theory that traces back to Mirrlees original work.

Equipped with the above results, we now turn to the dynamics of the relative wedges, and in particular how such dynamics are influenced by LBD. Fixing the rule χ, hereafter, we will refer to the term

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

as to the period-t "handicap," where

$$\epsilon_\theta^\psi(y_t, \theta_t) \equiv -\frac{\theta_t \psi_\theta(y_t, \theta_t)}{\psi(y_t, \theta_t)}.$$

Note that, while the wedges measure marginal distortions in the allocations, the handicaps measure the welfare losses due to asymmetric information, under the rule χ. Then let

$$\hat{W}_t^{RN}(\theta^t) \equiv \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_1(\theta_1)\theta_t} \epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)$$

$$\Omega_t(\theta^t) \equiv \frac{\frac{\partial}{\partial y_t} \mathbb{E}^{\lambda[\chi]|\theta^t, y_t(\theta^t)}[\sum_{\tau=t+1}^T \delta^{\tau-t} h_\tau(\theta^\tau, y^\tau(\theta^\tau))]}{\psi_y(y_t(\theta^t), \theta_t)}$$

and

$$RA_t(\theta^t) \equiv \frac{\widehat{\mu}_t(\theta^t) v'(c_t(\theta_t))}{\frac{\theta_1}{\theta_t} \frac{I_1^t(\theta^t, y_{t-1}(\theta^{t-1}))}{\theta_1 \eta(\theta_1)}}$$

As we explain below, the term $\hat{W}_t^{RN}$ is the period-t (relative) wedge the principal would choose at productivity history θ^t in the absence of LBD, when the agent is risk neutral and the principal has a Rawlsian objective, which is the benchmark case often considered in the literature. The term RA_t, in turn, is a correction that accounts for the agent's risk-aversion and/or to the principal's lower aversion to inequality. This term is equal to one when the agent is risk neutral and the principal is Rawlsian, and different than one otherwise. Finally, and most interesting to us, the term Ω_t summarizes the effects of LBD on the wedges when the agent is risk neutral and the principal has a Rawlsian objective, and is equal to zero when the type process is exogenous.

The following result provides a useful representation for the wedge which summarizes the interactions among all driving forces. This representation extends (8) in the general model.

Theorem 1 *At any period $t \geq 1$, with $\lambda[\chi]$-probability one, the relative wedge is given by*

$$\hat{W}_t(\theta^t) = RA_t(\theta^t) \left[\hat{W}_t^{RN}(\theta^t) + \Omega_t(\theta^t) \right].$$

To appreciate the above decomposition, note that, by using the law of motion of the shadow cost $\widehat{\mu}_t(\theta^t)$, (23), and the definition $\widehat{\mu}_0(\theta^0) = 0$, we have that, when the agent is risk neutral and the planner has a Rawlsian objective,

$$\widehat{\mu}_t(\theta^t) = \frac{\theta_1}{\theta_t} \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\theta_1 \eta(\theta_1)}$$

In this case the correction due to the agent's risk-aversion and/or to the planner's lower aversion to inequality is absent, i.e. $RA_t(\theta^t) = 1$. Furthermore, in the absence of LBD, $\Omega_t(\theta^t) = 0$, in which case the formula for the period-t wedge becomes[18]

$$\hat{W}_t(\theta^t) = \hat{W}_t^{RN}(\theta^t) \equiv \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_1(\theta_1)\theta_t} \epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)$$

[18]Note that, even if there is no LBD here, for simplicity, we continue to express the impulse responses as functions of past decisions.

In particular, note that, when $t = 1$, the formula for the wedge under risk neutrality and a Rawlsian objective reduces to the familiar one in the early public finance literature (e.g. Diamond (1998)):

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

For any period $t > 1$, instead, the formula for $\hat{W}_t^{RN}(\theta^t)$ is adjusted by taking into account the intertemporal informational linkage between type θ^t and type θ_1, as captured by the impulse response $I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))$. In this case, the dynamics of the wedges are entirely driven by the (exogenous) dynamics of the impulse responses, as discussed at length in the dynamic mechanism design literature with quasilinear payoffs (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 there is no LBD and that the agent is risk neutral. Again, using the law of motion for the shadow cost $\widehat{\mu}_t(\theta^t)$, (23), and the definition $\widehat{\mu}_0(\theta^0) = 0$, we have that

$$\widehat{\mu}_t(\theta^t) = [1 - \widehat{q}_1(\theta_1, 1)]\, \frac{\theta_1}{\theta^t} \frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\theta_1 \eta(\theta_1)}$$

in which case the correction due to the planner's lower aversion to inequality is equal to $RA_t(\theta^t) = 1 - \widehat{q}_1(\theta_1, 1)$ in all periods. The wedges are then equal to

$$\hat{W}_t(\theta^t) = RA_t(\theta^t)\hat{W}_t^{RN}(\theta^t) = \frac{[1 - \widehat{q}_1(\theta_1, 1)]\, I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_1(\theta_1)\theta_t}\epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)$$

and their dynamics continue to be driven by the exogenous dynamics of the impulse responses. Furthermore, when $t = 1$, the above formula for the wedge under risk neutrality reduces to the familiar one in Diamond (1998).

Next, consider the case in which the agent is risk averse, but continue to assume no LBD. Once again, by (23), we have that

$$\widehat{\mu}_t(\theta^t) = \sum_{\tau=2}^{t} \left[\frac{\theta_\tau}{\theta_t} I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1}))\right] \frac{p_\tau(\theta^{\tau-1})\,[\widehat{v}_\tau(\theta^\tau) - 1]}{\theta_\tau \eta_\tau(\theta_\tau|\theta_{\tau-1}, y_{\tau-1}(\theta^{\tau-1}))} - \left[\frac{\theta_1}{\theta_t} I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))\right] \widehat{\mu}_1(\theta_1)$$

The last term, and hence the correction $RA_t(\theta^t)$, depends on the inequality aversion of the principal, in a similar manner to the one discussed above. A special case of interest for the new dynamic public finance literature is that of a planner with an Utilitarian objective (e.g., Farhi and Werning, 2013, and Kapicka, 2013). In this case, (23) for $t = 1$ combined with $\widehat{\mu}_0(\theta^0) = 0$ becomes

$$\widehat{\mu}_1(\theta_1) = \frac{p_1\,[\widehat{v}_1(\theta_1) - 1]}{\theta_1 \eta_1(\theta_1)}$$

so that the correction due to the joint combination of the agent's risk aversion with the lower (here zero) planner's aversion to inequality becomes

$$RA_t(\theta^t) = v'(c_t(\theta^t)) \sum_{\tau=1}^{t} p_\tau(\theta^{\tau-1})\,[\widehat{v}_\tau(\theta^\tau) - 1]\, \frac{I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_\tau(\theta_\tau|\theta_{\tau-1}, y_{\tau-1}(\theta^{\tau-1}))} \left[\frac{I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_1(\theta_1)}\right]^{-1} \tag{26}$$

which is in general different than 1. In this case, the relative wedge becomes

$$\hat{W}_t(\theta^t) \equiv v'(c_t(\theta^t)) \sum_{\tau=1}^{t} p_\tau(\theta^{\tau-1}) \left[\widehat{v}_\tau(\theta^\tau) - 1\right] \frac{I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1}))}{\eta_\tau(\theta_\tau|\theta_{\tau-1}, y_{\tau-1}(\theta^{\tau-1}))} \frac{\epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)}{\theta_t}.$$

Another special case of interest with reference to the dynamic mechanism design (and managerial compensation) literature is that of a principal with Rawlsian preferences. In this case,

$$\widehat{\mu}_1(\theta_1) = \frac{p_1 \widehat{v}_1(\theta_1)}{\theta_1 \eta_1(\theta_1)}$$

and thereby

$$\hat{W}_t(\theta^t) \equiv$$

$$v'(c_t(\theta^t)) \left[\sum_{\tau=1}^{t} I_\tau^t(\theta^t, y^{t-1}(\theta^{t-1})) \frac{p_\tau(\theta^{\tau-1}) \left[\widehat{v}_\tau(\theta^\tau) - 1\right]}{\eta_\tau(\theta_\tau|\theta_{\tau-1}, y_{\tau-1}(\theta^{\tau-1}))} + I_1^t(\theta^t, y^{t-1}(\theta^{t-1})) \frac{p_1}{\eta_1(\theta_1)} \right] \frac{\epsilon_\theta^{\psi_y}(y_t(\theta^t), \theta_t)}{\theta_t}.$$

Interestingly, note that, contrary to the quasi-linear case examined in most of the dynamic mechanism design literature, the dynamics of the wedges are now driven by all (current and) past impulse responses.

In the presence of LBD, the terms $\hat{W}_t^{RN}(\theta^t)$ and $RA_t(\theta^t)$ 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 $\hat{W}_t^{RN}(\theta^t)$ and $RA_t(\theta^t)$ should now be interpreted as the ones corresponding to the process induced by the rule χ. The novel effects of LBD on the wedges are thus the ones captured by the terms $\Omega_t(\theta^t)$ and not the mere adjustment of the levels of $\hat{W}_t^{RN}(\theta^t)$ and $RA_t(\theta^t)$ due to the endogeneity in the process. The implications of the novel term $\Omega_t(\theta^t)$ for the dynamics of the wedges are discussed next.

Observe that the expectational term at the numerator of $\Omega_t(\theta^t)$ is the expected discounted sum of all future handicaps, with the latter capturing the welfare losses due to asymmetric information, in a world in which the agents are risk neutral and the planner is Rawlsian. In other words, the expectation of the net present value of future handicaps measures the expected social cost of leaving the agents information rents.

As mentioned already, LBD affects the expectation of future handicaps through two channels. The first one is by shifting (in the sense of first-order stochastic dominance) the distribution of future types towards higher levels, holding fixed all future handicaps. Under the specifications typically considered in the literature, the handicaps are increasing in the agents' types. When this is the case, by shifting the distribution towards levels for which the handicaps are higher, LBD contributes to *higher* expected future losses. Other things equal, this channel thus contributes to higher wedges in the present period aimed at curtailing the enhanced expected value of future losses brought in by LBD.

The second channel operates via 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 given distribution of future types. 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. When current skills interact with current output in a complementary manner in shaping future skills, impulse responses are increasing in current output. In this case, this second channel adds to the first one and LBD 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 LBD on expected future losses and hence to lower wedges.

As a result, when the first channel dominates, LBD may contribute to wedges that are higher than in the absence of LBD. Moreover, it can contribute to wedges declining over time. This is because the negative effect of LBD on expected future information rents is stronger when there are many periods ahead. Finally, note that, when sufficiently strong, these novel effects may also impact on the progressivity of the wedges, that is, the extent to which wedges increase with current productivity. In particular, LBD may introduce progressivity in the wedges under the same conditions that would have led to regressivity in the absence of LBD. The reason is that the benefit of distorting downwards current labor supply in order to economize on future informational rents is stronger for higher types, given that these are the types that expect, on average, larger rents in the future. The simple model in the previous section was an example if such environment.

We conclude by illustrating how LBD interacts with risk aversion in shaping the level, dynamics, and progressivity of the wedges. For this purpose, consider the same $T = 2$ economy as in Section 2 above, but now assume that the agent's preferences over consumption are CRRA with coefficient of relative-risk aversion $\hat{\eta}$. Risk-aversion affects wedges via two channels. First, due to the presence of skill-uncertainty, the allocation of consumption across time has an impact on the agent's incentives. Second, the cost of compensating the agent for his labor is typically higher for higher productivity types, who typically have lower marginal utility of consumption compared to lower productivity types.

Nevertheless, the qualitative nature of the effects of LBD on the progressivity and on the dynamics of the wedges discussed above appear robust to the possibility that the agent is risk averse. Figure 3 depicts the wedges in the case the agent's utility over consumption is CRRA for four different levels of the coefficient of relative risk aversion, namely for $\hat{\eta} = 0$, 0.2, 0.5, and 0.8. To isolate the effects of the agent's risk aversion, the figure assumes that the period-1 productivity θ_1 and the second-period shock ε_2 are drawn from the same Pareto-Lognormal distribution as in Figure 2 above and that the principal has the same Rawlsian objective as above. As the figure illustrates, risk aversion contributes

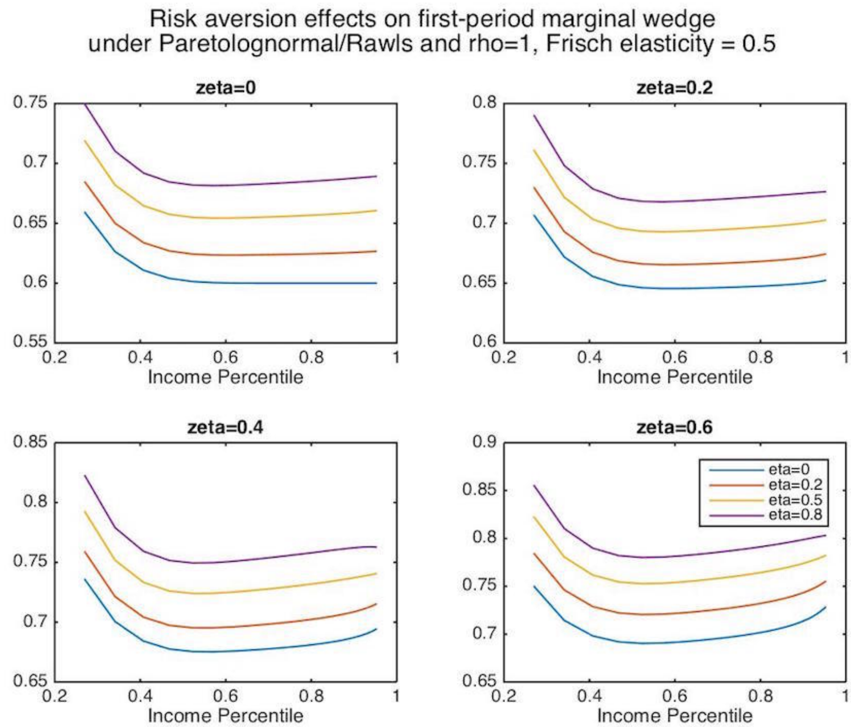

Figure 3: The CRRA Rawlsian Paretolognormal case

to an amplification of the period-1 wedge and to more progressivity. The reason is that, when the agent is risk-averse, the cost of compensating him for his marginal disutility of effort is higher as it now takes $\psi_y(\theta_t, y_t)/v'(c_t)$ units of consumption. Hence, risk aversion, by itself, contributes to higher wedges. What the figure illustrates is that the increase in the cost of compensation also amplifies the effects of LBD on the progressivity of the wedges. This is because, by increasing the cost of future information rents, risk aversion also increases the benefits of shifting the distribution of future types towards levels that command lower informational rents, and, as in the risk-neutral case, this effect is most pronounced at the top of the period-1 type distribution where the expectation of future rents is the highest. Similar results hold under less extreme inequality aversion on the principal's side. Figure 4 depicts the analogs of the results in Figure 3 for a planner with an Utilitarian objective—all other elements are the same as in Figure 3.

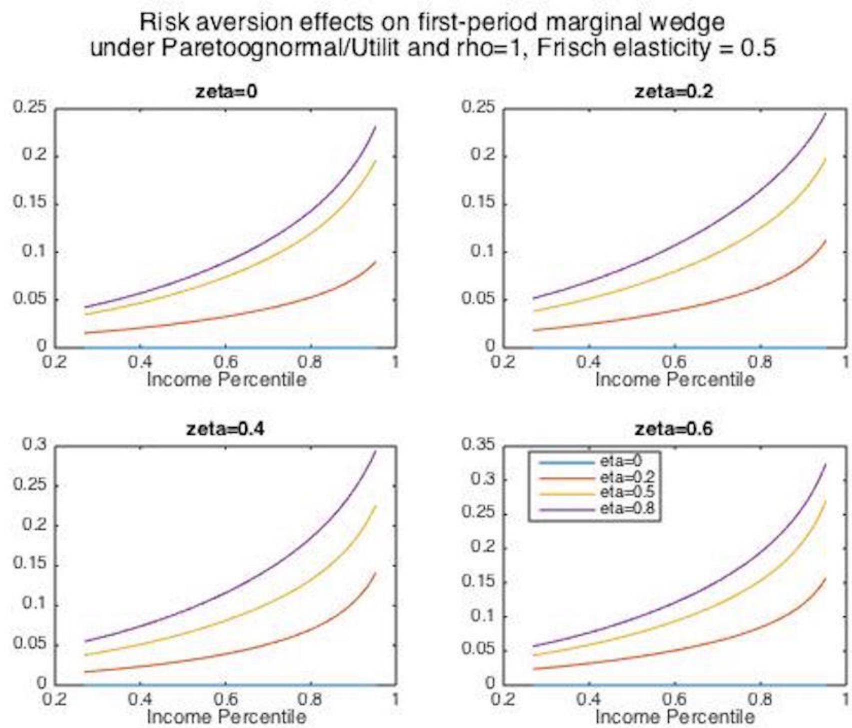

Figure 4: The CRRA Utilitarian Paretolognormal case

4 Conclusions

We study the effects of learning-by-doing on the dynamics of distortions in an economy in which the agents' private information evolves endogenously over time. We show that learning-by-doing may contribute to higher wedges (i.e., to higher marginal tax rates) that are declining over time. Furthermore, learning-by-doing can lead to wedges that are increasing in current types (equivalently, marginal tax rates increasing in earnings). These results are established for the same primitive conditions typically considered in the received taxation literature, where the theory, by abstracting from learning-by-doing, often predicts regressivity in the optimal tax code, and lower wedges that are increasing over time. Interesting extensions of our model could incorporate hidden savings and/or limited commitment, to mention few.

References

[1] Arrow, K., (1962): "The Economic Implications of Learning by Doing", *Review of Economic Studies*, 29(3), 155-173.

[2] Altugw, S., and R. Miller, (1998): "The Effect of Work Experience on Female Wages and Labour Supply", *Review of Economic Studies*, 65(1), 45-85.

[3] Baron, D. and D. Besanko (1984): "Regulation and Information in a Continuing Relation-ship," *Information Economics and Policy*, 1(3), 267–302.

[4] Battaglini, M. (2005): "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 95(3), 637-658.

[5] Bergemann, D. and A. Pavan (2015): "Introduction to JET Symposium on Dynamic Contracts and Mechanism Design," *Journal of Economic Theory,* Volume 159, 679-701.

[6] Besanko, D. (1985): "Multi-Period Contracts between Principal and Agent with Adverse Selection," *Economics Letters*, 17, 33-37.

[7] Best, M.C., and H.J. Kleven, (2013): "Optimal Income Taxation with Career Effects of Work Effort," mimeo, LSE.

[8] Courty, P. and H. Li (2000): 'Sequential Screening,' *Review of Economic Studies,* 67, 697–717.

[9] 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.

[10] Dustmann, C., and K. Meghir, (2005): "Wages, Experience and Seniority", *Review of Economic Studies*, 72, 77-108.

[11] Eso, P. and B. Szentes (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74, 705-731.

[12] Farhi, E., and I. Werning, (2013): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80, 596-635.

[13] Garrett, D., Pavan, A., (201: "Dynamic Managerial Compensation: a Variational Approach," *Journal of Economic Theory* 159, 775–818.

[14] Golosov, M., M. Troshkin, and A. Tsyvinski, (2015): "Optimal Dynamic Taxes," *American Economic Review,* forthcoming.

[15] Golosov, M., A. Tsyvinski, and I. Werning, (2006): "New Dynamic Public Finance: A User's Guide", NBER Macroeconomic Annual 2006, MIT Press.

[16] Jacobson L., R. LaLonde, and D. Sullivan, (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83(4), 685-709.

[17] 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.

[18] Kapicka, M., (2013b): "Optimal Mirrleesean Taxation with Unobservable Human Capital Formation," mimeo, UC Santa Barbara.

[19] Kapicka, M. and J. Neira (2014): "Optimal Taxation in a Life-Cycle Economy with Endogenous Human Capital Formation," mimeo, UC Santa Barbara.

[20] Kocherlakota, N. R., (2010): The New Dynamic Public Finance. Princeton University Press, USA.

[21] Krause, A., (2009): "Optimal Nonlinear Income Taxation with Learning-by-doing", *Journal of Public Economics*, 93(9-10), 1098-1110.

[22] Laffont, J-J and J. Tirole (1986): "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94, 614-641.

[23] Lucas, R.E., (1988): "On the Mechanics of Economic Development" *Journal of Monetary Economics*, 22(1), 33-42.

[24] Mirrlees, J., (1971): "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, 38, 175-208.

[25] Myerson, R. B., (1986): "Multistage Games with Communication," *Econometrica*, 54(2), 32358.

[26] Rogerson, W.P, (1985): "Repeated Moral Hazard," *Econometrica*, 53(1), 69-76.

[27] Stantcheva, S., (2014): "Optimal Taxation and Human Capital Policies over the Life Cycle," MIT, mimeo.

[28] Topel, R., (1991): "Specific Capital, Mobility and Wages: Wages Rise with Seniority," *Journal of Political Economy*, 99(t), 145-176.

[29] Willis, R., (1986): "Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions," in Handbook of Labor Economics, ed. by O. Ashenfelter and R. Layard, New York, Elsevier Science Publisher, Vol. 1, 525-602.

5 Appendix: Omitted Proofs

Proof of Proposition 1. Let $y_1(\theta_1)$ be the unique solution to the following equation

$$\left[1+\hat{W}_1^{RN}(\theta_1)\right]^{-1}\theta_1^{1+\phi}+\delta\zeta\bar{\varepsilon}(\phi)\theta_1^{\frac{(1+\phi)^2}{\phi}}\left[1+\hat{W}_1^{RN}(\theta_1)\right]^{-\frac{1+\phi}{\phi}}y_1^{\frac{\zeta(1+\phi)-\phi}{\phi}}-y_1^{\phi}=0 \tag{27}$$

with $\bar{\varepsilon}(\phi)\equiv\mathbb{E}\left[\varepsilon_2^{\frac{1+\phi}{\phi}}\right]$. Observe that the assumption that $\zeta\leq\phi/(1+\phi)$ implies that the left-hand-side of (27) is strictly concave. This, in turn, implies that the unique solution $y_1(\theta_1)$ to (27) is nondecreasing in θ_1 whenever $\hat{W}_1^{RN}(\theta_1)$ is nonincreasing. Finally, let

$$\begin{aligned}\Gamma(\theta_1) &\equiv \frac{1}{\bar{\varepsilon}(\phi)}\left[1+\hat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}\left\{\epsilon^{\hat{W}_1^{RN}}(\theta_1)\left[1-\frac{\hat{W}_1^{RN}(\theta_1)}{\phi(1+\hat{W}_1^{RN}(\theta_1))})\right]+\frac{1+\phi}{\phi}\right\}\\ &\quad+\frac{\zeta(1+\phi)-\phi}{\phi}\frac{1}{\bar{\varepsilon}(\phi)}\left[1+\hat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}\epsilon^{y_1}(\theta_1)+\delta\zeta\epsilon^{\hat{W}_1^{RN}}(\theta_1)\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}\end{aligned}$$

with $y_1(\theta_1)$ implicitly defined by (27),

$$\epsilon^{\hat{W}_1^{RN}}(\theta_1)\equiv\frac{d\hat{W}_1^{RN}(\theta_1)}{d\theta_1}\frac{\theta_1}{\hat{W}_1^{RN}(\theta_1)},$$

and

$$\epsilon^{y_1}(\theta_1)\equiv\frac{\theta_1 y_1'(\theta_1)}{y_1(\theta_1)}.$$

The proof proceeds in four steps. Step 1 shows that the period-1 wedge is given by formula (11) in the main text (which we repeat here for convenience)

$$\hat{W}_1(\theta_1)=\hat{W}_1^{RN}(\theta_1)\left\{1+\frac{\frac{1}{\phi}\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{1}{\delta\zeta\bar{\varepsilon}(\phi)}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}+\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}\right\}$$

where $y_1(\theta_1)$ is implicitly defined by (27) above. Given that at optimum $y_1(\theta_1)>0$, it is then immediate that LBD contributes to a higher period-1 wedge for all θ_1 and to first-period wedges that are higher than the second-period wedges (that is, $\hat{W}_1(\theta_1)>\hat{W}_1^{RN}(\theta_1)=\hat{W}_2^{RN}(\theta)$ all θ). These prove parts (i) and (ii) of the proposition. Step 2 shows that LBD contributes to higher progressivity if and only if condition (10) in the proposition holds, which establishes part (iii) in the proposition. Step 3 establishes the result in part (iv). Finally, Step 4 establishes part (v) by showing that, under the Pareto distribution, the allocation (y_1,y_2) satisfies the integral monotonicity conditions in (20).

Step 1. Recall from (9) that the effects of LBD on the period-1 wedge are summarized by the term

$$\Omega_1(\theta_1)=\frac{\delta}{\psi_y(y_1(\theta_1),\theta_1)}\widehat{W}_1^{RN}(\theta_1)\frac{\partial}{\partial y_1}\mathbb{E}^{\lambda[\chi]|\theta_1,y_1(\theta_1)}\left[\psi\left(y_2(\theta),\theta_2\right)\right]$$

Next use (**??**) for $t=1$ to verify that $y_1(\theta_1)$ is given by the first-order-condition

$$1+LD_1^{\chi}(\theta_1)=\psi_y(y_1(\theta_1),\theta_1)\left[1+\widehat{W}_1^{RN}(\theta_1)\right] \tag{28}$$

with

$$LD_1^{\chi}(\theta_1) = LD_1^{FB;\chi}(\theta_1) - \delta \widehat{W}_1^{RN}(\theta_1) \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right]$$

and (recalling (15))

$$\begin{aligned} LD_1^{FB;\chi}(\theta_1) &= \delta \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[y_2(\theta) - \psi(y_2(\theta), \theta_2)\right] \\ &= \delta \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[y_2(\theta)\right] - \delta \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right] \end{aligned} \tag{29}$$

Now, use (**??**) for $t = 2$ to see that $y_2(\theta)$ is given by the first-order-condition:

$$1 = \psi_y(y_2(\theta), \theta_2) \left[1 + \widehat{W}_1^{RN}(\theta_1)\right]$$

which, when ψ is isoelastic, can be rewritten as

$$y_2(\theta) = (1 + \phi) \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \psi(y_2(\theta), \theta_2). \tag{30}$$

Replacing (30) into (29), we have that

$$LD_1^{FB;\chi}(\theta_1) = \delta \left\{(1 + \phi) \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] - 1\right\} \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right]$$

and hence that

$$LD_1^{\chi}(\theta_1) = \delta\phi \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right] \tag{31}$$

Using (28), in turn we have that

$$\frac{1}{\psi_y(y_1(\theta_1), \theta_1)} = \frac{1 + \widehat{W}_1^{RN}(\theta_1)}{1 + \delta\phi \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right]} \tag{32}$$

Replacing (32) into (9), we have that

$$\Omega_1(\theta_1) = \frac{\delta \widehat{W}_1^{RN}(\theta_1) \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi\left(y_2(\theta), \theta_2\right)\right]}{1 + \delta\phi \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right]}$$

or, equivalently,

$$\Omega_1(\theta_1) = \frac{\delta \widehat{W}_1^{RN}(\theta_1) \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \Lambda^{\chi}(\theta_1, y_1(\theta_1))}{1 + \delta\phi \left[1 + \widehat{W}_1^{RN}(\theta_1)\right] \Lambda^{\chi}(\theta_1, y_1(\theta_1))}$$

where we used the shortcut notation

$$\Lambda^{\chi}(\theta_1, y_1(\theta_1)) \equiv \frac{\partial}{\partial y_1} \mathbb{E}^{\lambda[\chi]|\theta_1, y_1(\theta_1)} \left[\psi(y_2(\theta), \theta_2)\right]$$

Next, use (30) to observe that, when ψ is isoelastic,

$$y_2(\theta) = \theta_2^{\frac{1+\phi}{\phi}} \left[1 + \widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1}{\phi}} \tag{33}$$

and hence

$$\psi\left(y_2(\theta),\theta_2\right)=\frac{1}{1+\phi}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1+\phi}{\phi}}\cdot\theta_2^{\frac{1+\phi}{\phi}}$$

It follows that

$$\Lambda^{\chi}(\theta_1,y_1(\theta_1))=\frac{1}{1+\phi}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1+\phi}{\phi}}\frac{\partial}{\partial y_1}\mathbb{E}\left[\theta_2^{\frac{1+\phi}{\phi}}|\theta_1,y_1(\theta_1)\right]$$

Using

$$\overline{\varepsilon}(\phi)=\mathbb{E}\left[\varepsilon_2^{\frac{1+\phi}{\phi}}\right]$$

we have that

$$\begin{aligned}
&\frac{\partial}{\partial y_1}\left\{\mathbb{E}\left[\theta_2^{\frac{1+\phi}{\phi}}|\theta_1,y_1\right]\right\}=\frac{1+\phi}{\phi}\mathbb{E}\left[\theta_2^{\frac{1}{\phi}}\left(-\frac{F_{2,y}(\theta_2|\theta_1,y_1)}{f_2(\theta_2|\theta_1,y_1)}\right)|\theta_1,y_1\right]\\
&=\frac{1+\phi}{\phi}\mathbb{E}\left[\theta_2^{\frac{1}{\phi}}\left(\frac{\theta_2\zeta}{y_1}\right)|\theta_1,y_1\right]=\frac{1+\phi}{\phi}\mathbb{E}\left[\theta_2^{\frac{1+\phi}{\phi}}\left(\frac{\zeta}{y_1}\right)|\theta_1,y_1\right]\\
&=\frac{\zeta(1+\phi)}{\phi}\frac{1}{y_1}(\theta_1 y_1^{\zeta})^{\frac{1+\phi}{\phi}}\overline{\varepsilon}(\phi).
\end{aligned}$$

This implies that

$$\Lambda^{\chi}(\theta_1,y_1(\theta_1))=\left\{\frac{1}{1+\phi}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1+\phi}{\phi}}\frac{\zeta(1+\phi)}{\phi}\overline{\varepsilon}(\phi)\theta_1^{\frac{1+\phi}{\phi}}\right\}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}\tag{34}$$

Replacing the formula for $\Lambda^{\chi}(\theta_1,y_1(\theta_1))$ into the formula for $\Omega_1(\theta_1)$ above, we then have that the latter can be expressed as

$$\Omega_1(\theta_1)=\frac{\frac{1}{\phi}\widehat{W}_1^{RN}(\theta_1)\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{1}{\delta\zeta\overline{\varepsilon}(\phi)}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}+\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}\tag{35}$$

Replacing (35) into (8) for $t=1$, permits us to establish the formula for $\widehat{W}_1(\theta_1)$ in (11) above.

We now show that, at the optimum, $y_1(\theta_1)$ is implicitly given by equation (27) in the proposition. To see this, use again (28), (31), the definition of $\Lambda^{\chi}(\theta_1,y_1(\theta_1))$ and (34) to verify that $y_1(\theta_1)$ must solve equation (27) above.

Step 2. Taking the derivative of $\Omega_1(\theta_1)$ in (35) and simplifying it using the fact that, at the optimum, $y_1(\theta_1)>0$, we have that $\Omega_1(\theta_1)$ is increasing in θ_1 if and only if

$$\begin{aligned}
&\epsilon^{\widehat{W}_1^{RN}}(\theta_1)\left\{\frac{1}{\delta\zeta\overline{\varepsilon}(\phi)}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}\left[1-\frac{\widehat{W}_1^{RN}(\theta_1)}{\phi(1+\widehat{W}_1^{RN}(\theta_1))})\right]+\theta_1^{\frac{1+\phi}{\phi}}y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}\right\}\\
&+\left[\frac{1+\phi}{\phi}+\frac{\zeta(1+\phi)-\phi}{\phi}\epsilon^{y_1}(\theta_1)\right]\frac{1}{\delta\zeta\overline{\varepsilon}(\phi)}\left[1+\widehat{W}_1^{RN}(\theta_1)\right]^{\frac{1}{\phi}}\geq 0.
\end{aligned}$$

Rearranging, we have that $\Omega_1(\theta_1)$ is increasing in θ_1 if and only if $\Gamma(\theta_1)\geq 0$, thus establishing part (iii) in the proposition.

Step 3. Now observe that, when F_1 is Pareto, $\eta_1(\theta_1)\theta_1 = \lambda$, in which case

$$\widehat{W}_1^{RN}(\theta_1) = \widehat{W}_2^{RN}(\theta) = \frac{1+\phi}{\lambda} \text{ all } \theta,$$

and equation (27) reduces to

$$T_1(\theta_1) + T_2(\theta_1) y_1^{\frac{\zeta(1+\phi)-\phi}{\phi}} - y_1^{\phi} = 0 \tag{36}$$

where

$$\begin{aligned} T_1(\theta_1) &\equiv \left[1 + \frac{1+\phi}{\lambda}\right]^{-1} \theta_1^{1+\phi} \\ T_2(\theta_1) &\equiv \delta\zeta\overline{\varepsilon}(\phi)\theta_1^{\frac{(1+\phi)^2}{\phi}} \left[1 + \frac{1+\phi}{\lambda}\right]^{-\frac{1+\phi}{\phi}} \end{aligned}$$

Furthermore, in this case,

$$\Gamma(\theta_1) = \frac{1}{\overline{\varepsilon}(\phi)} \left[1 + \frac{1+\phi}{\lambda}\right]^{\frac{1}{\phi}} \frac{1+\phi}{\phi} + \frac{\zeta(1+\phi)-\phi}{\phi} \frac{1}{\overline{\varepsilon}(\phi)} \left[1 + \frac{1+\phi}{\lambda}\right]^{\frac{1}{\phi}} \epsilon^{y_1}(\theta_1).$$

Now use (36) to obtain that

$$\frac{dy_1(\theta_1)}{d\theta_1} = -\frac{\frac{dT_1(\theta_1)}{d\theta_1} + \frac{dT_2(\theta_1)}{d\theta_1} y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{\zeta(1+\phi)-\phi}{\phi} T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-2\phi}{\phi}} - \phi y_1(\theta_1)^{\phi-1}}. \tag{37}$$

Using the fact that, for all θ_1, $y_1(\theta_1) > 0$, we then have that

$$\frac{dy_1(\theta_1)}{d\theta_1} = -\frac{\frac{dT_1(\theta_1)}{d\theta_1} y_1(\theta_1) + \frac{dT_2(\theta_1)}{d\theta_1} y_1(\theta_1)^{\frac{\zeta(1+\phi)}{\phi}}}{\frac{\zeta(1+\phi)-\phi}{\phi} T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}} - \phi y_1(\theta_1)^{\phi}} \tag{38}$$

Replacing

$$y_1(\theta_1)^{\phi} = T_1(\theta_1) + T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}$$

into (38), we then have that

$$\frac{dy_1(\theta_1)}{d\theta_1} = -\frac{\frac{dT_1(\theta_1)}{d\theta_1} y_1(\theta_1) + \frac{dT_2(\theta_1)}{d\theta_1} y_1(\theta_1)^{\frac{\zeta(1+\phi)}{\phi}}}{\frac{\zeta(1+\phi)-\phi}{\phi} T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}} - \phi \left[T_1(\theta_1) + T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}\right]}$$

Rearranging, we have that

$$\frac{dy_1(\theta_1)}{d\theta_1} = -\frac{\frac{dT_1(\theta_1)}{d\theta_1} y_1(\theta_1) + \frac{dT_2(\theta_1)}{d\theta_1} y_1(\theta_1)^{\frac{\zeta(1+\phi)}{\phi}}}{\frac{(\zeta-\phi)(1+\phi)}{\phi} T_2(\theta_1) y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}} - \phi T_1(\theta_1)} \tag{39}$$

Now note that

$$\frac{dT_1(\theta_1)}{d\theta_1} = (1+\phi)\frac{T_1(\theta_1)}{\theta_1}$$

$$T_2(\theta_1) = \delta\zeta\overline{\varepsilon}(\phi)\theta_1^{\frac{(1+\phi)}{\phi}}\left[1+\frac{1+\phi}{\lambda}\right]^{-\frac{1}{\phi}}T_1(\theta_1)$$

$$\frac{dT_2(\theta_1)}{d\theta_1} = \frac{(1+\phi)^2}{\phi}\frac{T_2(\theta_1)}{\theta_1}$$

Replacing these functions into (39), and letting $n(\theta_1) \equiv \delta\zeta\overline{\varepsilon}(\phi)\theta_1^{\frac{(1+\phi)}{\phi}}\left[1+\frac{1+\phi}{\lambda}\right]^{-\frac{1}{\phi}}$, we then have that

$$\epsilon^{y_1}(\theta_1) = -\frac{1+\phi+\frac{(1+\phi)^2}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{(\zeta-\phi)(1+\phi)}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}-\phi}. \tag{40}$$

It follows that

$$\Gamma(\theta_1) = \frac{1}{\overline{\varepsilon}(\phi)}\left[1+\frac{1+\phi}{\lambda}\right]^{\frac{1}{\phi}}\frac{1+\phi}{\phi}\left\{1+\left(\frac{\phi}{1+\phi}-\zeta\right)\left[\frac{1+\phi+\frac{(1+\phi)^2}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{(\zeta-\phi)(1+\phi)}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}-\phi}\right]\right\}$$

Hence $\Gamma(\theta_1) > 0$ if and only if

$$1+\left(\frac{\phi}{1+\phi}-\zeta\right)\left[\frac{1+\phi+\frac{(1+\phi)^2}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}}{\frac{(\zeta-\phi)(1+\phi)}{\phi}n(\theta_1)y_1(\theta_1)^{\frac{\zeta(1+\phi)-\phi}{\phi}}-\phi}\right] > 0. \tag{41}$$

Now fix θ_1 and observe that the left-hand side of (41) is increasing in $y_1(\theta_1)$. A sufficient condition for $\Gamma(\theta_1) > 0$ is thus that the inequality in (41) holds when $y_1(\theta_1) = 0$. It is easy to see that, when $y_1(\theta_1) = 0$, the left-hand side of (41) reduces to

$$\zeta\frac{1+\phi}{\phi}$$

which is obviously positive. The result in part (iv) then follows from the property above, along with the result in part (iii).

Step 4. First use (33) to observe that $y_2(\theta_1,\theta_2)$ is nondecreasing in θ_2 for any θ_1. Next note that (4) is equivalent to

$$\int_{\hat{\theta}_1}^{\theta_1}\left[\frac{y_1(s)^{1+\phi}}{s^{2+\phi}}+\frac{\delta\rho}{s}\int_0^{+\infty}\left(\frac{y_2(s,z)}{z}\right)^{1+\phi}f_2(z\mid s,y_1(s))dz\right]ds \geq$$

$$\int_{\hat{\theta}_1}^{\theta_1}\left[\frac{y_1\left(\hat{\theta}_1\right)^{1+\phi}}{s^{2+\phi}}+\frac{\delta\rho}{s}\int_0^{+\infty}\left(\frac{y_2\left(\hat{\theta}_1,z\right)}{z}\right)^{1+\phi}f_2(z\mid s,y_1(\hat{\theta}_1))dz\right]ds$$

for any θ_1 and $\hat{\theta}_1 < \theta_1$. Define now the variable $e_2(s,\varepsilon)$ according to

$$e_2(s,\varepsilon) = \frac{y_2(s,s^\rho y_1(s)^\zeta\varepsilon)}{s^\rho y_1(s)^\zeta\varepsilon}$$

Using this definition, the change of variables $z = s^{\rho} y_1(s)^{\zeta} \varepsilon$ in the left integral, the change of variables $z = s^{\rho} y_1(\hat{\theta}_1)^{\zeta} \varepsilon$ in the right integral, and noting that

$$e_2\left(\hat{\theta}_1, \left(\frac{s}{\hat{\theta}_1}\right)^{\rho} \varepsilon\right) = \frac{y_2(\hat{\theta}_1, s^{\rho} y_1(\hat{\theta}_1)^{\zeta} \varepsilon)}{s^{\rho} y_1(\hat{\theta}_1)^{\zeta} \varepsilon}$$

we have that the above inequality can be rewritten as:

$$\int_{\hat{\theta}_1}^{\theta_1} \left[\frac{y_1(s)^{1+\phi}}{s^{2+\phi}} + \frac{\delta\rho}{s} \int_0^{+\infty} \left[e_2(s,\varepsilon)^{1+\phi}\right] g(\varepsilon) d\varepsilon\right] ds \geq$$

$$\int_{\hat{\theta}_1}^{\theta_1} \left[\frac{y_1\left(\hat{\theta}_1\right)^{1+\phi}}{s^{2+\phi}} + \frac{\delta\rho}{s} \int_0^{+\infty} \left[e_2\left(\hat{\theta}_1, \left(\frac{s}{\hat{\theta}_1}\right)^{\rho} \varepsilon\right)^{1+\phi}\right] g(\varepsilon) d\varepsilon\right] ds$$

Clearly, then, the inequality above is satisfied if for all θ_1, $\hat{\theta}_1 < \theta_1$ and ε, both 1 and 2 below hold:

1. $y_1(\theta_1)$ is nondecreasing
2. $e_2(s,\varepsilon) \geq e_2\left(\hat{\theta}_1, \left(\frac{s}{\hat{\theta}_1}\right)^{\rho} \varepsilon\right)$ for all $\hat{\theta}_1 \leq s \leq \theta_1$

Recall the definitions above for $e_2(s,\varepsilon)$ and $e_2\left(\hat{\theta}_1, \left(\frac{s}{\hat{\theta}_1}\right)^{\rho} \varepsilon\right)$ we have that the inequality 2. above can be expressed equivalently as

$$\frac{y_2(s, s^{\rho} y_1(s)^{\zeta} \varepsilon)}{y_1(s)^{\zeta}} \geq \frac{y_2(\hat{\theta}_1, s^{\rho} y_1(\hat{\theta}_1)^{\zeta} \varepsilon)}{y_1(\hat{\theta}_1)^{\zeta}}$$

Recalling again (33), observe now that

$$\frac{y_2(\theta_1, s^{\rho} y_1(\theta_1)^{\zeta} \varepsilon)}{y_1(\theta_1)^{\zeta}} = [s^{\rho}\varepsilon]^{\frac{1+\phi}{\phi}} \left[1 + \widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1}{\phi}} y_1(\theta_1)^{\frac{\zeta}{\phi}}$$

Clearly, then, properties 1. and 2. above are satisfied if for all θ_1, $\hat{\theta}_1 < \theta_1$ and ε, both 1 and 2 below hold:

1. $y_1(\theta_1)$ is nondecreasing
2. $\left[1 + \widehat{W}_1^{RN}(\theta_1)\right]^{-\frac{1}{\phi}} y_1(\theta_1)^{\frac{\zeta}{\phi}}$ is nondecreasing

It is immediate then that the above sufficient conditions for the period-2 integral constraint to be satisfied hold if $y_1(\theta_1)$ is nondecreasing *and* $\left[1 + \widehat{W}_1^{RN}(\theta_1)\right]$ is nonincreasing. The desired result follows then directly from the fact, as we have mentioned immediately above Proposition 1, $y_1(\theta_1)$ as given by (27) is nondecreasing whenever $\left[1 + \widehat{W}_1^{RN}(\theta_1)\right]$ is nonincreasing. Q.E.D.

Proof of Proposition 2. For any $t > 1$, let $\pi_t(\theta^{t-1})$ be the multiplier of the period-t, with $t > 1$, promise keeping constraint (13). 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 constraint (14), which is an integral constraint. At optimum, the following necessary conditions with respect to $y_t(\theta^t)$, $V_t(\theta^t)$ and $\Pi_{t+1}(\theta^t)$ must hold with probability one:

$$v^{P\prime}(y_t^*(\theta^t)) - \frac{\psi_y(y_t^*(\theta^t), \theta_t)}{v'(c_t^*(\theta^t))} + \delta \frac{\partial Q_{t+1}^{FB}(\theta^t, y_t^*(\theta^t), \Pi_{t+1}^*(\theta^t))}{\partial y_t} = 0 \text{ for any } t = 1, ..., T$$

$$\frac{1}{v'(c_t^*(\theta^t))} + \pi_t^*(\theta^{t-1}) q_t'(V_t^*(\theta_t)) = 0 \text{ for any } t = 1, ..., T$$

$$\frac{1}{v'(c_t^*(\theta^t))} + \frac{\partial Q_{t+1}^{FB}(\theta^t, y_t^*(\theta^t), \Pi_{t+1}^*(\theta^t))}{\partial \Pi_{t+1}} = 0 \text{ for any } t < T$$

with $\pi_1(\theta^0) \equiv \pi_1$. Now use the envelope theorem to establish that

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

while for $t < T$

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

and

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

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

Turn now to the case of $r = 0$, and let $\hat{\pi}_1(\theta_1)$ be the (Kuhn-Tucker) multiplier of the redistribution constraint (14), which is now a constraint on the control variable $V_1^A(\theta_1)$. At optimum, the above necessary conditions hold, except the second one for $t = 1$ which is replaced by

$$\frac{1}{v'(c_1^*(\theta_1))} - \hat{\pi}_1^*(\theta_1) = 0$$

Moreover Q.E.D.

Proof of Proposition 3. To solve the principal's problem under asymmetric information, recall first from the discussion of first-best that

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

denotes the agent's (on-path) continuation payoff starting from period t, when the skills history is θ^t, where for any $t < T$,

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

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

Next, for any $t < T$, let

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

with $Z_{T+1}(\theta) \equiv 0$ if $T < +\infty$. Using this notation, the local incentive-compatibility constraint (referred to hereafter as ICFOC) constraints (19) in period t can be conveniently rewritten as

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

Using (19) for $t+1$, along with the law of motion of the impulse responses, we also have that

$$Z_{t+1}(\theta^t) = \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y_t(\theta^t)) dF_{t+1}(\theta_{t+1} \mid \theta_t, y_t(\theta^t)).$$

The property of the impulse responses described in (18), in turn, implies that

$$Z_{t+1}(\theta^t) = \int V_{t+1}(\theta^{t+1}) dF_{t+1,\theta}(\theta_{t+1}|\theta_t, y_t(\theta^t)).$$

Equipped with this notation, the principal's problem can be conveniently rewritten recursively as follows

$$\begin{aligned} & Q_t(\theta^{t-1}, y_{t-1}(\theta^{t-1}), \Pi_t(\theta^{t-1}), Z_t(\theta^{t-1})) \\ \equiv & \max_{y_t(\theta^{t-1},\cdot), V_t(\theta^{t-1},\cdot), \Pi_{t+1}(\theta^{t-1},\cdot), Z_{t+1}(\theta^{t-1},\cdot)} \mathbb{E}[\hat{Q}_t(\theta^t, y_t(\theta^t), V_t(\theta^t), \Pi_{t+1}(\theta^t), Z_{t+1}(\theta^t))|\theta^{t-1}, y_{t-1}(\theta^{t-1})] \end{aligned}$$

subject to[19]

$$\Pi_{T+1}(\theta) = Z_{T+1}(\theta) = 0, \text{ all } \theta \text{ (if } T < +\infty) \tag{42}$$

$$\kappa = (1-r)V_1(\underline{\theta}_1) + r\int q(V_1(\theta_1)) dF_1(\theta_1) \tag{43}$$

and

$$\frac{\partial V_t(\theta^{t-1}, \theta_t)}{\partial \theta_t} = -\psi_\theta(y_t(\theta^t), \theta_t) + \delta Z_{t+1}(\theta^t) \text{ all } t, \text{ all } \theta^t \tag{44}$$

$$\Pi_t(\theta^{t-1}) = \int V_t(\theta^t) dF_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})) \text{ all } t > 1, \text{ all } \theta^{t-1} \tag{45}$$

$$Z_t(\theta^{t-1}) = \int V_t(\theta^t) dF_{t,\theta}(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})) \text{ all } t > 1, \text{ all } \theta^{t-1} \tag{46}$$

[19]Strictly speaking, the equality constraint $\kappa = \int q(V_1(\theta_1))dF_1(\theta_1)$ should be replaced by its inequality counterpart $\int q(V_1(\theta_1))dF_1(\theta_1) \geq \kappa$. However, because such constraint always binds at the optimum, we treat it as an equality.

where[20]

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

In words, the principal's problem consists in choosing, in each period t, and for each current type θ_t, the output $y_t(\theta^{t-1}, \theta_t)$, the continuation utility, $V_t(\theta^{t-1}, \theta_y)$ from period t (included) onwards, the agent's future expected continuation utility, $\Pi_{t+1}(\theta^{t-1}, \theta_t)$, and the derivative of the agent's future expected continuation utility, $Z_{t+1}(\theta^{t-1}, \theta_t)$. All these functions must be selected jointly to maximize the combination of the principal's current and future payoff (taking into account that the principal will face a similar optimization problem in future periods, as is always the case in dynamic programming), and subject to the requirement that the current policies be consistent with what promised in previous period (i.e. with $\Pi_t(\theta^{t-1})$ and $Z_t(\theta^{t-1})$), when $t > 1$. In period-1, instead, the principal is constrained only by (43). This is a weaker version of the redistribution constraint (12). Its use is more convenient here and it comes without loss of generality because the necessary envelope condition for incentive compatibility implies that the agent's continuation payoff is increasing in the current type. Note that the key difference with respect to the corresponding recursive problem in the first-best benchmark is twofold: (a) the the constraint that the derivative of the agent's continuation payoff satisfies the necessary envelope condition for incentive compatibility, which, in recursive form corresponds to (44), and (b) the constraint that the derivative of the agent's future expected continuation utility equals $Z_{t+1}(\theta^{t-1}, \theta_t)$, as described by (46).

In any period t, the principal is, in effect, 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 integral constraints (45) and (46) arising from having set the "promised" expected utility $\Pi_t(\theta^{t-1})$ and "marginal promised" expected utility $Z_t(\theta^{t-1})$ in the period$-t-1$ problem. Also note that constraint (43) can be conveniently rewritten as

$$\Pi_1 = (1-r)V_1(\underline{\theta}_1) + r\int q(V_1(\theta_1))dF_1(\theta_1) \tag{47}$$

by letting $\Pi_1 \equiv \kappa$. Then let π_1 be the multiplier associated to constraint (47) and set $\xi_1 \equiv 0$. Define also with some abuse of notation $q'_t(V;r)$ by $q'_1(V;r) = rq'(V)$ and, for $t > 1$, $q'_t(V;r) = 1$ Moreover, to accommodate for moving support let $\mathbf{1}_{\underline{\theta}_t}$ be the index function that takes the value 1 if and only if $\underline{\theta}_t$ depends non-trivially on history. Similarly, let $\mathbf{1}_{\overline{\theta}_t}$ be the index function that takes the value 1

[20]We ignore the constraint $Z_{t+1}(\theta^t) \geq 0$, because it is going to be slack at an optimum whenever period-$t+1$ production does take place for some period-$t+1$ types (i.e., when $y_{t+1}(\theta^{t+1}) > 0$ for some θ^{t+1}. To see this, note that if $Z_{t+1}(\theta^t) = 0$ was the case, then the definition of $Z_{t+1}(\theta^t)$ (see (5) in the main text) and that $I_t^\tau(\theta^\tau, y^{\tau-1}(\theta^{\tau-1})) > 0$ and $\psi_\theta(y, .) < 0$ for $y > 0$ and $\psi_\theta(0, .) = 0$ would imply zero production for every period-$t+1$ productivity.

if and only if $\overline{\theta}_t$ depends non-trivially on history. Assume that if $\underline{\theta}_t$ and $\overline{\theta}_t$ depend non-trivially on history, then they are differentiable functions of history.

Start with the case of $r = 1$. The necessary conditions of the principal's problem are then that for period-t the following conditions hold with probability one: condition (22) in the proposition, the following conditions for all t

$$\frac{\partial \mu_t(\theta^t)}{\partial \theta_t} = f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))\cdot \tag{48}$$

$$\cdot \left\{ \frac{1}{v'(c_t(\theta^t))} + \pi_t(\theta^{t-1}) q_t'(V_t(\theta^t; r)) + \xi_t(\theta^{t-1}) \frac{\partial f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))/\partial \theta_{t-1}}{f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))} \right\}$$

$$\mu_t(\underline{\theta}_t, \theta^{t-1}) = \xi_t(\theta^{t-1}) I_{t-1}^t(\underline{\theta}_t, \theta^{t-1}, y_{t-1}(\theta^{t-1})) f_t(\underline{\theta}_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})) \mathbf{1}_{\underline{\theta}_t} \tag{49}$$

$$\mu_t(\overline{\theta}_t, \theta^{t-1}) = \xi_t(\theta^{t-1}) I_{t-1}^t(\overline{\theta}_t, \theta^{t-1}, y_{t-1}(\theta^{t-1})) f_t(\overline{\theta}_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})) \mathbf{1}_{\overline{\theta}_t} \tag{50}$$

and the following conditions for all $1 \leq t < T$

$$\frac{1}{v'(c_t(\theta^t))} + \pi_{t+1}(\theta^t) = 0 \tag{51}$$

$$\mu_t(\theta^t) + \xi_{t+1}(\theta^t) f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1})) = 0 \tag{52}$$

where we have used that $\frac{\partial Q_{t+1}}{\partial \Pi_{t+1}} = \pi_{t+1}$ and $\frac{\partial Q_{t+1}}{\partial Z_{t+1}} = \xi_{t+1}$. The solution to the principal's problem is then given by the above necessary conditions, along with the laws of motions for the agent's continuation payoff (44), the constraints (45), (46), and (47), 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

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

while for $t < T$

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

We start the analysis of the above necessary conditions by noting that after using (51) and (52) to eliminate $\pi_t(\theta^{t-1})$ and $\xi_t(\theta^{t-1})$ from the first necessary condition (48) for $t > 1$, the derived law of motion for the co-state variable, alongside the boundary conditions (49) and (50) for $t > 1$, yields

the familiar Rogerson-Euler condition (21) in the proposition (recalling that $q'_{t+1}(V_{t+1}(\theta^{t+1};r)) = 1$ for $t = 1, ..., T-1$) and that

$$\frac{-\mu_t(\theta^t)}{f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))} = \frac{1}{\eta_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \int_{\theta_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \tilde{\theta}_t))} \frac{dF_t(\tilde{\theta}_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{1 - F_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}$$
$$- \frac{p_t(\theta^{t-1})}{\eta_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} - \frac{\mu_{t-1}(\theta^{t-1})}{f_{t-1}(\theta_{t-1} \mid \theta_{t-2}, y_{t-2}(\theta^{t-2}))} I_{t-1}^t(\theta^t, y^{t-1}(\theta^{t-1}))$$

where

$$\eta_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1})) \equiv \frac{f_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}{1 - F_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))}$$

and

$$p_t(\theta^{t-1}) \equiv \int_{\underline{\theta}_t}^{\bar{\theta}_t} \frac{1}{v'(c_t(\theta^{t-1}, \theta_t))} dF_t(\theta_t \mid \theta_{t-1}, y_{t-1}(\theta^{t-1}))$$

A straightforward rearrangement of terms leads to condition (23) for $t > 1$ in the proposition.

Next, recalling the definition $\xi_1 = 0$, we have that the necessary condition (48) for $t = 1$, alongside the boundary conditions (49) and (50) for $t = 1$, yields

$$\frac{-\mu_1(\theta_1)}{f_1(\theta_1)} = \frac{1}{\eta_1(\theta_1)} \left[\int_{\theta_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\widetilde{\theta}_1))} \frac{dF_1(\tilde{\theta}_1)}{1 - F_1(\theta_1)} + \pi_1 \int_{\theta_1}^{\bar{\theta}_1} q'(V_1(\widetilde{\theta}_1)) \frac{dF_1(\tilde{\theta}_1)}{1 - F_1(\theta_1)} \right]$$

where

$$\pi_1 = -\frac{\int_{\underline{\theta}_1}^{\bar{\theta}_1} \frac{1}{v'(c_1(\theta_1))} dF_1(\theta_1)}{\int_{\underline{\theta}_1}^{\bar{\theta}_1} q'(V_1(\theta_1)) dF_1(\theta_1)} < 0$$

A straightforward rearrangement of terms leads to condition (23) for $t = 1$ in the proposition.

For the case of $r = 0$, we only need to observe that the boundary condition (49) for $t = 1$ is now not relevant. Recalling that $\xi_1 = 0$, we have that the necessary condition (48) for $t = 1$, alongside the boundary condition (50) for $t = 1$, leads to condition (23) for $t = 1$ in the proposition. Q.E.D.

Proof of Theorem 1. Recall from the previous proof the formula that gives $LD_t^\chi(\theta^t)$. For $t < T$, eliminate the multipliers $\pi_{t+1}(\theta^t)$ and $\xi_{t+1}(\theta^t)$ by using the optimality conditions with respect to $\Pi_{t+1}(\theta^t)$ and $Z_{t+1}(\theta^t)$ stated in the above proof ((51) and (52)) to arrive at

$$LD_t^\chi(\theta^t) = \delta \int [v^P(y_{t+1}(\theta^{t+1})) - c_{t+1}(\theta^{t+1})$$
$$+\delta Q_{t+2}(\theta^{t+1}, y_{t+1}(\theta^{t+1}), \Pi_{t+2}(\theta^{t+1}), Z_{t+2}(\theta^{t+1}))] dF_{t+1,y}(\theta_{t+1}|\theta_t, y_t(\theta^t))$$
$$+\frac{\delta}{v'(c_t(\theta^t))} \int V_{t+1}(\theta^{t+1}) dF_{t+1,y}(\theta_{t+1}|\theta_t, y_t(\theta^t))$$
$$+\delta \frac{\mu_t(\theta^t)}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \frac{\partial}{\partial y_t} \int V_{t+1}(\theta^{t+1}) dF_{t+1,\theta}(\theta_{t+1}|\theta_t, y_t(\theta^t))$$

Using the general definition of $LD_t^{FB;\chi}(\theta^t)$ in (15) recognizing that

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

is the principal's period-$t+1$ continuation payoff at truthful history θ^{t+1}, and using the properties of impulse responses, we then have that

$$LD_t^{\chi}(\theta^t) = LD_t^{FB;\chi}(\theta^t)$$
$$+\delta \frac{\mu_t(\theta^t)}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))} \frac{\partial}{\partial y_t} \int \frac{\partial V_{t+1}(\theta^{t+1})}{\partial \theta_{t+1}} I_t^{t+1}(\theta^{t+1}, y^t(\theta^t)) dF(\theta_{t+1}|\theta_t, y_t(\theta^t))$$

Using the local IC constraint (44) along with the fact that $I_1^\tau = I_1^s I_s^\tau$ for any (τ, s), $\tau > s$, and the definition of period-t handicap in the main text, we then have that

$$LD_t^{FB;\chi}(\theta^t) - LD_t^{\chi}(\theta^t) = \frac{-\mu_t(\theta^t)\eta(\theta_1)}{f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))I_1^t(\theta^t, y^{t-1}(\theta^{t-1}))} \cdot \tag{53}$$
$$\cdot \delta \frac{\partial}{\partial y_t} \mathbb{E}^{\lambda[\chi]|\theta^t, y_t(\theta^t)} \left[\sum_{\tau=t+1}^{T} \delta^{\tau-t-1} l_\tau(\theta^\tau, y^\tau(\theta^\tau)) \right]$$

Recall the definition

$$\widehat{\mu}_t(\theta^t) = \frac{-\mu_t(\theta^t)}{\theta_t f_t(\theta_t|\theta_{t-1}, y_{t-1}(\theta^{t-1}))}$$

Using then the definition of wedges, (22) and (53), one can then show, after some straightforward rearrangement of terms, that the (relative) marginal wedges in period t, can be expressed as in the statement of the theorem. Q.E.D.