Optimal Dynamic Redistribution with Learning-by-doing

Miltiadis Makris

March 15, 2012

Department of Economics, University of Southampton

E-mail: M.Makris@soton.ac.uk.

Preliminary and Incomplete

Abstract

Focusing on a Markov environment with standard time-separable preferences, and assuming pre-commitment on the part of a planner with per-period redistributional motives, we investigate optimal dynamic redistribution when labour experience impacts on stochastic future productivity. We first derive an envelope condition that must characterise incentive-compatible allocations. This condition is similar to the one derived recently in Kapicka (2008) and Pavan et al. (2009), but we show it for a more general class of productivity distributions that allow us to incorporate learning-by-doing in skill distributions such as Pareto and Lognormal, which have been central in discussions of optimal income taxation. Echoing the approaches in Kapicka (2008), Pavan et al (2009), we then derive the optimal allocations from the perspective of the planner, if this envelope condition was fully characterising incentive-compatible allocations. We also derive various conditions under which this "first-order approach" is a valid means of characterising second-best allocations. The validity of the first-order-approach and the implied Relaxed problem faced by the planner is also confirmed by numerical methods. Given our discussion on sufficiency conditions, we then discuss the allocations, and the associated cross-sectional and intertemporal tax wedges, that are characterised by the necessary conditions of the Relaxed problem. We identify a new intertemporal trade-off between current and future inequality. As in the static model, a positive labour wedge (and hence lower effective labour) on some skill level reduces current inequality (at the margin). However, under learning-by-doing, such a wedge can also lead to an increase in future inequality by increasing future information rents from mimicking this skill. This trade-off implies that, for all skills above a certain (small) threshold, the labour distortion (relative to the first-best) is more progressive under learning-by-doing, ceteris paribus, when the distribution of skills is the Lognormal fitted at the modal skill level with a Pareto tail with sufficiently high Pareto index, and the mean of the Lognormal is increasing in current skill and labour with negative cross-derivative. The latter captures an environment where the positive returns from experience are higher for low skill levels, which is consistent with empirical evidence. We also identify how the tax system is affected when private utility is non-separable and/or the planner has non-Utilitarian preferences. We show that the second-best intertemporal wedge will be higher (resp. lower), all other things equal, to that set by a Utilitarian planner (and analysed by the New Dynamic Public Finance Literature), if at optimum the (conditional) expected marginal social welfare is decreasing (resp. increasing) over time, for given current skill. The reason is that in this case, current consumption is valued more (resp. less) socially. Moreover, in an economy where, at optimum, marginal utility of consumption and social marginal weights are non-increasing with current skill, and there is net skill persistence, a non-Utilitarian planner sets a higher labour wedge, across all skills, all other things equal. Under an empirically relevant condition on the skill distribution, and substitutability between leisure and consumption, learning-by-doing pushes, all other things equal, the labour wedge downwards for all interior skills.

Keywords: private information, redistribution, learning-by-doing, dynamic nonlinear taxation, first-order approach
JEL: C73, D80, D82, E60, H21

1 Introduction

We study a dynamic economy where skills/productivity evolve stochastically over time and, crucially, are subject to learning-by-doing. The latter term is used here to describe an environment where the higher current labour/effort is the more likely the worker is to have higher skill/productivity next period.

Stochastic over time skills seems to be a very natural assumption when a dynamic economy is under scrutiny, and is in line with the recent literature on optimal dynamic nonlinear taxation.[1] Our focus on learning-by-doing is motivated by the fact that the labour economics literature has established a significant effect of labour experience on wages,[2] and hence on human capital accumulation (insofar, variations in earnings reflect - to some extent - variations in skills/productivities). These effects can potentially have profound implications for *dynamic* redistribution policy. For instance, if the learning-by-doing effect is stronger for low-skilled workers (i.e. if a decrease in labour supply is more detrimental in terms of future productivity for low-skilled workers), neglecting learning-by-doing would underestimate the positive effect of low marginal taxes on low-skilled workers. Relatedly, because policy will influence labour supply, and hence the evolution of productivity *over time*, assuming away learning-by-doing may underestimate the scope for *future* redistribution. Interestingly, there is evidence (see, Dustmann and Meghir, 2005) that suggests that returns to experience and firm tenure may be higher for low-skilled workers.[3] Thus, by neglecting this dynamic aspect, the model would tend to predict a less progressive tax.

In our work, the planner can pre-commit on allocations/policies, but is constrained by the fact that workers' skills are their private information. When the planner chooses the optimal allocations, she takes into account that workers' (current) skills are their private information, and that there is scope for insurance due to idiosyncratic future skill uncertainty. The tax wedges that are associated with the optimal allocations have implications for the optimal design of labour

[1] See for instance the surveys by Golosov et al. (2006) and Kocherlakota (2010), and the more recent contributions by Farhi and Werning (2011) and Golosov et al. (2011).

[2] For instance, Willis (1986), Altug and Miller (1998) or Cossa, Heckman and Lance (2000) 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) show that job tenure affects wage profiles.

[3] Dustmann and Meghir (2005) provide evidence that the annual wage growth rate is higher for low-skilled workers after one year in employment (with the same firm), which suggests that returns to experience may be higher for low-skilled workers. They also find that the annual wage growth rate is decreasing with the number of years in employment (with the same firm) and vanishing over time. Furthermore, they find that the annual wage growth rate is *eventually* higher for workers who *started* as high-skilled workers. The latter finding, however, should not be interpreted as suggesting that the returns of longer tenure are generally higher for high-skilled workers. The reason is that the annual wage growth rates *after* the first year of employment are not estimated conditional on the annual wage growth rate in the previous year, but conditional on the skill (as approximated by apprenticeship education) when hired *and* that the workers have been in employment since. In other words, the sample of workers, over whom the returns to experience are estimated, includes also workers who have stayed in employment one more year *and* have experienced a drop (or a very small increase) in their wages. In the presence of skill persistence, however, the human capital (and hence wage) growth in the sample of workers with apprenticeship education will be higher on average over time, even if experience does not change. Therefore, the identification strategy in question may underestimate the returns to experience when being in employment for more than two years for workers who *started* as low-skill workers.

and capital income taxes. Discussing these allocations and associated wedges is the main task of this paper.

We will discuss related literature in more detail after the presentation of our model. Here we note only the major differences between the received literature on optimal nonlinear taxation and our work. The received literature consists of static models; dynamic environments where skills do not evolve stochastically over time; and dynamic models where skills do change stochastically over time. Of course, static models exclude, by construction, any learning-by-doing effects.[4] Importantly for our purposes, in the third class of existing models there is also no learning-by-doing: past employment does not affect the stochastic process of skills. Most of the models in the second class of environments also assume away learning-by-doing effects. The only exceptions we know are Krause (2009) and Balle and Spadaro (2006) where allocations are dynamic (the studied economy is multi-period) *but* skills do not evolve stochastically over time. Specifically, only initial skills are uncertain with their realisation being private information of workers, while a worker's future skill is a *deterministic function* of the initial skill and past employment. Therefore, in terms of the incentive constraints faced by the planner, the environment is effectively static - the asymmetric information is one-dimensional.

We focus on a Markov environment with standard time-separable preferences, and first derive an envelope condition that must characterise individual decisions at optimum. This condition is similar to the one derived recently in Kapicka (2008) and Pavan et al. (2009), though our condition holds for more general skill stochastic processes (conditional on a Markov set-up), and generalises the necessary condition in Mirrlees (1971). Echoing the approaches in Kapicka (2008) and Pavan et al. (2009) we then derive the allocations that would be optimal if this envelope condition was fully characterising individual decisions at optimum. We also derive various conditions under which this "first-order approach" is a valid means of characterising optimal allocations. The validity of the first-order-approach and the implied "Relaxed" problem faced by the planner is also checked by means of a numerical analysis. Given our discussion on sufficiency conditions, we then move to discuss the optimal allocations as these are characterised by the necessary conditions of the Relaxed problem of the planner.

We identify a new intertemporal trade-off between current and future (information rents and) inequality. As in the static model, a positive labour wedge (and hence lower labour) on some skill level reduces current information rents associated with this skill. However, under learning-by-doing, such a wedge can also lead to an increase in future information rents from mimicking this skill. This trade-off implies that, for all skills above a certain (small) threshold, the labour distortion (relative to the first-best) is more progressive under learning-by-doing, ceteris paribus, when the distribution of future skills is the Lognormal fitted at the modal skill level with a Pareto tail with sufficiently high Pareto index, and the mean of the Lognormal is increasing in current skill and labour with negative cross-derivative. The latter captures an environment where the positive returns from experience are higher for low skill levels, which is consistent with Dustmann

[4] See, for instance, Mirrlees (1971), Diamond (1998) and Saez (2001).

and Meghir (2005).

We also identify how the tax system is affected when private utility is non-separable and the planner has non-Utilitarian preferences. We show that the second-best intertemporal wedge will be higher (resp. lower), all other things equal, to that set by a Utilitarian planner (and analysed by the New Dynamic Public Finance Literature), if at optimum the (conditional) expected marginal social welfare is decreasing (resp. increasing) over time, for given current skill. The reason is that in this case, current consumption is valued more (resp. less) socially. In an economy where, at optimum, marginal utility of consumption and social marginal weights are non-increasing with current skill, and there is net skill persistence, a non-Utilitarian planner sets a higher labour wedge, across all skills, all other things equal. Under an empirically relevant condition on the skill distribution, and substitutability between leisure and consumption, learning-by-doing pushes, all other things equal, the labour wedge upwards for all interior skills.

The organisation of the paper is as follows. In Section 2 and 3 we present the model, and in Section 4 we discuss incentive compatibility and the approach we will deploy to characterise the constrained efficient allocations. Sections 5 and 6 derive the constrained optimum and discuss the theoretical implications for tax wedges. Section 7 investigates numerically a case that has often been discussed in the literature, in order to help us improve our understanding of the effects of learning-by-doing on optimal dynamic redistribution, and in particular on the path over time of (average) cross-sectional distortions. In Section 7 we also compute the welfare loss from resorting instead to a linear (either time-independent or time-dependent) tax and compare it with related findings in the received literature. Section 8 concludes and points to areas for future research.

2 The Model

2.1 The Environment

We consider an economy consisting of overlapping generations. Each generation consists of a continuum of agents. Each adult member of any given generation works for T_L periods and is a pensioner for the next $T - T_L$ periods.[5] Normalise, without loss of generality, the size of each generation to a unit mass. Members of a given generation differ only in terms of their production skills/types. Productivity types are independent across individuals. By the Law of Large Numbers, the realised proportion of agents of a given productivity type coincides with the prior probability that the typical agent is of this particular productivity type. Thus, in our model there is aggregate certainty, but there is individual uncertainty. Furthermore, let us restrict attention to an environment where the net contribution of generations to the fiscal budget is pre-determined. These assumptions allow us to treat every typical generation in isolation. Let thus $t = 1, ..., T$ be an index of period t in adulthood of the typical adult member of any given generation.

Assume that the gross rate of return to savings and the wage per effective unit of labour are

[5] Children are treated passively here.

exogenously given, with the latter being normalised to one. Denote the former with R. These assumptions are made to simplify the analysis and are often used in the literature.[6] Note also that there are a number of environments that feature exogenously given (time-invariant) rates of return to capital and effective labour. One example is with a production technology which is linear in capital and effective labour, and perfectly competitive markets. Another example is with a constant returns to scale production technology, perfectly competitive markets and a small open economy (with the rest of the world in steady state).

We will use the following notational conventions. Subscript t will be used to denote variables of period t, while (unless indicated otherwise) superscript t will be used to denote history of variables up to and including period t; that is, $x^t = \{x_1, .., x_t\}$. All products of measurable spaces encountered in what follows are endowed with the product sigma algebra. Furthermore, for any set S and variable x, we have $S^0 = \{\emptyset\}$ and $x^0 = \{\emptyset\}$. Finally, we will be using the conventions that $\prod_{i=k}^{k'} x_i = 1, \sum_{i=k}^{k'} x_i = 0$ for $k' < k$.

Let θ_t denote productivity in period t, with $\theta_t \in \Theta \equiv (\underline{\theta}, \bar{\theta}) \subset \mathbb{R}_{++}$, where with some abuse of notation $\bar{\theta}$ is allowed to represent infinity and $\underline{\theta}$ is allowed to be zero (in which case $\Theta \equiv \mathbb{R}_{++}$). We implicitly endow Θ_t with the Borel sigma-algebra. Let $\Theta^T \equiv \prod_{\tau=1}^{T} \Theta$, with a generic element θ. Let also $\Theta^t \equiv \prod_{\tau=1}^{t} \Theta$ denote the set of *feasible* period-t histories of productivities/skills.

Let $y_t = \{c_t, z_t\} \in Y_t \equiv [\underline{c}_t, \infty) \times Z_t$ where c_t and z_t denote, respectively, the (typical) agent's consumption and effective labour/effort (or before-tax income). Effective labour/effort, or before-tax income (used interchangeably hereafter), is the product of the period's productivity and exerted effort (or labour supply). Here, $\underline{c}_t \geq 0$. Moreover, for $t = 1, ..., T_L$ we have $Z_t \equiv \mathbb{R}_+$ while $Z_t = 0$ for $t = T_L + 1, ..., T$. This captures an implicit assumption that effort is represented by a non-negative real number and that during pensionable age individuals do not work. We implicitly endow Y_t with the Borel sigma-algebra. Let $Y^T \equiv \prod_{\tau=1}^{T} Y_\tau$ denote the set of *feasible* histories of consumption and effective labour/effort, with a generic element y. Let also $Y^t \equiv \prod_{\tau=1}^{t} Y_\tau$ be the set of feasible period-t histories of pre-tax income and consumption.

Agents' preferences, which we call *lifetime* utility, are represented by the expectation of a utility function. The lifetime utility of an agent of productivity θ given a feasible allocation y is[7]

$$U(\theta, y) \equiv \sum_{t=1}^{T} \beta^{t-1} u(c_t, z_t/\theta_t)$$

where $e_t \equiv z_t/\theta_t$ represents the necessary period-t "effort/labour" to ensure before-tax income z_t given period-t productivity θ_t, and β is the discount factor. $u(c, e)$ is a (measurable) *per-period* utility function, which is assumed to be strictly increasing in c, strictly decreasing in e, and strictly concave. We will also restrict attention to the case of $u_{ec} \geq 0$, which postulates that leisure and

[6] See for instance Golosov et. al. (2006), Golosov et al. (2011), Farhi and Werning (2011).

[7] Note that we assume that there is only one consumption good. This allows us to focus on the consumption-labour tax wedge implied by an implemented allocation while economising on notation. One could extend our set up to allow for more than one consumption goods and thereby study also commodity taxation by following steps similar to those in Golosov et al. (2003). We leave this for future work.

consumption are (weak) substitutes, which is generally considered to be the empirically relevant case (see also Golosov et al. 2011).

We focus on a Markov environment: the distribution of the period$-t$ skill depends on the period$-t-1$ skill and effort/labour level. In particular, assume that the distribution of θ_t is governed by a history-dependent probability measure $\hat{F}_t(. \mid \theta_{t-1}, \frac{z_{t-1}}{\theta_{t-1}})$ on Θ_t, where $F_1(\theta_1) \equiv \hat{F}_1(\theta_1 \mid \theta_0, \frac{z_0}{\theta_0})$, such that $\hat{F}_t(S \mid .)$ is measurable for all measurable sets $S \subset \Theta_t$. With some abuse of notation, $\hat{F}_t(. \mid \theta_{t-1}, \frac{z_{t-1}}{\theta_{t-1}})$ will also denote the period-t cumulative distribution function (cdf) corresponding to the above measure. Assume that the cdf $\hat{F}_t(\theta_t \mid \theta_{t-1}, \frac{z_{t-1}}{\theta_{t-1}})$ is non-increasing in θ_{t-1} *and* $\frac{z_{t-1}}{\theta_{t-1}}$ for any $t = 2, ..., T$. This distribution captures an environment where more productive agents are more likely to be more productive in the next period (i.e. where skills are *persistent*), *and* where *learning-by-doing* is present in the sense that more effort/labour now increases the likelihood that next period's productivity will be high. Let $F_t(. \mid \theta_{t-1}, z_{t-1}) \equiv \hat{F}_t(. \mid \theta_{t-1}, \frac{z_{t-1}}{\theta_{t-1}})$. Note that $F_{t+1}(\theta_{t+1} \mid \theta_t, z_t)$ is also non-increasing in z_t. However, the monotonicity of $F_{t+1}(\theta_{t+1} \mid ., z_t)$ depends on the relative strength of the "skill-persistence" and "learning-by-doing" effects, i.e. of $\mid \partial \hat{F}_{t+1}(\theta_{t+1} \mid \theta_t, e_t)/\partial \theta_t \mid$ and $-[\partial \hat{F}_{t+1}(\theta_{t+1} \mid \theta_t, e_t)/\partial e_t]\frac{z_t}{\theta_t^2}$, respectively. That is, $F_{t+1}(\theta_{t+1} \mid \theta_t, z_t)$ is non-increasing in θ_t if the skill-persistence effects (weakly) dominates the learning-by-doing effect (i.e. $\partial \hat{F}_{t+1}(\theta_{t+1} \mid \theta_t, e_t)/\partial \theta_t \leq [\partial \hat{F}_{t+1}(\theta_{t+1} \mid \theta_t, e_t)/\partial e_t]\frac{z_t}{\theta_t^2}$). We will assume hereafter the latter case, i.e. that there is *net skill persistence*.[8]

Let $y_t(\theta^t)$ denote the period-t allocation given skill-history θ^t, and $y^t(\theta^t) \equiv (y_1(\theta_1), ..., y_t(\theta^t))$. For a given allocation rule $y : \Theta \to Y$, where $y(\theta) \equiv y^T(\theta^T)$, define

$$W_F(y) \equiv EU(\theta, y(\theta)) \equiv$$

$$\int [\sum_{t=1}^{T} \beta^{t-1} u(c_t(\theta^t), z_t(\theta^t)/\theta_t)] dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_1(\theta_1)$$

the expected lifetime utility of the typical consumer before she learns her first-period skill. By the Law of Large Numbers, this is also equal to the average welfare of the typical generation given the allocation rule $y(\theta)$. Moreover, we have the (generational) intertemporal resource constraint[9]

$$G_F(y) \equiv \int [\sum_{t=1}^{T} \frac{z_t(\theta^t) - c_t(\theta^t)}{R^{t-1}}] dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_1(\theta_1) \geq 0$$

[8] To understand this, note that an increase in current skill, for given effective labour, implies that effort is now lower reducing thus the likelihood that next-period's productivity will be high through the learning-by-doing effect. However, due to skill persistence, the increase in the current skill makes it more likely, for given effort, that next-period's skill will be higher. We will be saying that there is net skill persistence when the latter effect dominates.

[9] This constraint could be extended to incorporate fixed revenue requirements across time easily. We choose not to do so for expositional simplicity only.

Let also

$$W_F(y \mid \theta^t) \equiv \sum_{\tau=1}^{t} \beta^{\tau-1} u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau)$$

$$+ \int [\sum_{\tau=t+1}^{T} \beta^{\tau-1} u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau)] dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

be the (expected) utility of the typical consumer given that her period-t history of productivity is θ^t and the allocation rule is $y(\theta)$.

An allocation gives rise to intertemporal and (consumption-) labour (tax) wedges defining implicitly labour and capital income taxes.[10] In our set up, the intertemporal wedge (relative to the rate of return R) is

$$\zeta_t^k(\theta^t) \equiv 1 - \frac{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)}{\beta R \int u_c(c_{t+1}(\theta^{t+1}), z_{t+1}(\theta^{t+1})/\theta_{t+1}) dF(\theta_{t+1} \mid \theta_t, z_t(\theta^t))} \tag{1}$$

while the consumption-labour wedge is

$$\zeta_t^z(\theta^t) \equiv 1 + \frac{u_e(c_t(\theta^t), z_t(\theta^t)/\theta_t)}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)\theta_t} + \frac{\hat{\delta}(\theta^t, z^t(\theta^t))}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)} \tag{2}$$

where

$$\hat{\delta}(\theta^t, z_t(\theta^t)) \equiv$$

$$\frac{\partial}{\partial z_t(\theta^t)} \int [\sum_{\tau=t+1}^{T} \beta^{\tau-t} u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau)] dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

Note that the latter wedge takes into account that a marginal change in period-t labour supply has a learning-by-doing effect on the distribution of future skills and hence future expected welfare for given continuation allocation $(y_{t+1}(\theta^{t+1}), ..., y_T(\theta^T))$ and skill-history θ^t.

We assume that the typical agent's productivity/skill in period t, θ_t, is her private information. This will be referred to as the agent's period-t (productivity/skill) type, while θ^T will be referred to as simply her (productivity/skill) type. We assume that the planner can verify/enforce effective labour and consumption.[11] In the contract theory jargon, allocations in Y are *contractible*.

Finally, we assume that the planner can commit to an allocation in Y at the start of the

[10] The intertemporal and labour wedges are the marginal taxes consumer would face if the allocation was implemented via appropriately set per-unit non-linear savings and labour income taxes. The intertemporal (resp. labour) wedge is the difference between the (absolute values of the) rate of transformation and the marginal rate of substitution between current consumption and future consumption (resp. effective labour).

[11] This requires that the planner can also verify/enforce the rate of return to savings and, crucially, private savings. We leave the study of an environment as the one here but with private savings being out of the planner's control for future work.

first period, after the first-period types have been realised. This is certainly not without loss of generality. As is known from contract theory, a planner could find in its interests to alter the allocation if some of agents' information is revealed. Nevertheless, echoing the dynamic Ramsey literature and most of the literature on dynamic nonlinear taxation, the full-commitment case is a useful benchmark (if not more), and we leave the study of the non-commitment case for future work.

2.2 Incentive Compatibility

Given that a generation consists of a continuum of consumers and skills are independent across consumers we can without loss of generality cast the problem of determining optimal allocations from the point of view of the planner (where optimality is defined shortly) in terms of a principal-agent model where the principal is the planner and the agent is the typical consumer of the generation under scrutiny. Suppose thus that the planner/principal implements an *outcome* by means of committing at the start of the first period, after the first-period types have been realised, on a deterministic[12] *mechanism* that assigns in each period t to the agent a set of possible *messages* M_t and a *policy rule* $\Omega_t : M^t \times Y^{t-1} \to Y_t$, *with* $M \equiv < M_t >_{t=1}^{T}$ and $\Omega \equiv < \Omega_t >_{t=1}^{T}$, such that (i) $M^t \equiv \prod_{\tau=1}^{t} M_\tau$, (ii) $(\Omega_t(m^t, y^{t-1}), y^{t-1}) \in Y^t$ with $m^t \in M^t$, and (iii) for all t, $\Omega_t(.)$ is measurable.

The timing of events in each period t, given a mechanism $< M, \Omega >$ and any history $(\theta^{t-1}, m^{t-1}, y^{t-1})$ is: first, θ_t is realised and observed by the agent; subsequently, the agent reports her message $m_t \in M_t$; and, finally, the mechanism selects a period-t allocation according to $y_t = \Omega_t(m^t, y^{t-1}) \in Y_t$.

Given a history $(\theta^t, m^{t-1}, y^{t-1})$, the agent thus responds with a period$-t$ *reporting strategy* $\sigma_t : \Theta^t \times M^{t-1} \times Y^{t-1} \to M_t$.[13] Call $\sigma \equiv < \sigma_t >_{t=1}^{T}$ the agent's reporting strategy. Let also $\sigma^t \equiv < \sigma_\tau >_{\tau=1}^{t}$.

The Revelation Principle still applies here (adapted by Myerson, 1986), and so for any standard optimality concept, any equilibrium outcome that can be induced by an arbitrary mechanism (as defined above) can also be induced as an equilibrium outcome of a *direct mechanism* in which the agent is asked to report the current type (i.e. $M_t = \Theta_t$), and where, in each period, the agent reports truthfully *conditional* on having reported truthfully in the past (i.e. $\sigma_t(\theta^t, \theta^{t-1}, y^{t-1}) = \theta_t$ for all t and all histories $(\theta^t, \theta^{t-1}, y^{t-1})$). Therefore it comes without loss of generality to restrict attention to such *incentive-compatible direct mechanisms*. We do so in what follows and we call any such mechanism $< \Theta, \Omega >$, with some abuse of terminology, as simply a *mechanism* Ω. Call

[12] In general, mechanisms could be stochastic, specifying lotteries over allocations. We will however restrict attention (following the received literature on optimal taxation) to deterministic mechanisms. This can be justified with reference to complexity costs in administering in reality in large societies a randomised allocation.

[13] Restricting attention to pure strategies is without loss of generality. By the Revelation Principle we can restrict attention to the agent reporting truthfully in equilibrium, while if there was a profitable mixed strategy deviation then there would also be a profitable deviation to a pure strategy in the support of the mixed strategy deviation. Therefore, the absence of profitable pure-strategy deviations ensures also the absence of profitable mixed-strategy deviations along the equilibrium play.

truthful strategy the strategy with $\sigma_t(\theta^t, \theta^{t-1}, y^{t-1}) = \theta_t$ for all t and all histories $(\theta^t, \theta^{t-1}, y^{t-1})$. Denote this strategy with an asterisk.

Define with H the set of all histories from the agent's perspective when she contemplates the reporting strategy: $H \equiv \cup_{t=0}^{T} H_t$ with $H_t \equiv \{\Theta^t \times M^{t-1} \times Y^{t-1}\} \cup \{\Theta^{t-1} \times M^{t-1} \times Y^{t-1}\}$. Denote a generic element of H_t with h_t and of H with h. A history $h_t = (\theta^s, m^{t-1}, y^{t-1})$, $s = t, t-1$ is a *truthful history* if $m^{t-1} = \theta^{t-1}$. Note that a mechanism Ω, a strategy σ and a history $h_t = (\theta^t, m^{t-1}, y^{t-1})$ determine a continuation allocation rule (indexed by $\Omega, \sigma | h_t$) $y^{\Omega,\sigma|h_t} : h_t \times \prod_{\tau=t+1}^{T} \Theta \to \prod_{\tau=t}^{T} Y_\tau$. Let then (with some abuse of notation), $W_F(\Omega \circ \sigma \mid h_t)$ denote the associated expected lifetime utility of the typical consumer given history h_t, the mechanism Ω and the strategy σ. [14]

In the most of the received literature on dynamic taxation under idiosyncratic skills uncertainty, it is ex-ante IC that is often used, focusing thus on optimal risk sharing. A recent exception is Battaglini and Coate (2008) where the focus is on IC-θ_1 for all (truthful) histories (θ_1). In the "static" Mirrlees tradition, on the other hand, where there is only one period, it is IC-θ_1 for all (truthful) histories $h_1 = \theta_1$ that is adopted.[15] Echoing the latter tradition, we adopt the more stringent optimality criterion of IC$-h_t$ for all *truthful* histories h_t *for all periods* t. Apart from its natural appeal given our focus on dynamic redistribution, the technical reason for this choice is that it allows us to determine optimal allocations recursively through the solutions of a sequence of optimal control problems that are similar, but not the same, to the problem studied in Mirrlees (1971). Note however that these two criteria differ only on a set of truthful histories of a zero measure.[16]We will refer hereafter to an allocation in Y that can be induced by a a mechanism which is IC$-h_t$ for all truthful histories h_t for all periods t as simply an *incentive compatible* allocation.

Following the terminology in Pavan et al. (2009), a direct mechanism Ω is *incentive compatible* at history h (IC-h hereafter) if

$$W_F(\Omega \circ \sigma^* \mid h) = sup_\sigma W_F(\Omega \circ \sigma \mid h)$$

That is, if truth-telling is optimal for the agent given history h. *An IC-h_0* mechanism will be referred to as *ex-ante incentive compatible* (ex-ante IC hereafter).

In most of the received literature on dynamic taxation under idiosyncratic skills uncertainty, it is ex-ante IC that is used, focusing thus on optimal risk sharing. A recent exception is Battaglini and Coate (2008) where the focus is on IC-θ_1 for all (truthful) histories (θ_1). In the "static" Mirrlees tradition, on the other hand, where there is only one period, it is IC-θ_1 for all (truthful) histories $h_1 = \theta_1$ that is adopted.[17] Echoing the latter tradition we adopt the more stringent

[14]Thus, $W_F(\Omega \circ \sigma^* \mid \theta_1) = W_F(y^{\Omega,\sigma^*|\theta_1} \mid \theta_1)$. Moreover, after noting that h_0 is the *null* history (i.e. the history prior to the first-period type being realised), we have $W_F(\Omega \circ \sigma^* \mid h_0) = W_F(y^{\Omega,\sigma^*|h_0})$.

[15]See, for instance, Diamond (1998) and Saez (2001), where $T = T_L = 1$.

[16]The Principle of Optimality implies that a mechanism that is ex-ante IC is IC for almost all truthful histories and all periods, while a mechanism that is IC for all truthful histories and all periods is (trivially) also IC for the null history and thereby ex-ante IC. See also Kapicka (2008) and Pavan et al. (2009).

[17]See, for instance, Diamond (1998) and Saez (2001) for $T = T_L = 1$.

optimality criterion of IC$-h_t$ for all *truthful* histories h_t *for all periods* t. Apart from its natural appeal given our focus on redistribution, the technical reason for this choice is that it allows us to determine optimal allocations recursively through the solutions of a sequence of optimal control problems that are similar, but not the same, to the problem studied in Mirrlees (1971). Note however that these criteria differ only on a set of truthful histories of a zero measure.[18] We will refer hereafter to an allocation in Y that can be induced by a a mechanism which is IC$-h_t$ for all truthful histories h_t for all periods t as simply an *incentive compatible* allocation.

2.3 Statement of the Problem

In this paper we are concerned with the *constrained efficient* (second-best) allocations. These are defined here as the allocations that maximise the social welfare function

$$\int \sum_{t=1}^{T} \beta^{t-1} Q(u(c_t(\theta^t), z_t(\theta^t)/\theta_t)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}))...dF_1(\theta_1)$$

subject to the resource constraint $G_F(y) \geq 0$ and the incentive-compatibility constraints that for all truthful histories h_t and all periods t, $W_F(\Omega \circ \sigma^* | h_t) \geq W_F(\Omega \circ \sigma | h_t)$ for all strategies σ.[19]

In the above, $Q(.)$ is a strictly increasing and (weakly) concave function. Note that the above social welfare function includes as a special case the Utilitarian criterion $W_F(y)$ (when $Q(v) = v$). In the Utilitarian case, every agent is treated symmetrically from an ex ante perspective, and the focus is, in effect, on providing efficient insurance to consumers given that (future) stochastic skills are the workers' private information. In general, $Q(.)$ captures the redistributional considerations of the planner. The more concave Q, the higher the aversion of the planner to inequality is.[20] One way to justify our assumption that the planner is averse to inequality on a *period-per-period* basis is with reference to work on political economy of taxation and redistribution such as Piketty (1995) who shows that society's attitudes toward redistribution evolve across time as a result of the interaction between citizens' ideologies and learning based on past histories. Another reason is that this form of redistributional concerns allows us to isolate the role of learning-by-doing for optimal redistribution by facilitating a direct comparison with studies of static environments that have used a similar welfare function (see, for instance, Mirrlees, 1971, Diamond, 1998, and Saez, 2001).

[18] The Principle of Optimality implies that a mechanism that is ex-ante IC is IC for almost all truthful histories and all periods, while a mechanism that is IC for all truthful histories and all periods is (trivially) also IC for the null history and thereby ex-ante IC. See also Kapicka (2008) and Pavan et al. (2009). Similarly for the comparison between our criterion and IC-θ_1 for all θ_1.

[19] Note than in general there will be more than one savings (cross-section and across time) profiles of private savings, public debt holdings and tax revenues that are consistent with a second-best allocation. The approach of determining the profiles of second-best consumption and effective labour, and subsequently the implications for tax wedges, we adopt here is common in the literature on optimal taxation, see for instance Chari and Kehoe (1999) and Golosov et al. (2003), and the discussions therein.

[20] As an example, consider the function $Q(u) = u^\delta/\delta$, $\delta \neq 0, \delta \leq 1$. This implies the Utilitarian criterion when $\delta = 1$, the Rawlsian criterion (of maximising the utility of the individuals with the lowest productivity) when $\delta \to -\infty$, while it implies a logarithmic social welfare when $\delta \to 0$.

2.4 Technical Assumptions

Denote with $C(u,z,\theta)$ the consumption required in period t to ensure a period-t utility of u, given period-t skill θ, and before-tax income z - that is, $C(u,z,\theta)$ is the expenditure function defined as the unique solution with respect to C of $u = u(C, z/\theta)$.

We will be adopting at various stages of our analysis the following assumptions:

AU1 $\frac{z}{\theta^2}u_e(c,\frac{z}{\theta})$ exists and is uniformly bounded in (θ, y).

AU2 u is twice continuously differentiable.

AU3 $\frac{z}{\theta^2}u_e(C(u,z,\theta),z/\theta)$ and $-C(u,z,\theta)$ are concave in (u,z) for any $\theta \in \Theta$.

AF1 $\int \theta_t dF_t(\theta_t \mid \theta_{t-1}, z_{t-1}) < \infty$ for any $(\theta_{t-1}, z_{t-1}) \in \Theta \times Z_t$ and any t.

AF2 $F_t(. \mid \theta_{t-1}, z_{t-1})$ is strictly increasing on Θ, for any $(\theta_{t-1}, z_{t-1}) \in \Theta \times Z_t$ and any t.

AF3 $\partial F_t(. \mid \theta_{t-1}, z_{t-1})/\partial\theta_{t-1}$ is a.e. differentiable on Θ, for any $(\theta_{t-1}, z_{t-1}) \in \Theta \times Z_t$ and any t.

AF4 $f_t(. \mid \theta_{t-1}, .)$ and $\partial f_t(. \mid \theta_{t-1}, .)/\partial\theta_{t-1}$ are twice continuously differentiable on $\Theta \times Z_t$, for any $\theta_{t-1} \in \Theta$ and any t.

AF5 $\partial F_t(\theta_t \mid \theta_{t-1}, z_{t-1})/\partial\theta_{t-1}$ exists, for any $(\theta_t, \theta_{t-1}, z_{t-1}) \in \Theta \times \Theta \times Z_t$ and any t.

A6(i) $\mid \partial F_t(\theta_t \mid \theta_{t-1}, z_{t-1})/\partial\theta_{t-1} \mid \leq B_t(\theta_t)$ for some integrable function $B_t : \Theta \to \mathbb{R}_+ \cup \{\infty\}$, for any $(\theta_t, \theta_{t-1}, z_{t-1}) \in \Theta \times \Theta \times Z_t$ and any t.

AF6(ii) (ii) $\partial F_t(\theta_t \mid \theta_{t-1}, z_{t-1})/\partial\theta_{t-1}$ is monotone in θ_{t-1}, for any $(\theta_t, \theta_{t-1}, z_{t-1}) \in \Theta \times \Theta \times Z_t$ and any t.

AF6(iii) $\int \theta_t dF_t(\theta_t \mid \theta_{t-1}, z_{t-1})$ is locally equi-Lipschitz continuous in θ_{t-1}. That is, for all $z_{t-1} \in Z_{t-1}$, there exists L such that $\mid \int \theta_t d[F_t(\theta_t \mid \theta_{t-1}, z_{t-1}) - F_t(\theta_t \mid \eta_{t-1}, z_{t-1})] \mid \leq L \mid \theta_{t-1} - \eta_{t-1} \mid$ for all θ_{t-1} and η_{t-1} in every compact subset of Θ.

AF6(iv) $lim_{\theta \to \overline{\theta}} \theta \frac{\partial F_t(\theta|\theta_{t-1}, z_{t-1})}{\partial \theta_{t-1}} = 0$ for all $z_{t-1} \in Z_{t-1}$ and all θ_{t-1} in every compact subset of Θ.

The following remarks on the above assumptions are in order here. Note that AU1 implies that $u(c, z/\theta)$ is equi-Lipschitz continuous (and hence almost everywhere differentiable) in θ. AF2 is a full-support assumption and AF5 is a smoothness assumption. AU1 and AF1, AF6(i)-AF6(iii) are technical assumptions needed, alongside AF2 and AF5, to prove an envelope theorem that echoes the ones proved in Milgrom and Segal (2002), Kapicka (2008) and Pavan et al. (2009). Note that AF1 is also present in Mirrlees (1971) and is trivially satisfied if $\overline{\theta}$ is finite, and AF2 and AF5 are commonly used in the literature. AU2, AU3 AF3 and AF4 impose sufficient additional smoothness, and, alongside AF6(iv), will enable us to use recursive optimal control methods to characterise the optimum in a way reminiscent of that in Kapicka (2008), Farhi and Werning (2011) and Golosov et al. (2011).

We note here that Pavan et al. (2009) prove their envelope theorem under AU1, AF1, AF2, AF5 and AF6(i) (and without AF6(ii)-AF6(iv)).[21] However, AF6(i) is not satisfied by many parametarised distributions like the Lognormal and Pareto distributions, when their mean is given by a function of θ_{t-1} and z_{t-1} that has an unbounded partial derivative with respect to θ_{t-1} as $\theta_{t-1} \to 0$. Lognormal and Pareto distribution distributions play a central role in discussions of optimal income taxation and do/can satisfy AF1-AF6(iv) with the appropriate choice of parameters. Kapicka (2008), on the other hand, proves his envelope theorem under AU1, AF1, AF2, AF5 and AF6(ii)-(iv),[22] but his proof cannot be applied here because it relies on the absence of learning-by-doing. Here, by combining insights from these two papers we succeed in proving an envelope theorem that can also be used under the Lognormal and Pareto distributions with mean - for example - $\Xi_t(\theta_{t-1}, z_{t-1}) = \xi + \rho^{-1}\theta_{t-1}^{\rho} + log(\theta_{t-1} + 2P(\frac{z_{t-1}}{\theta_{t-1}})), \xi \geq 1, \rho \in (0,1), \rho \to 0$, where $P(e)$ is a non-negative, increasing, and bounded function.[23]

2.5 Related Literature

In addition to the work cited already, we can now clarify how our work is related to existing work on optimal dynamic non-linear taxation. Note here that with $T_L = T$, and $F_t(. \mid \theta_{t-1}, z_{t-1})$ being the Dirac measure for any $t > 1$, our model becomes similar to those in Balle and Spadaro (2006) and Krause (2009). Balle and Spadaro (2006) study a continuous time model with quasi-linear in consumption per-period preferences and where the planner is allowed to have different time preferences than consumers. Krause (2009) study a two-period model with two possible (initial) productivity types, additively-separable per-period utility and a Utilitarian planner. Moreover, both models do not allow for private savings (and Krause (2009) does not allow also for public debt).

Our model reduces to a model similar to those used in Kapicka (2008), Battaglini and Coate (2008), Golosov et al. (2011) and Farhi and Werning (2011) when $\partial F_t(. \mid \theta_{t-1}, z_{t-1})/\partial z_{t-1} = 0$ for

[21] Note that AF1 and A6(i) imply that $\int \theta_t dF_t(\theta_t \mid \theta_{t-1}, z_{t-1})$ is equi-Lipschitz continuous in θ_{t-1}, which in turn implies A6(iii). However, as we discuss shortly, A6(ii) is needed in our case for when A6(i) is not satisfied (see Appendix A for more details).

[22] In fact, AF6(iv) is never explicitly stated in Kapicka (2008), but is needed to derive the last step in the proof of the first part of Lemma 2 in that paper. Moreover, he uses a slightly more stringent assumptions than AF1 and AF6(iii) (see his Assumption 3 and 7 - or 3 and 8 in the 2011 version of the paper), but he can also derive his results with AF1 and AF6(iii) instead.

[23] Note that with $P(x) \equiv \frac{1}{1+e^{-x}} \in [1/2, 1)$ - the standard logistic function, defined on $\mathbb{R}_+$ - the derivative of $\rho^{-1}\theta^{\rho} + log(\theta + 2P(\frac{z_{t-1}}{\theta}))$ with respect to θ is equal to $\theta^{\rho-1} + \frac{\theta - (z/\theta)2P'(z/\theta)}{\theta^2 + \theta 2P(\frac{z}{\theta})}]$ where $P'(x) = P(x)(1 - P(x)) \in [0, 0.25]$. Clearly, then, for $\rho \to 0$, the derivative in question is positive, and tends to infinity as $\theta \to 0$ (note that $P(x) - xP'(x) \in [0.5, 1)$. One can also after some tedious algebra that $\rho^{-1}\theta^{\rho} + log(\theta + 2P(\frac{z}{\theta}))$ is concave with respect to θ for $\rho \to 0$. [The second partial derivative with respect to θ converges to $-\theta^{-2}\{1 - \frac{2(z/\theta)P'(z/\theta)[2+(z/\theta)(1-2P(z/\theta))]}{\theta + 2P(z/\theta)}\} - [\frac{1-2(z/\theta)P'(z/\theta)}{\theta+2P(z/\theta)}]^2$. A diagram of this function suffices then to show that it is negative for any $(\theta, \frac{z}{\theta}) \in \mathbb{R}_+^2$.] This implies that AF1-AF5 and AF6(ii)-AF6(iv) can be satisfied for the Pareto (with sufficiently high index $\alpha > 1$) and Lognormal distributions, with their mean being identified with $\xi + \rho^{-1}\theta_{t-1}^{\rho} + log(\theta_{t-1} + P(\frac{z_{t-1}}{\theta_{t-1}})), \xi \geq 0, \rho \to 0^+$ (for fixed scale parameters). [Satisfaction of A6(ii) follows from the fact that (for fixed scale parameters) the Lognormal cdf is decreasing and convex in its mean, while the Pareto cdf is decreasing in its mean and the ratio of its second to first derivative with respect to its mean equals $-(1-\alpha)$ divided by the mean.]

all t. All these models use a Markov environment. In the first two works, we have $T_L = T \equiv \infty$, in the third work $T_L = T < \infty$ while in the last work $T_L \leq T < \infty$. In the second paper there are only two possible types. Moreover, in all papers except Golosov et al. (2011) the planner has Utilitarian preferences, while in the latter the planner may be averse only in initial inequality; that is, the planner's preferences are given by $\int Q(W_F(y, \theta_1))dF_1(\theta_1)$ (in terms of our notation). Note that for $Q(v) = v$ the latter preferences reduce to the Utilitarian preferences as well. Finally, in Battaglini and Coate (2008) per-period private preferences are quasi-linear in consumption, while in Kapicka (2008) they can be re-written as an affine function of the period's allocation after an appropriate reparametarisation of per-period types and allocations. The latter would also be the case here under additively-separable private utility with iso-elastic labour cost function.

3 A Recursive Representation of the Resource Constraint and the Full-Information Tax Wedges

Given an allocation y, let us denote by $\widehat{G}_{t+1}(\theta^t)$ the period-$t+1$ net (generational) liability[24] that is financed in period $t+1$ by all consumers of past skill-history θ^t (of the generation under scrutiny). An admissible path of these net liabilities, given type θ^T, is such that

$$\widehat{G}_t(\theta^{t-1}) =$$

$$\int [z_t(\theta^t) - c_t(\theta^t) + \frac{\widehat{G}_{t+1}(\theta^t)}{R}]dF_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) \tag{3}$$

with $\widehat{G}_{T+1}(\theta^T) \leq 0$ and $\widehat{G}_1(\theta^0) \equiv 0$. In fact, note that backward iteration of the above law-of-motion implies the resource constraint $G_F(y) \geq 0$. In addition, given an allocation y *that satisfies the resource constraint* $G_F(y) \geq 0$, one can set $\overline{G}_{t+1}(\theta^t), t = 0, ..., T$, such that

$$\frac{\overline{G}_{T+1}(\theta^T)}{R^{T+1}} = -G_F(y)$$

and, for any $t = 0, ..., T-1$,

$$\overline{G}_{t+1}(\theta^t)/R = \int \{\sum_{\tau=t}^{T-1} [z_{\tau+1}(\theta^{\tau+1}) - c_{\tau+1}(\theta^{\tau+1})][R^{\tau-t+1}]^{-1} +$$

$$+ \frac{\overline{G}_{T+1}(\theta^T)}{R^{T-t+1}}\} dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

and thereby ensure that (3) above is satisfied with $\widehat{G}_{t+1}(\theta^T) = \overline{G}_{t+1}(\theta^t)$, for any $t = 0, ..., T$ (note that by construction $\overline{G}_{T+1}(\theta^T) \leq 0$). This two-sided mapping implies, alongside the definition of optimum, that y is a constrained efficient allocation if and only if it is the solution of the optimisation problem that emerges after replacing the resource constraint in the second-best

[24] That is, the value of public debt minus net asset holdings (by the generation in question).

problem with the net-liability law-of-motion (3).[25] It will be convenient to determine the second-best allocations through the latter problem. Hereafter, we will be referring to $z_t(\theta^t) - c_t(\theta^t) + \frac{\widehat{G}_{t+1}(\theta^t)}{R}$ as the planner's period-t *revenues* from individuals with period-t skill-history θ^t.

3.1 The Full-Information Tax Wedges

It will be instructive to discuss at this stage the benchmark case of skills being common knowledge. The full-information allocations are defined as the allocations that maximise the social welfare function subject to the the net-liability law-of-motion (3) (or equivalently to the resource constraint).

One can then show that[26]

$$\frac{1}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)} - \frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} =$$

$$+ \frac{1}{pq(t)}\{Q'(u(c_t(\theta^t), z_t(\theta^t)/\theta_t) - Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})\} \tag{4}$$

holds.

With a Utilitarian planner ($Q' \equiv 1$), this reduces to the standard Euler condition. Note that the inverse of the marginal utility of consumption represents the increase in fiscal revenues generated by reducing utility marginally for given earnings. The Euler condition simply ensures that, at optimum, reducing (for given earnings profile) the utility of a consumer with skill history θ^{t-1} by $\beta\varepsilon$ (and saving the generated revenues) in exchange of increasing her future utility under skill θ_t by ε, with $\varepsilon \to 0$, does not free up resources in period t, and hence is not welfare improving. It implies a zero intertemporal tax wedge: $\zeta_t^k(\theta^t) = 0$ (for $1 < t \leq T$), as intuition would suggest.

With a non-Utilitarian planner, the reallocation of utils in question leads (in terms of period-t revenues) to a net decrease in period-t revenues (equal to the left-hand side of (4)), as well as a net increase in social welfare in period-t units (equal to the right-hand side of (4)).[27] Thus, at optimum, the marginal social benefit in revenue-units must be equal to the marginal revenues cost, which is what the above condition states. Note now that this condition implies the (gross) intertemporal wedge:

[25] The proof is by contradiction, and uses that any feasible allocation in one problem corresponds to some feasible allocation in the other problem.

[26] The problem analysed in Appendix C reduces to the full-information problem after setting $\mu(.) = \gamma(.) \equiv 0$. Thus, (4) and (5) in the main text are simply the counterparts of conditions (36) and (37) in Appendix C.

[27] Reducing period-$t-1$ utility by β units lead to a decrease in social welfare by $\beta Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})$. Dividing it by β expresses this in terms of period-t utils, and dividing the curly brackets in (4) by the shadow price of period$-t$ revenues expresses the net social welfare change in terms of period$-t$ revenues.

$$\frac{\zeta_t^k(\theta^t)}{1-\zeta_t^k(\theta^t)} = -\frac{1}{pq(t)}\times$$

$$\int_{\underline{\theta}}^{\overline{\theta}} u_c(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)Q'(u(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)dF_{t+1}(x \mid \theta_t, z_t(\theta^t))$$

$$-Q'(u(c_t(\theta^t), z_t(\theta^t)/\theta_t)\} \int_{\underline{\theta}}^{\overline{\theta}} u_c(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)dF_{t+1}(x \mid \theta_t, z_t(\theta^t))$$

(for $1 < t < T$). Therefore, the implied full-information intertemporal wedge is positive if and only if

$$Q'(u(c_t(\theta^t), z_t(\theta^t)/\theta_t) > \int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)dF_{t+1}^u(x \mid \theta_t, z_t(\theta^t)) \tag{5}$$

and vice versa, where

$$f_{t+1}^u(x \mid \theta_t, z_t(\theta^t)) = \frac{u_c(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)f_{t+1}(x \mid \theta_t, z_t(\theta^t))}{\int_{\underline{\theta}}^{\overline{\theta}} u_c(c_{t+1}(x,\theta^t), z_{t+1}(x,\theta^t)/x)dF_{t+1}^u(x \mid \theta_t, z_t(\theta^t))}$$

is the conditional density of period-$t+1$ skill x adjusted by the wealthiness of type (x, θ^t) as this is measured by the corresponding marginal utility of consumption relative to the average marginal utility of consumption. In other words, the implied full-information intertemporal wedge in period t, given skill-history θ^t, is positive (pushing for lower savings) if and only if (at the margin) the period-t social welfare loss (in terms of period$-t+1$ utils) is greater than the period-$t+1$ (wealthiness-adjusted) average social welfare gain, which arise due to a decrease in period-t utility by β units and an increase in period-$t+1$ utility (regardless of period-$t+1$ skill) by one unit.

This is in itself an interesting result. The common intuition that seems to emerge from the received literature is that the optimal intertemporal wedge should be zero except in a second-best allocation with idiosyncratic skill/productivity uncertainty. In fact, Chamley (1986) shows in a representative-agent model that at steady state the (second-best) optimal linear capital income tax is zero. Building on insights from Diamond and Mirrlees (1978) and Rogerson (1985), the New Dynamic Public Finance literature has shown that the second-best intertemporal wedge is likely to be positive when individuals face idiosyncratic skill uncertainty; see, for instance, Golosov et al. (2003), Kocherlakota (2005) and the discussion in the survey paper by Golosov et al. (2006). However, this literature assumes a Utilitarian planner.[28] Moreover, Judd (1985), who allows for non-Utilitarian redistributional motives, focuses on (steady-state) linear taxes in a model with

[28] Or, to the same effect, as in Golosov et al. (2003), a planner who weighs individuals' expected utilities (conditional on their initial skill). Indeed, setting $Q'(.) \equiv \chi$ (for all skill histories with initial skill θ_1) in (4) we can see that our model delivers the same prediction as well.

two productivity types, where all uncertainty is resolved in the first period, and also delivers the zero-capital-tax result. Here, however, we show that *non-Utilitarian* redistributional incentives may lead to non-zero intertemporal wedges, even when skills are commonly known and there is idiosyncratic skill uncertainty. Note that this does not rely on the presence of learning-by-doing.

Regarding the full-information effective labour, we have that

$$-\frac{\partial \hat{Q}_{t+1}(\theta_t, z_t(\theta^t), .)}{\partial z}\frac{\beta}{pq(t)} = 1 + \frac{u_e(.)}{\theta u_c(.)}$$

holds. This condition states that the period-t effective labour of skill history θ^t must be such that reducing it marginally leads to a net reduction in planner's revenues from all taxpayers with skill-history θ^t (the right-hand side) which equals the expected discounted net social marginal welfare gain (expressed in terms of period-t public funds) from making low future skills more likely for them (the left-hand side).

This condition implies that the labour wedge is

$$\zeta_t^z(\theta^t) = \frac{\hat{\delta}(\theta^t, z_t(\theta^t))}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)} - \frac{\partial \hat{Q}_{t+1}(\theta, , z(\theta), .)}{\partial z}\frac{\beta}{pq(t)}$$

As we show in Appendix E, with a Utilitarian planner,

$$\frac{\partial \hat{Q}_{t+1}(\theta, ., z(\theta), .)}{\partial z}\frac{\beta}{pq(t)} = \frac{\hat{\delta}(\theta^t, z_t(\theta^t))}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)}$$

$$+\frac{\partial}{\partial z_t(\theta_t)}\int \sum_{\tau=t+1}^{T} \beta^{\tau-t-1}\frac{Q(u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau))}{pq(t)/\beta} dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta_{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta_t))$$

$$-\frac{Q'(u(c_t(\theta_t), .))}{pq(t)/\beta}\int v_{t+1}(\theta^{t+1})\frac{\partial f_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial z_t(\theta_t)} d\theta_{t+1}$$

$$+R^{-1}\int [z_{t+1}(\theta^{t+1}) - c_{t+1}(\theta^{t+1}) + \frac{\widehat{G}_{t+2}(\theta^{+1})}{R}]\frac{\partial f_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial z_t(\theta_t)} d\theta_{t+1}$$

To see the intuition behind this note that a marginal increase in period-t (effective) labour affects the conditional distribution of future skills and, specifically, makes future skills more likely to be high. This has the following effects (measured in terms of period-t revenues) for given future and past per-period utilities/consumption. First, it increases expected revenues by the fourth term in the right-hand side of the above equation. Second, it increases expected discounted social welfare the third term. Third, it increases the expected discounted continuation utility of consumers of type θ^t by $\beta\frac{\partial}{\partial z_t(\theta_t)}\int v(\theta^{t+1})dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$, which is equal to $\widehat{\delta}(\theta_t, z_t(\theta^t))$. This allows the planner to generate more revenues in period t by reducing their period-t utility by $\beta\frac{\partial}{\partial z_t(\theta_t)}\int v(\theta^{t+1})dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$ units by means of an appropriate reduction in *current* (i.e. period$-t$) consumption. This, in turn, implies the following two effects. On the one hand, there is a decrease in current social welfare that equals the second term above. On the other

hand, there is an increase in period-t revenues by the first term at the right-hand side of the equation above.

Consequently, the full-information labour wedge is equal to

$$\zeta_t^z(\theta^t) =$$

$$-(\beta/pq(t)) \times \tag{6}$$

$$\frac{\partial}{\partial z_t(\theta_t)}\{\int \sum_{\tau=t+1}^{T} \beta^{\tau-t-1} Q(u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta_{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta_t))$$

$$-Q'(u(c_t(\theta_t), .)) \int v_{t+1}(\theta^{t+1}) dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta_t))\}$$

$$-R^{-1} \int [z_{t+1}(\theta^{t+1}) - c_{t+1}(\theta^{t+1}) + \frac{\widehat{G}_{t+2}(\theta_t, \theta_{t+1})}{R}] \frac{\partial}{\partial z_t(\theta_t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

Thus, as intuition would suggest, a Utilitarian planner would set a positive labour wedge if the conditional expected present value of planner's revenues from individuals of skill-history θ^t is decreasing in its effective labour, and vice versa. This is represented by the second term above, and highlights an intertemporal redistribution trade-off faced by the planner when learning-by-doing is present: current and future revenues may be substitutes calling for a positive labour wedge. An additional intertemporal trade-off arises under non-Utilitarian preferences *and* learning-by-doing, and is represented by the first term above. Therefore, the full-information labour wedge will be higher under non-Utilitarian preferences, all other things equal, if the net social welfare gain from higher skills becoming more likely and exploiting the associated increase in expected private utility to generate revenues in period t is negative.

We now move to the focus of our study: the case of skills being workers' private information. Moving to the private information case brings up additional aspects of the intertemporal redistribution tensions that emerge when learning-by-doing is present.

4 Necessary and Sufficient Conditions for Incentive Compatibility

Incentive-compatibility constraints are significantly many to render a full characterisation of second-best allocations very difficult. We will therefore follow the so-called first-order approach whereby we study a relaxed problem and identify conditions that ensure that the derived allocations satisfy the constraints of the full problem. The *Relaxed Program* is the one that emerges from the planner's full problem after replacing all incentive constraints with *all* the *local incentive-compatibility constraints*.[29] These constraints are discussed next. To simplify exposition, drop hereafter the superscripts Ω and σ^* whenever there is no risk of confusion.

[29] See also Kapicka (2008), Pavan et al. (2009), Golosov et al. (2011) and Farhi and Werning (2011). A related approach is also adopted in Battaglini and Coate (2008).

4.1 Necessary Conditions

Let, for all t and all truthful histories $h_t = (\theta^t, \theta^{t-1}, y^{t-1})$,

$$v(h_t) \equiv u(c_t(\theta^t), z_t(\theta^t)/\theta_t) +$$

$$\int \sum_{\tau=t+1}^{T} \beta^{\tau-t} u(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t)) \quad (7)$$

or

$$v(h_t) \equiv u^t(c_t(\theta^t), z_t(\theta^t)/\theta_t) + \beta \int v(h_{t+1}) dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t)) \quad (8)$$

That is, $v(h^t)$ is the period-t value function, given truthful history h^t. Echoing similar results in Pavan et al. (2009) and Kapicka (2008), we show in Appendix A that:

Proposition 1 *Under AU1, AF1, AF2, AF5, and AF6(i) and/or AF6(ii)-(iii), if a mechanism is IC–h_t for a truthful history $h_t = ((\theta_t, \theta^{t-1}), \theta^{t-1}, y^{t-1})$, then $v((\theta_t, \theta^{t-1}), \theta^{t-1}, y^{t-1})$ is equi-Lipschitz continuous in θ_t, and for all truthful histories $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$ and almost all θ_t, the following* local incentive-compatibility constraint *(local IC hereafter) is satisfied:*

$$\frac{\partial v((\theta_t, \theta^{t-1}), \theta^{t-1}, y^{t-1})}{\partial \theta_t} = -\frac{z_t(\theta^t)}{\theta_t^2} u_e(c_t(\theta^t), z_t(\theta^t)/\theta_t)$$

$$- \beta \int \frac{\partial v((\theta_{t+1}, \theta^t), \theta^t, y^t(\theta^t))}{\partial \theta_{t+1}} \frac{\partial F_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t)))}{\partial \theta_t} d\theta_{t+1} \quad (9)$$

This envelope condition generalises the corresponding condition in Mirrlees (1971). It describes the information rents, along the equilibrium path, that are associated with skill-history θ^t. Specifically, the right-hand side of the above condition is the additional (net) utility a worker of type $(\theta, \theta^{t-1}), \theta \to \theta_t$, should enjoy relative to type θ^t in order not to mimic type θ^t. This additional utility takes into account of current gains as well as future gains from mimicking type θ^t and reporting truthfully in the future. The current gains arise from supplying labour of $z_t(\theta^t)/\theta_t$ units instead of $z_t(\theta, \theta^{t-1})/\theta$ units, while the future gains arise from influencing the future allocations the worker will enjoy.

4.2 Sufficiency Conditions

The Relaxed Program is valid as a way of determining the second-best allocations if the derived allocations are incentive compatible. We discuss next conditions under which this is the case.

First consider the following assumption, that will be used hereafter.

AF7 $F_t(. \mid \theta_{t-1}, z_{t-1})$ is absolutely continuous with density $f_t(\theta_t \mid \theta_{t-1}, z_{t-1}) > 0$ for almost all $\theta_t \in \Theta$, and all $(\theta_{t-1}, z_{t-1}) \in \Theta \times Z_t$ and any t.

Forward iteration of (9) allows us to rewrite the local IC as

$$\frac{\partial v((\theta_t, \theta^{t-1}), \theta^{t-1}, y^{t-1})}{\partial \theta_t} = -\frac{z_t(\theta^t)}{\theta_t^2} u_e(c_t(\theta^t), z_t(\theta^t)/\theta_t)$$

$$- E[\sum_{\tau=t+1}^{T_L} J_t^\tau(\theta^\tau, z^{\tau-1}(\theta^{\tau-1}))\beta^{\tau-t}\frac{z_\tau(\theta^\tau)}{\theta_\tau^2} u_e(c_\tau(\theta^\tau), z_\tau(\theta^\tau)/\theta_\tau) \mid \theta_t, z_t(\theta^t)] \tag{10}$$

where,

$$E[X \mid \theta_t, z_t(\theta^t)] \equiv$$

$$\int X dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

and, for $\tau > t$,

$$J_t^\tau(\theta^\tau, z^{\tau-1}) \equiv \prod_{k=t+1}^{\tau} I_{k-1}^k(\theta^k, z_{k-1})$$

with

$$I_{k-1}^k(\theta^k, z_{k-1}) \equiv -\frac{\partial F_k(\theta_k \mid \theta_{k-1}, z_{k-1})/\partial \theta_{k-1}}{f_k(\theta_k \mid \theta_{k-1}, z_{k-1})}$$

In the above we have used the fact that $z_\tau = 0$ for any $\tau > T_L$. The term J_t^τ incorporates all the ways (direct and indirect through its effect on interim skills) in which θ_t can affect the distribution of θ_τ. Note that when $t = T_L$, the above becomes Mirrlees's static envelope theorem formula (Mirrlees, 1971).

Pavan et al. (2009) provides us with the following sufficiency result (which is valid in our case as well):[30] A mechanism is IC-h_t for all truthful histories $h_t = \{\theta^t, \theta^{t-1}, y^{t-1}\}$ and all periods t, if (10) holds and

$$D(m_t, h_t) \text{ is non-decreasing in } m_t$$

for every truthful history h_t and all t, where

$$D(m_t, h_t) \equiv$$

$$-E[\sum_{\tau=t}^{T_L} \beta^{\tau-t} J_t^\tau(\theta^\tau, z^{\tau-1}(\theta^{\tau-1}(m_t))\frac{z_\tau(\theta^\tau(m_t))}{\theta_\tau^2} u_e(c_\tau(\theta^\tau(m_t)), z_\tau(\theta^\tau(m_t))/\theta_\tau) \mid \theta_t, z_t(m_t, \theta^{t-1})]$$

with

$$E[X \mid \theta_t, z_t(m_t, \theta^{t-1})] \equiv$$

$$\int X dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1}(m_t)))...dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(m_t, \theta^{t-1}))$$

and, for $T_L \geq \tau \geq t$,

$$\theta^\tau(m_t) \equiv (\theta_{t+1}^\tau, m_t, \theta^{t-1})$$

[30] For completeness, we provide the proof in Appendix B. For a related result see also Theorem 4 in Kapicka (2008) (or Theorem 5 in the 2011 version).

where

$$\theta_{t+1}^{k} \equiv (\theta_{t+1}, ..., \theta_k)$$

with the convention that

$$\theta_{t+1}^{t} \equiv \{\emptyset\}$$

and hence

$$(\theta_{t+1}^{t}, m_t, \theta^{t-1}) \equiv (m_t, \theta^{t-1})$$

In the above we have again taken into account that $z_t = 0$ for all $t > T_L$.

To understand this, note first that in a Markov environment, incentive compatibility for all period-t truthful histories h_t implies incentive compatibility for all (possibly non-truthful) period-t histories because once θ_t is realised (and observed by the agent) the history of the agent's *true* types θ^{t-1} does not matter for her incentives in period t.[31] Therefore, for any given history h^t, truth-telling is optimal from period $t+1$ and onwards even if the agent has misreported her period-t type. In other words, multi-period deviations from truth-telling are not profitable: we can restrict attention to one-stage deviations without loss of generality. Note next that $D(m_t, h_t)$ is the change in period-t expected utility from a marginal increase in period-t skill, given that the consumer misreports the period-t skill and reverts to truth-telling in the future. The above monotonicity of $D(m_t, h_t)$ is thus a single-crossing property on the period-t expected utility. Next, note that $D(\theta_t, h_t)$ is the right-hand side of the local incentive-compatibility constraint (10). The postulated monotonicity of $D(m_t, h_t)$ then implies that truth-telling is also optimal in period t (using familiar arguments from static environments).[32]

To improve further the understanding of the above monotonicity (sufficiency) condition, note that, for $T_L = 1$, $D(m_1, h_1) = -\frac{z_1(m_1)}{\theta_1^2} u_e(c_1(m_1), z_1(m_1)/\theta_1)$. In the case of additively-separable per-period utility,[33] the monotonicity of $D(m_1, h_1)$ reduces to the monotonicity of $z_1(m_1)$. Similarly, if $T_L > 1$, but there is no net persistence of skills, i.e. $\partial F_{t+1}(\theta_{t+1} \mid \theta_t, z_t)/\partial \theta_t \equiv 0$ (which includes as special case the environment where there is no learning-by-doing and skills are serially independent), then $I_t^{t+1}(\theta^{t+1}, z_t) = 0$ and hence the monotonicity of $D(m_t, h_t)$ reduces to the monotonicity of $-\frac{z_1(m_1)}{\theta_1^2} u_e(c_1(m_1), z_1(m_1)/\theta_1)$ in m_t. The monotonicity of $D(m_t, h_t)$ can thus be thought of as the generalisation of the well-known "monotonicity condition" in static environments, by taking into account that in a dynamic environment (unilateral) deviations from truth-telling have multi-period implications for the implemented allocations. [34] We will return to the issue of the validity of the first-order approach at the end of the next Section.

[31] Specifically, past types/reports affect incentives only through the reports-induced allocation, given the allocation rule. Thus, if a mechanism is not incentive compatible at non-truthful history $(\theta^t, (m_{t-1}, m^{t-2}), (y_{t-1}(m_{t-1}, m^{t-2}), y^{t-2}))$ then it will also not be incentive compatible at the truthful history $(\theta^t, (m_{t-1}, m^{t-2}), (y_{t-1}(m_{t-1}, m^{t-2}), y^{t-2}))$, i.e. when the actual period-$t-1$ type is m_{t-1}, given past skills m^{t-2}.

[32] See Appendix B for more details, where, for completeness, we replicate the proof in Pavan et al. (2009) in our Markov set-up.

[33] This is the case most often adopted in the literature, and will also be used in our numerical analysis in Section 7.

[34] Note here that Balle and Spadaro (2006) simply assume that the first-order approach is valid and do not discuss conditions for the latter to be true. For a related assumption see also Krause (2009).

5 Solving the Relaxed Program

The Relaxed Program we have discussed in the previous Section is a discrete-time dynamic optimisation problem where the planner's choice variable (the allocation of that period) *in each and every time period* is *infinite dimensional*. We show next how this problem can be reduced to a sequence of optimal control problems, with each of these problems being conditional on the solution of the previous-period optimal control problem.

To lighten notation, refer hereafter, unless stated otherwise, to a truthful history $h_t = (\theta^t, \theta^{t-1}, y^{t-1})$ as history θ^t. It will prove useful for what follows to rewrite the local IC, (9), by invoking AF3, alongside AF6(iii) and AF6(iv), as:[35]

$$\frac{\partial v(\theta^t, \theta^{t-1})}{\partial \theta_t} = -\frac{z_t(\theta^t)}{\theta_t^2} u_e(c_t(\theta^t), z_t(\theta^t)/\theta_t)$$

$$+ \beta \int \frac{\partial^2 F_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial \theta_t \partial \theta_{t+1}} v(\theta^{t+1}) d\theta_{t+1} \tag{11}$$

We now proceed to the following reparametarisation. Let

$$\hat{U}_{t+1}(\theta_t, \theta^{t-1}) \equiv \int v(\theta_{t+1}, \theta^t) dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t)) \tag{12}$$

$$\hat{V}_{t+1}(\theta_t, \theta^{t-1}) = \int v(\theta_{t+1}, \theta^t) \frac{\partial F^2_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial \theta_t \partial \theta_{t+1}} d\theta_{t+1} \tag{13}$$

with $\hat{U}_{T+1}(\theta^T) = \hat{V}_{T+1}(\theta^T) \equiv 0$. $\hat{U}_{t+1}(\theta^t)$ denotes the promised expected utility from period $t+1$ and onwards in terms of period-$t+1$ utils, referred to as promised utility hereafter. $\hat{V}_{t+1}(\theta^t)$ denotes the change in promised expected utility from period $t+1$ and onwards (in terms of period-$t+1$ utils) from a marginal change in the current skill, referred to as marginal promised utility hereafter. Given this reparametarisation and the definition of $C(.)$ in Section 2.4, (8), (11) and the net-liability law-of-motion (3) can be re-written, respectively, as

$$v(\theta^t) = u(C(v(\theta^t) - \beta \hat{U}_{t+1}(\theta^t), z_t(\theta^t), \theta_t), z_t(\theta^t)/\theta_t) + \beta \hat{U}_{t+1}(\theta^t) \tag{14}$$

$$\frac{\partial v(\theta^t)}{\partial \theta_t} = -\frac{z_t(\theta^t)}{\theta_t^2} u_e(C(v(\theta^t) - \beta \hat{U}_{t+1}(\theta^t), z_t(\theta^t), \theta_t), z_t(\theta^t)/\theta_t) + \beta \hat{V}_{t+1}(\theta^t) \tag{15}$$

and

$$\int [z_t(\theta^t) - C(v(\theta^t) - \beta \hat{U}_{t+1}(\theta^t), z_t(\theta^t), \theta_t)] dF_t(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$= \hat{G}_t(\theta^{t-1}) - \int R^{-1} \hat{G}_{t+1}(\theta^t) dF_t(\theta_{t+1} \mid \theta_t, z_t(\theta^t)) \tag{16}$$

[35] See end of Appendix A for the details.

with $\hat{G}_1(\theta^0) = \hat{U}_{T+1}(\theta^T) = \hat{V}_{T+1}(\theta^T) \equiv 0$ and $\hat{G}_{T+1}(\theta^T) \leq 0$.

The above change of variables allows us to solve the Relaxed Program as a *recursive* sequence of optimal control problems.[36] In more detail, we first express the Relaxed Program, using (14), as $max \int \sum_{t=1}^{T} \beta^{t-1} Q(v(\theta^t) - \beta \hat{U}_{t+1}(\theta^t)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_1(\theta_1)$, subject to (12), (13), (15) and (16). Subsequently, using backward induction, we have that the solution of this problem can be found by considering the recursive problem $\mathcal{R}_t(\theta^{t-1})$, $t = 1, ..., T$: for a fixed period t, $T > t > 1$, fixed period-$t-1$ skill and level of effective labour $(\theta_{t-1}, z_{t-1}(\theta^{t-1}))$, and for given pre-determined variables $\hat{G}_t(\theta^{t-1})$, $\hat{U}_t(\theta^{t-1})$ and $\hat{V}_t(\theta^{t-1})$,

$$\hat{Q}_t(\theta_{t-1}, z_{t-1}(\theta^{t-1}), \hat{U}_t(\theta^{t-1}), \hat{V}_t(\theta^{t-1}), \hat{G}_t(\theta^{t-1})) \equiv$$

$$max \int [Q(v(\theta) - \beta \hat{U}(\theta)) + \beta \hat{Q}_{t+1}(\theta, z(\theta), \hat{U}(\theta), \hat{V}(\theta), \hat{G}(\theta))] dF_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

subject to

$$\int [z(\theta) - C(v(\theta) - \beta \hat{U}(\theta), z(\theta), \theta) + R^{-1} \hat{G}(\theta)] dF_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) \geq \hat{G}_t(\theta^{t-1})$$

$$\dot{v}(\theta) = -\frac{z(\theta)}{\theta^2} u_e(C(v(\theta) - \beta \hat{U}(\theta), z(\theta), \theta), z(\theta)/\theta) + \beta \hat{V}(\theta)$$

$$\hat{U}_t(\theta^{t-1}) = \int v(\theta) dF_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

and

$$\hat{V}_t(\theta^{t-1}) = \int v(\theta) \frac{\partial^2 F_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{\partial \theta_{t-1} \partial \theta} d\theta$$

where to lighten up exposition we have suppressed the dependence of $(v(\theta), z(\theta), \hat{U}(\theta), \hat{V}(\theta), \hat{G}(\theta))$ on θ^{t-1}. For period $t = 1$ the corresponding problem, $\mathcal{R}_1(\theta^0)$, does not have the last two constraints whereas $\hat{G}_1(\theta^0) \equiv 0$, while the problem for period $t = T$, $\mathcal{R}_T(\theta^{T-1})$, is with $\widehat{Q}_{T+1}(.) = 0$, and has the additional restrictions $\hat{U}(\theta) = \hat{V}(\theta) \equiv 0$ and $\hat{G}(\theta) \leq 0$. Note that the pre-determined variables $z_{t-1}(\theta^{t-1})$, $\hat{G}_t(\theta^{t-1})$, $\hat{U}_t(\theta^{t-1})$ and $\hat{V}_t(\theta^{t-1})$ are part of the solution of Program $\mathcal{R}_{t-1}(\theta^{t-2})$. These pre-determined variables will be referred to as the period-t state (variables). The period$-t$ value function $v(\theta)$ (given past skills θ^{t-1}) will be referred to as the period-t phase (variable).

We invoke hereafter AU2 and AF4. Assume for all $t, \theta_{t-1}, z_{t-1}(\theta^{t-1})$, and all θ, that problem

[36] Note here that if there was no net skill persistence then (11) and (15) would reduce to the standard local IC in the static (or dynamic but without skill persistence) models. Recall also from our discussion in Section 4.2 that in this case the first-order approach would be valid if private utility was additively-separable and before-tax labour income was non-decreasing in the current shock. It follows then that in the absence of learning-by-doing and if the principal cannot issue roll-over liabilities (i.e. $\hat{G}_t(.) \equiv 0$), the promised utility is the single state variable that summarises past reporting history as in the well-known models of Green (1987), Thomas and Worrall (1990), Atkenson and Lucas (1992). In fact, the approach in these papers works because the promised utility is independent of the current shock, which is not the case in the presence of (net) shock persistence. This highlights the differences of our set up from the literature on dynamic contracting under hidden information with i.i.d. across time idiosyncratic shocks. For the differences, under shock persistence, of our approach with the recursive formulation in Fernandes and Phelan (2000) see the excellent discussion in Kapicka (2008).

$\mathcal{R}_{t+1}(\theta,\theta^{t-1})$ has a unique solution for any period-$t+1$ state, and that $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ is concave in $(z,\hat{U},\hat{V},\hat{G})$.[37] It follows that Problem $\mathcal{R}_t(\theta^{t-1})$ is a standard optimal control problem (with iso-perimetric constraints). Moreover, after letting $\mu_t(\theta^t)$ be the co-state variable in period$-t$ for skill-type θ^t and p the shadow-price of an extra unit of tax revenues in the first period, we have at an optimum with $z_t(\theta^t)>0$ for any t, that $\hat{G}_{T+1}(\theta^T)=0$ (and hence the intertemporal resource constraint is binding) and, for any t, the following conditions hold:

$$-\frac{\partial\hat{Q}_{t+1}(\theta_t,z_t(\theta^t),.)}{\partial z_t}\beta f_t(\theta_t\mid .)+\mu_t(\theta^t)\frac{1}{\theta_t^2}\{u_e(.)+\frac{z_t(\theta^t)}{\theta_t}u_{ee}(.)\}=$$

$$\mu_t(\theta^t)\frac{z_t(\theta^t)}{\theta_t^2}\frac{u_{ec}(.)u_e(.)}{\theta_t u_c(.)}+pq(t)f_t(\theta_t\mid .)[1+\frac{u_e(.)}{\theta_t u_c(.)}] \tag{17}$$

and

$$\frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}=\frac{1}{pq(t)}Q'(u(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1}))$$

$$+\{\int_{\underline{\theta}}^{\overline{\theta}}\frac{1}{u_c(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x)}\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx \tag{18}$$

$$-\frac{1}{pq(t)}\int_{\underline{\theta}}^{\overline{\theta}}Q'(u(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x))\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx\}$$

$$-\frac{1}{pq(t)}\frac{\mu_{t-1}(\theta^{t-1})}{f_{t-1}(\theta^{t-1}\mid .)}\times\{\int_{\underline{\theta}}^{\overline{\theta}}\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$$

$$+\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}{u_c(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}\}$$

and

$$\mu_t(\theta^t)/pq(t)= \tag{19}$$

$$-\frac{1}{pq(t)}\times\int_{\theta_t}^{\overline{\theta}}\{\int_{\underline{\theta}}^{\overline{\theta}}Q'(u(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x))\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$$

$$-Q'(u(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x))\}\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$$

[37] Unfortunately, we have not been able to push these assumptions further back into the fundamentals. These assumptions ensure the differentiability of $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ and that Kuhn-Tucker first-order conditions are necessary and sufficient for optimality (see Appendices C and D for the details). Similar assumptions, to the same effect, have also been used by Golosov et al. (2011) and (implicitly by) Farhi and Werning (2011) in related problems, but without learning-by-doing. Kapicka (2009) show (in the absence of learning-by-doing) that the concavity of the counterpart of $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ is indeed satisfied if per-period private utility is affine in the allocation of the period in question; this case is equivalent to the case of a per-period utility which is additively-separable with isoelastic cost of labour function.

$$+\int_{\theta_t}^{\overline{\theta}}\{\int_{\underline{\theta}}^{\overline{\theta}}\frac{\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}{u_c(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x)}dx-\frac{1}{u_c(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x)}\}\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},.)dx$$

$$-\frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1}\mid.)}R\beta\int_{\theta_t}^{\overline{\theta}}\{\int_{\underline{\theta}}^{\overline{\theta}}\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\widehat{f}_t(x\mid\theta_{t-1},.)dx$$

$$-\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\}\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$$

where

$$q(t)\equiv(R\beta)^{1-t}$$

$$\Gamma(x,\theta^t)\equiv exp[-\int_{\theta_t}^{x}\frac{z_t(\vartheta,\theta^{t-1})}{\vartheta^2}\frac{u_{ec}(c_t(\vartheta,\theta^{t-1}),z_t(\vartheta,\theta^{t-1})/\vartheta)}{u_c(c_t(\vartheta,\theta^{t-1}),z_t(\vartheta,\theta^{t-1})/\vartheta)}d\vartheta] \tag{20}$$

$$\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))\equiv\frac{\Gamma(x,\underline{\theta},\theta^{t-1})f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}{\int_{\underline{\theta}}^{\overline{\theta}}\Gamma(x,\underline{\theta},\theta^{t-1})f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx} \tag{21}$$

$$p=\frac{\int_{\underline{\theta}}^{\overline{\theta}}Q'(u(c_1(x),z_1(x)/x))\widehat{f}_1(x)dx}{\int_{\underline{\theta}}^{\overline{\theta}}[u_c(c_1(x),z_1(x)/x)]^{-1}\widehat{f}_1(x)dx}>0 \tag{22}$$

and $\frac{\partial\hat{Q}_{T+1}(\theta_T,z_T(\theta^T),.)}{\partial z_T}\equiv$ and $\mu_0(\theta^0)\equiv 0$.

Leaving this section, we note that the above necessary conditions reduce to those found in the static models of Mirrlees (1971), Diamond (1998) and Saez (2001), to mention few, when $t=T=1$.[38]

5.1 Validity of First-Order Approach

Leaving this Section, we note that an alternative approach to checking/imposing the sufficiency condition of Section 4.2 is to verify directly the incentive compatibility of any proposed allocation (rule). This is a common approach in static screening models, and could in principle be used in our dynamic context as well.[39]

Specifically, let $c_t^*(r,\theta^{t-1})\equiv C(v(r)-\beta\widehat{U}(r),z(r),r)$, $z_t^*(r,\theta^{t-1})\equiv z(r)$ and $\widehat{U}_{t+1}^*(r,\theta^{t-1})\equiv\widehat{U}(r)$, where $v(r),z(r)$ and $\widehat{U}(r)$ are part of the solution of the period-t Program $\mathcal{R}_t(\theta^{t-1})$. Therefore, $\{c_t^*(r,\theta^{t-1}),z_t^*(r,\theta^{t-1})\}_{r\in\Theta}$ is the period-t allocation rule given history θ^{t-1}. Let also $v^*(\theta_{t+1},r,\theta^{t-1})\equiv u(c_{t+1}^*(\theta_{t+1},r,\theta^{t-1}),z_{t+1}^*(\theta_{t+1},r,\theta^{t-1})/\theta_{t+1})+\beta\widehat{U}_{t+2}^*(\theta_{t+1},r,\theta^{t-1})$ be the solution for the period-$t+1$ value function (associated with period-$t+1$ skill history $(\theta_{t+1},r,\theta^{t-1})$) that arises from the period-$t+1$ Program $\mathcal{R}_{t+1}(\theta^t)$. We then have that the solution of the Re-

[38] Or, when $t=T>1$ and there is zero net skill persistence, i.e. $\partial F_t(x\mid\theta_{t-1},z_{t-1})/\partial\theta_{t-1}=0$ for all (x,θ_{t-1},z_{t-1}).

[39] See also Farhi and Werning (2011) and Kapicka (2011).

laxed problem satisfies incentive compatibility (and thus the first-order approach is valid and the conditions above characterise second-best allocations) if, for given past skill-history θ^{t-1}:

$$\theta = argmax_{r\in\Theta}\{u(c_t^*(r,\theta^{t-1}), z_t^*(r,\theta^{t-1})/\theta) + \beta \int v^*(\theta_{t+1}, r, \theta^{t-1}) dF_{t+1}(\theta_{t+1} \mid \theta, z_t^*(r,\theta^{t-1}))$$

We assume hereafter that either the sufficiency condition of Section 4.2 or the above condition for any t and θ^{t-1} and any $\theta \in \Theta$ are satisfied (we will also verify this in our numerical analysis in Section 7). Therefore, we use the necessary conditions of the Relaxed Program to discuss in the next section second-best allocations and the implies tax wedges.

6 Optimal Tax Wedges

Understanding optimal dynamic redistribution and implied tax wedges amounts to studying (17)-(19).[40] To discuss these, we will restrict attention hereafter on the case of a *regular* economy:

Definition An economy is *regular* if, at optimum, (a) marginal utility of consumption is non-increasing in current skill, and (b) the social marginal utility is non-increasing in current skill.

Restricting attention to regular economies facilitates the interpretation of the necessary conditions of the Relaxed Program. Interpretation of policies in economies that are not regular is very cumbersome. Note also that property (a) can be proved in the static model under the "Mirrlees condition", normality of consumption and substitutability between consumption and leisure. However, this is no longer guaranteed in a dynamic framework. Finally, we note that property (b) is also postulated in static models during the interpretation of results.

It will help our understanding to start our discussion with simpler case of additively-separable private per-period utility.

6.1 The case of additively-separable utility

Here $u_{ec} = 0$ (and thus $\Gamma = 1$ and $\widehat{f} = f$).

6.1.1 The savings tax wedge

In this case, condition (18) becomes

$$\frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} =$$

[40] Note that after using (17) for period $t-1$ to eliminate $\mu_{t-1}(\theta^{t-1})$ from (19), and the resulting equation to eliminate $\mu_t(\theta^t)$ from (17) for period t, and combining the result with the definition of the labour wedge, (2), one can derive a law-of-motion for the labour tax wedge along the lines of Farhi and Werning (2011). The interpretation would, however, be very cumbersome here due to the presence of learning-by-doing. Therefore, to interpret our results, we follow steps that are closer to the traditional labour income taxation literature.

$$\Delta_t(\theta^{t-1}) + \int_{\underline{\theta}}^{\overline{\theta}} \frac{1}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))dx \tag{23}$$

where

$$\Delta_t(\theta^{t-1}) \equiv \frac{1}{pq(t)} \times$$

$$\{Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})) - \int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)) f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))dx\}$$

Under Utilitarian preferences, we have $\Delta_t(\theta^{t-1}) = 0$ and the above reduces to the standard inverse Euler condition,[41] which, in turn, implies a positive intertemporal tax wedge (given Jensen's inequality).

The inverse Euler condition simply ensures that reducing the utility of a consumer with skill history θ^{t-1} by $\beta\varepsilon$ in exchange of increasing by ε, with $\varepsilon \to 0$, her future utility *regardless* of the realised future skill, does not free up resources and hence is not welfare improving. Note here that, in contrast to the full-information case, the increase in future utility must be uniform across period-t skills in order to be incentive-compatible. The intuition behind the positive intertemporal wedge is that it reduces to incentive to save more in order to shirk next period.

The standard inverse Euler condition may no longer hold when we move to a non-Utilitarian planner. The intuition for this is analogous to the one behind the right-hand side of the generalised Euler condition (4) in Section 3.1. Simply note that when decreasing period-$t-1$ utility by β units in exchange of higher period-t utility by one unit, the latter has to take place regardless of the realised period-t skill, due to incentive compatibility, and that $\Delta_t(\theta^{t-1})$ is the negative of the conditional expectation of the right-hand side of (4).

Consequently, the second-best intertemporal wedge will be higher (resp. lower), all other things equal, than that set by a Utilitarian planner if $\Delta_t(\theta^t) > 0$ (resp. $\Delta_t(\theta^t) < 0$); that is, if at optimum the (conditional) expected marginal social welfare is decreasing (resp. increasing) over time, for given current skill. The reason is that, in this case, the social marginal welfare gain from higher current (consumption and, thereby, for given earnings) private utility is higher (resp. lower) than the social marginal expected welfare gain from higher future private utility (regardless of the realised future skill). Thus, the planner has an incentive to reduce (resp. increase) savings and thereby set a higher (resp. lower) intertemporal wedge. This finding extends the result in Golosov et al. (2003) for the case of a non-Utilitarian planner (and Markov stochastic processes for skills).[42]

[41] By adopting techniques in Rogerson (1985) one can derive the above condition as a necessary condition for optimality in the Full-Program of the planner, as it has also been shown in the New Dynamic Public Finance literature.

[42] Note that $\Delta_t(\theta^t) < 0$ could be the case even if (5) holds, as the measure used to take conditional expectations is different in the two conditions.

We turn to the labour tax wedge.

6.1.2 The labour tax wedge

We have from (22) that

$$p = \frac{\int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_1(x), z_1(x)/x))f_1(x)dx}{\int_{\underline{\theta}}^{\overline{\theta}} [u_c(c_1(x), z_1(x)/x)]^{-1} f_1(x)dx} \tag{24}$$

To increase the utility of all consumers by one unit the planner needs to raise revenues equal to the denominator in (24). The above condition then says that the social cost of such a raise must equal the impact on social welfare (the numerator in (24)).

Furthermore, condition (17) becomes (after denoting $\eta(e) \equiv 1 + e\frac{u_{ee}(e)}{u_e(e)}$)

$$\frac{\mu_t(\theta^t)}{pq(t)} \frac{u_e(.)\eta(e_t(\theta^t))}{\theta_t^2} = [1 + \frac{u_e(.)}{\theta_t u_c(.)} + \frac{\partial \hat{Q}_{t+1}(\theta_t, z_t(\theta^t), .)}{\partial z} \frac{\beta}{pq(t)}] f_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) \tag{25}$$

where $e_t(\theta^t) = z_t(\theta^t)/\theta_t$. This simply says that at optimum a reduction in the period-t effective labour of workers with skill history θ^t must balance at the margin (with units expressed in terms of period-t public funds) the social gain from the resulting reduction in the information rents associated with period-t type θ_t given history θ^{t-1} (the term in the left-hand side) with the social cost from the resulting net reduction in planner's revenues from all taxpayers with skill-history θ^t (the first term in the right-hand side) and from making low future skills more likely for them (the second term in the right-hand side). Recalling the definition of the labour tax wedge, (2), we have that the latter can be re-written as:

$$\zeta_t^z(\theta^t) =$$

$$\frac{\hat{\delta}(\theta^t, z_t(\theta^t))}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)} - \frac{\partial \hat{Q}_{t+1}(\theta_t, z_t(\theta^t), .)}{\partial z} \frac{\beta}{pq(t)} + \frac{[\mu_t(\theta^t)/pq(t)]}{\theta_t f_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} \frac{u_e(.)\eta(e_t(\theta^t))}{\theta_t} \tag{26}$$

with $\mu_t(\theta^t)/pq(t)$ given implicitly by condition (19).

The third term of (26), which arises because of asymmetric information, has been explored extensively in the received literature for the case of no learning-by-doing. When there is no learning-by-doing (as, for example, at $t = T$), we have $\hat{\delta}(\theta^t, z_t(\theta^t)) = \frac{\partial \hat{Q}_{t+1}(\theta, ., z(\theta), .)}{\partial z} \equiv 0$ and hence the formula for the labour tax wedge coincides with the third term.

The first two terms, for $0 < t < T$, of (26) are novel and capture the direct impact of learning-by-doing on the labour wedge. As we will see, the second term depends also on $\mu_t(\theta^t)/pq(t)$. To facilitate understanding of the labour tax wedge, we discuss first $\mu_t(\theta^t)/pq(t)$ and then the first two terms in (26).

We will discuss first the case of Utilitarian planner and then the case of a non-Utilitarian planner.

Utilitarian Planner Note that when the planner has Utilitarian preferences, (19) becomes (for any $T \geq t > 0$)

$$\mu_t(\theta^t)/pq(t) =$$

$$\int_{\theta_t}^{\overline{\theta}} \{ \int_{\underline{\theta}}^{\overline{\theta}} \frac{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} dx - \frac{1}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} \} f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) dx$$

$$- \frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1} \mid .)} R\beta \frac{\partial F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{\partial \theta_{t-1}} \tag{27}$$

where $\mu_0(\theta^0) = 0$. $\mu_t(\theta^t)/pq(t)$ is the social marginal valuation of period-t value function associated with skill-history θ^t in terms of period-t public funds (at optimum). That is, it is the shadow price of period$-t$ value function of type θ^t, $v(\theta^t)$.

Lemma 1: *Suppose additively-separable private utility and a Utilitarian planner. Then, in a regular economy, $\mu_t(\theta^t)/pq(t) < 0$ for any interior skill θ_t, for all $t > 0$.*

Proof: Let

$$B_t(\theta_t) \equiv \frac{1}{1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} \times$$

$$\int_{\theta_t}^{\overline{\theta}} \{ \int_{\underline{\theta}}^{\overline{\theta}} \frac{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} dx - \frac{1}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} \} dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

and note that $B_t(\underline{\theta}) = 0$, and $B_t'(\theta_t) < 0$ when the equilibrium marginal utility is decreasing in current skill. Note now that $B_t(\theta_t)[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))]$ is the first term in (27), which implies, given net skill persistence, by backward induction on (27) and $\mu_0(\theta^0) = 0$, that $\mu(\theta^t) < 0$ for any $t > 0$ and any interior θ_t.■

Note that $B_t(\theta_t)$ describes the difference between the average revenues needed to marginally increase everyone's utility and the conditional average revenues needed to marginally increase utility of every individual whose skill is over θ_t. Therefore, the proof above implies also that in a regular economy, the lower the skill θ_t the higher $B_t(\theta_t)$ is. Thus, $B_t(\theta_t)$ can be thought of as a measure of the redistributional motive of a Utilitarian planner under additively-separable private per-period utility: the lower $B_t(\theta^t)$, the lower the incentives of the planner to redistribute towards individuals with-skill history θ^t; that is, the lower $\mu_t(\theta^t)/pq(t)$ and the higher the marginal tax imposed on these workers.

To understand why the social marginal valuation of period-t value function is equal, at optimum, to the right-hand side of (27), we evaluate the social value of increasing marginally period-t utility (and thereby value function) of all consumers with skill history θ^t (by changing appropriately consumption for given earnings). We start by noting that the local IC, (11), requires that the planner must do so for all consumers with skill-history (x, θ^{t-1}), $x > \theta_t$, too. Furthermore, these changes in utility do not affect current information rents (given $u_{ec} = 0$).

In a static environment, or - to the same effect - in the first period here, this would lead to a decrease in revenues from each and every such taxpayer and an increase in social welfare associated with every such taxpayer. The decrease in revenues equals the inverse of the taxpayer's marginal utility of period-t consumption. The increase in social welfare is equal to $1/p$ (in terms of first-period public funds). Recalling (24) with $Q' \equiv 1$, this is equal to the average of the inverse marginal utilities of first-period consumption. Putting these two effects together gives the right-hand side of (27) for $t = 1$ (recall also that $\mu_0(\theta_0) = 0$).

However, in a dynamic environment, the planner can also raise revenues in the previous period, and transfer them in period t (plus interest) to finance the increase in period-t utility, if $t > 1$. Note, however, that in period $t-1$ the planner cannot condition the period-$t-1$ allocation on the period-t skill. Suppose thus that, given skill history θ^{t-1}, the planner increases the period-t utility by one unit of all taxpayers with period-t skill (weakly) above θ_t, while decreasing by $[1-F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))]\beta$ units period-$t-1$ utility of all workers (by means of an appropriate change in consumptions, for given earnings). Note that this reallocation of utility over time leaves total discounted private utility unaffected, and, as in the $t = 1$ case, does not affect current information rents (given $u_{ec} = 0$).

The effects of this perturbation are as follows: First, as in the $t = 1$ case, it reduces period-t revenues from each and every such taxpayer by the inverse of the taxpayer's marginal utility of period-t consumption. This is captured by the second term in the curly brackets in the integrand in (27). Second, it increases period-$t-1$ revenues from each and every such taxpayer by $R\beta[1-F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))]$ times the inverse of the taxpayer's marginal utility of period-$t-1$ consumption (in units of period$-t$ public funds,). Recalling (23) with $Q' \equiv 1$, the latter equals the first term in the curly brackets in the integrand in (27). Finally, in a dynamic environment, the marginal increase in period-t utility of all consumers with period-t skill history $(x, \theta^{t-1}), x \geq \theta_t$, increases also the *future* information rents of all workers of period-$t-1$ skill-history θ^{t-1} *conditional* on their period-t skill turning out to be $x \geq \theta_t$. This increase, in terms of period-$t-1$ utils per worker with skill history (x, θ^{t-1}), equals $\beta \frac{\partial f_t(x|\theta_{t-1}, z_{t-1}(\theta^{t-1}))/\partial \theta_{t-1}}{f_t(x|\theta_{t-1}, z_{t-1}(\theta^{t-1}))}$ (recall the local IC, expressed as in (15), and the definition of marginal promised utility (13) for period $t-1$). The per-unit social value of this (in terms of period-$t-1$ public funds per worker of skill history θ^{t-1}) is equal to $\frac{\mu(\theta^{t-1})/pq(t-1)}{f_t(\theta_{t-1}|.)}$, and expressing the total effect over all workers with period-t skill-history $(x, \theta^{t-1}), x \geq \theta_t$, in terms of period-$t$ public funds gives rise to the second term in (27).

The above discussion highlights that moving to a dynamic environment affects the period-t, $t > 1$, social marginal valuation of value function in two ways. First, ignoring net persistence for a while, the net effect (of the aforementioned perturbations) associated with a taxpayer of skill-history (x, θ^{t-1}) is

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{dF_1(\theta)}{u_c(c_1(\theta), z_1(\theta)/\theta)} - [u_c(c_1(x), z_1(x)/x)]^{-1}$$

in the first-period, instead of

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{dF_t(\theta_t|\theta_{t-1}, z_{t-1}(\theta^{t-1}))}{u_c(c_t(\theta^t), z_t(\theta^t)/\theta_t)} - [u_c(c_t(x, \theta^{t-1}), z_t(x, \theta^{t-1})/x)]^{-1}$$

in some period $t > 1$. For the sake of argument, restrict attention to the case when marginal utilities in the first period and period $t > 1$ are the same for type (x, θ^{t-1}). It follows, then, that the fact that the planner can finance an increase in period-t consumption with a decrease in period-$t-1$ consumption raises (relative to the first period) the redistribution incentives vis-a-vis skill x given history θ^{t-1}, if $\int_{\underline{\theta}}^{\overline{\theta}} \frac{dF_1(\theta)}{u_c(c_1(\theta), z_1(\theta)/\theta)} < \int_{\underline{\theta}}^{\overline{\theta}} \frac{dF_t(\theta|\theta_{t-1}, z_{t-1}(\theta^{t-1}))}{u_c(c_t(\theta, \theta^{t-1}), z_t(\theta, \theta^{t-1})/\theta)}$ and vice versa. The latter condition simply states that the average revenues raised by uniformly reducing consumption by one unit are higher in period t than in the first period. Second, if there is net skill persistence and the marginal utility of consumption is decreasing in current skill (which would be the case in a regular economy), the second term in (27) is negative for any interior θ_{t-1} (recall Lemma 1), which highlights that skill persistence reduces the redistribution motive of the planner. As also discussed in Golosov et al. (2011), the intuition here is that in the presence of skill persistence, a worker who has a low income early in life and high income later in life is more likely to be one who has mimicked to be of low skill early in life. Therefore, increasing the labour distortions worsens the benefits from masquerading to be of low skill early in life.

We now move to the discussion of the first two terms in the right-hand side of (26), which represent the effect of learning-by-doing on the labour wedge when $t < T$. To facilitate understanding of this effect, we focus on the two-period version of the model.[43]

In Appendix E, we show that the learning-by-doing component of the first-period labour wedge (i.e. the first two terms in (26)) is

$$\widehat{\zeta}_1^z(\theta_1) \equiv \frac{\hat{\delta}(\theta_1, z_1(\theta_1))}{u_c(c_1(\theta_1), z_1(\theta_1)/\theta_1)} - \frac{\partial \hat{Q}_2(\theta_1, z_1(\theta_1), .)}{\partial z} \frac{\beta}{p} =$$

$$-\beta \times \{\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)} \int u(c_2(\theta_2, \theta_1), z_2(\theta_2, \theta_1)/\theta_2)) \frac{\partial^2}{\partial z_1(\theta_1) \partial \theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))$$

$$+[R\beta]^{-1} \int [z_2(\theta_2, \theta_1) - c_2(\theta_2, \theta_1)] \frac{\partial}{\partial z_1(\theta_1)} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))\}$$

This has two parts. The second part is the full-information wedge, which we have already discussed. The first part emerges due to skills being consumers' private information and leads,

[43] Insights extend in a straightforward manner to the case of more than two periods. Bearing in mind our earlier discussion of the social marginal valuation of period-t value function, the optimal labour distortion in an *interim* period $1 < t < T$ will still be driven, as we will shortly show in the main text, by the reduction of current and future information rents (and the possible tensions thereof) while taking into account the learning-by-doing effect. The multi-period model, however, can provide us with some insights on whether the average labour tax wedge decreases with age at optimum (see Kapicka (2008), Golosov et al. (2011) and Farhi and Werning (2011) for related discussions). We will return to this at our numerical analysis in Section 7.

ceteris paribus, to a lower wedge (relative to the no learning-by-doing case) if

$$0 < \mathcal{T}_1 \equiv$$

$$\frac{1-F(\theta_1)}{f_1(\theta_1)} B_1(\theta_1) \int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)) \frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))$$

and vice versa. A sufficient condition for $\mathcal{T}_1 < 0$ (and hence a higher wedge) is that the economy is regular (and hence $B_1(\theta_1) < 0$ from Lemma 1) *and* $\int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)) \frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1)) dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))) > 0$.

To understand the intuition behind $\mathcal{T}_1$ - and the first term in the learning-by-doing component of the first-period labour wedge, $-\beta\mathcal{T}_1$ - note from (6) with $Q' \equiv 1$ and (26) that for an interior skill, the difference between the second-best and full-information labour wedges equals

$$\widetilde{\zeta}_1^z(\theta_1) = -\frac{1-F(\theta_1)}{f_1(\theta_1)} B(\theta_1) \times \tag{28}$$

$$\{-\frac{u_e(c_1(\theta_1),.)\eta(e_1(\theta_1))}{\theta_1^2} + \beta \int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)) \frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))\}$$

Note that the second term is simply $-\beta\mathcal{T}_1$. The term in the curly brackets is the change in the information rents (i.e. in the right-hand side of the first-period local IC as expressed in (15) for $t = 1$) from a marginal increase in effective labour. This term drives whether the distortion in labour supply (to reduce information rents) is upward or downward, while the extend of the distortion depends on the "weight" $-\frac{1-F(\theta_1)}{f_1(\theta_1)} B(\theta_1)$. This weight represents two considerations. The "neediness" of skill θ_1, as captured (inversely) by $-B(\theta_1)$, and the planner's information problem associated with skill θ_1, as captured by the relative likelihood of higher skills $\frac{1-F(\theta_1)}{f_1(\theta_1)}$ vis-a-vis skill θ_1. Thus, if the term inside the curly brackets is positive, then the planner wants to distort downwards the labour supply of skill θ_1 to reduce information rents associated with this skill: this calls for a positive labour wedge, and vice versa.

The first part in (28) is familiar from the static model: it captures the change in current information rents from mimicking (at the margin) to be of skill θ_1, and calls indeed for a positive wedge. The reason is that if a consumer with higher skill mimicks to be of skill θ_1, then she will work less than individuals of skill θ_1 (in order to achieve effective labour $z_1(\theta_1)$) leading to a welfare gain, which (equals the first term in (15) for $t = 1$ and) is increasing in $z_1(\theta_1)$.

The second term in (28) is novel. It captures the *labour-induced* change in *future* rents, and may call, depending on its sign, for a positive, zero or negative wedge. In more detail, by mimicking to be of skill θ_1 (at the margin), one also faces future information rents equal to $\beta\widehat{V}_2(\theta_1) = \beta \int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)) \frac{\partial}{\partial\theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1)))$. In the presence of learning-by-doing, these rents depend on (effective) labour. The type of this dependence influences the planner's incentives to distort labour of skill θ_1. Thus, if future information rents are non-decreasing

in effective labour, the incentives to reduce current and future information rents associated with skill θ_1 are aligned, and the planner distorts effective labour of skill θ_1 further downwards. When future information rents are decreasing in effective labour, there is, however, a trade-off between current and future rents/inequality. As in the static model, a positive wedge on (and hence lower labour from) skill θ_1 reduces current information rents associated with skill θ_1. Nevertheless, under learning-by-doing, a positive wedge leads also to an increase in future information rents associated with skill θ_1. Therefore, the optimal wedge in this case can be thought of balancing the trade-off between current and future inequality.

To discuss further this potential tension, we would like to pin the dependence of future information rents on effective labour down on properties of the skills distribution. We have:

Proposition 2: *In addition to AF3, AF6(iii) and AF6(iv), assume AF4, and that* $\frac{\partial F_2(.|\theta_1,.)}{\partial \theta_1}$ *and* $\frac{\partial^2 F_2(.|\theta_1,.)}{\partial \theta_1 \partial z_2}$ *are continuous on* $\Theta \times Z_2$ *for any* $\theta_1 \in \Theta$. *If skill and effective labour are complements in making high future skills more likely, i.e.* $\frac{\partial^2}{\partial z_1(\theta_1)\partial \theta_1} F_2(\theta_2 \mid \theta_1, z_1(\theta_1)) < 0$, *then* $\int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)) \frac{\partial^2}{\partial z_1(\theta_1)\partial \theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))) > 0$, *and hence* $\mathcal{T}_1 < 0$, *and vice versa.*

Proof: By using first integration by parts - as in the proof of Lemma A3 in Appendix A - and then that $v(\theta^2) = u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2)$ alongside the last-period local IC (expressed as in (11)), we have

$$\int u(c_2(\theta_2,\theta_1), z_2(\theta_2,\theta_1)/\theta_2) \frac{\partial}{\partial \theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1)) =$$

$$\int \frac{z_2(\theta^2)}{\theta_2^2} u_e(c_2(\theta^2), z_2(\theta^2)/\theta_2) \frac{\partial F_2(\theta_2 \mid \theta_1, z_1(\theta_1))}{\partial \theta_1} d\theta_2$$

The result follows directly from differentiating both sides with respect to $z_1(\theta_1)$ (where AF4 and continuity of $\frac{\partial F_2(.|\theta_1,.)}{\partial \theta_1}$ and $\frac{\partial^2 F_2(.|\theta_1,.)}{\partial \theta_1 \partial z_2}$ allows us to use Leibniz integral rule to take the derivatives into the integrals).■

Under the assumptions of Proposition 2 we therefore have that when skill and effective labour are complements in making high future skills more likely, the planner has an incentive to further distort labour of skill θ_1 downwards.[44] When skill and effective labour are substitutes, there is,

[44] Recalling also footnote 23, the assumptions in Proposition 2 are satisfied if the skill distribution is the Pareto with mean $\Xi(\theta_1, \frac{z_1}{\theta_1}) = \xi + \rho^{-1}\theta_1^{\rho} + log(\theta_1 + \frac{1}{1+e^{-(z_1/\theta_1)}}), \xi \geq 0, \rho \to 0^+$ and Pareto index $\alpha > 1$ sufficiently high. One can then show that $\frac{\partial^2}{\partial z_1 \partial \theta_1} F_2(\theta_2 \mid \theta_1, z_1) \leq 0$ for any $z_1 \in \mathbb{R}_+$ and any θ_1 not very low (in fact the relevant threshold skill level is smaller than 1/2), if α is sufficiently high. To see this note that (a) the Pareto cdf is decreasing in its mean, (b) the ratio of the second and first derivatives of the Pareto cdf with respect to the mean is equal to $\frac{(\alpha-1)}{\Xi(\theta_1, \frac{z_1}{\theta_1})}$ and (c) there is a sufficiently high Pareto index $\alpha > 1$ to ensure that $\frac{\partial^2}{\partial \theta_1 \partial z_1}\Xi(\theta_1, \frac{z_1}{\theta_1}) + \frac{(\alpha-1)}{\Xi(\theta_1, \frac{z_1}{\theta_1})} \frac{\partial}{\partial \theta_1}\Xi(\theta_1, \frac{z_1}{\theta_1}) \frac{\partial}{\partial z_1}\Xi(\theta_1, \frac{z_1}{\theta_1}) > 0$ because this term is increasing and linear in α and $\frac{\partial^2}{\partial \theta_1 \partial z_1}\Xi(\theta_1, \frac{z_1}{\theta_1})$ is bounded for not very low skills (the relevant threshold is also below 1/2). The case of Pareto distributed skill is discussed in Kapicka (2008), and so, for a sufficiently high Pareto index, our model predicts higher labour wedges, compared to that work, for all period-1 skills that are not very low (not below 1/2 would suffice).

however, is a tension between current and future information rents, and the optimal wedge is lower, all other things equal, than under no learning-by-doing.[45]

Interestingly, we could also have a situation where $\frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1}F_2(\theta_2 \mid \theta_1, z_1(\theta_1)) < 0$ for large θ_1 and $\frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1}F_2(\theta_2 \mid \theta_1, z_1(\theta_1)) > 0$ for lower θ_1. For instance, this would be the case if (a) the (second-period) skill distribution was the Lognormal with a Pareto tail fitted at the modal skill level $\theta^m \equiv \Xi e^{-\frac{3}{2}\sigma^2}$, where Ξ denotes the mean and σ^2 the shape parameter of the Lognormal distribution,[46]with the Pareto index sufficiently high and the mean Ξ given by the function $\Xi(\theta_1, \frac{z_1}{\theta_1}) = \xi + \rho^{-1}\theta_1^\rho + log(\theta_1 + \frac{1}{1+e^{-(z_1/\theta_1)}}), \xi \geq 0, \rho \to 0^+$.[47]The above proposition would then imply (for given history $(\theta_1, z_1(\theta_1))$) a more progressive tax, compared to the one without learning-by-doing, in the sense that here taxes are lower below the modal skill level and higher above it. This result emphasises the importance of learning-by-doing for the design of redistribution.

The Non-Utilitarian Planner Here, (19) implies that (for any $T \geq t > 0$)

$$\mu_t(\theta^t)/pq(t) =$$

$$+\int_{\theta_t}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} \frac{dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

$$-\frac{1}{pq(t)}\int_{\theta_t}^{\overline{\theta}}\int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/\theta_{t-1}))dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

$$+\int_{\theta_t}^{\overline{\theta}} \frac{Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/\theta_{t-1}))}{pq(t)} dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

[45]Recalling also footnote 23, the assumptions in Proposition 2 are satisfied if the skill distribution is the Lognormal with mean $\Xi(\theta_1, \frac{z_1}{\theta_1}) = \xi + \rho^{-1}\theta_1^\rho + log(\theta_1 + \frac{1}{1+e^{-(z_1/\theta_1)}}), \xi \geq 0, \rho \to 0^+$. In this case, one can easily see after some tedious algebra that $\frac{\partial^2}{\partial z_1\partial\theta_1}F_2(\theta_2 \mid \theta_1, z_1) \geq 0$ for any $z_1 \in \mathbb{R}_+$ and any θ_1 not very small (in fact the relevant threshold skill level is smaller than 1/2). The Lognormal case is discussed in Mirrlees (1971) and Farhi and Werning (2011), and so, compared to these papers, our model predicts lower labour wedges for all period-1 skills that are not very low (not below 1/2 would suffice).

[46]See Diamond (1998) for a related construction. The skills distribution we consider here also fits the one used in Saez (2001) and Golosov et al. (2011), where the skills distribution is calibrated from the data.

[47]Recalling also the above two footnotes, one can easily see after some tedious algebra that in this case $\frac{\partial^2}{\partial z_1\partial\theta_1}F_2(\theta_2 \mid \theta_1, z_1) > 0$ for θ_2 below the mode, for any $z_1 \in \mathbb{R}_+$ and any $\theta_L \leq \theta_1$, where $\theta_L < 1/2$ is a threshold skill above which $\frac{\partial^2}{\partial\theta_1\partial z_1}\Xi(\theta_1, e_1) \leq 0$ and bounded for all $e_1 \in \mathbb{R}_+$. Moreover, there is a sufficiently high Pareto index $\alpha > 1$ to ensure that $\frac{\partial^2}{\partial\theta_1\partial z_1}\Xi(\theta_1, \frac{z_1}{\theta_1}) + \frac{(\alpha-1)}{\Xi(\theta_1, \frac{z_1}{\theta_1})}\frac{\partial}{\partial\theta_1}\Xi(\theta_1, \frac{z_1}{\theta_1})\frac{\partial}{\partial z_1}\Xi(\theta_1, \frac{z_1}{\theta_1}) > 0$ and very large, because the whole expression is increasing and linear in α. For such Pareto index we then have that $\frac{\partial^2}{\partial z_1\partial\theta_1}F_2(\theta_2 \mid \theta_1, z_1) \leq 0$ for θ_2 above the mode, for any $z_1 \in \mathbb{R}_+$ and any $\theta_1 \geq \theta_L$. Details are in Appendix G.

$$-\int_{\theta_t}^{\overline{\theta}} \frac{1}{u_c(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x)} dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

$$-\frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1} \mid .)} R\beta \frac{\partial F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{\partial \theta_{t-1}} \tag{29}$$

where, recall, $\mu_0(\theta^0) = 0$. We can obtain the intuition of the above equation in a similar manner to that when we obtained the intuition of (27) in the Utilitarian case. Suppose thus, given skill history θ^{t-1}, that the planner increases marginally period-t utility of all taxpayers with period-t skill (weakly) above θ_t, while, if $t > 1$, decreasing by $[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})]\beta$ units period-$t-1$ utility of all workers - by means of appropriate changes in consumption. The fourth and fifth terms above have been discussed at (27).

The first difference with the Utilitarian case is that now the increase in period-t utility of type $(x, \theta^{t-1}), x \geq \theta_t$, leads to an increase in period-t social welfare by $Q'(u(c_1(x), z_1(x)/\theta_1))/p$ (in units of period-t public funds), with the total effect being equal to the third term in (29). [Note also, recalling (24), that the first two terms in (29) cancel out when $t = 1$.]

The second difference, which is relevant only when $t > 1$, is that the private utility rearrangement in question leads here to a reduction in period-$t-1$ social welfare that equals, in terms of period-t public funds, $[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})] \frac{Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1}))}{pq(t)}$ - in addition to the increase in period-$t-1$ revenues of $[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})] \frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}$ units in terms of period-t public funds that we have already discussed at (27). Clearly, then, after recalling (23), we have that these two effects together, i.e.

$$\int_{\theta_t}^{\overline{\theta}} \{\frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} - \frac{Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1}))}{pq(t)}\} dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

reduce to the first two terms in (29).

Before we proceed, we state the following result:

Lemma 2: *Suppose additively-separable private utility and a non-Utilitarian planner. Then, in a regular economy, $\mu_t(\theta^t)/pq(t) < 0$ for any interior skill θ_t, and any t.*

Proof: Let

$$B_t^Q(\theta_t) \equiv$$

$$\int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/\theta_t) f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) dx$$

$$-\int_{\theta_t}^{\overline{\theta}} Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/\theta_{t-1}) \frac{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} dx$$

and note that the second and third terms in (29) together reduce to $-\frac{1}{pq(t)}[1 - F_t(\theta_t \mid$

$\theta_{t-1}, z_{t-1}(\theta^{t-1}))]B_t^Q(\theta_t)$. Note now, after differentiating $B_t^Q(\theta_t)$ with respect to θ_t, that $B_t^Q(\theta_t)$ is increasing in skill if the marginal social welfare is non-increasing with skill (at optimum). Note also that $B_t^Q(\underline{\theta}) = 0$ and hence $B_t^Q(\theta_t) > 0$ for any $\theta_t > \underline{\theta}$. This implies, given net skill persistence and Lemma 1, by backward induction on (29) and $\mu_0(\theta^0) = 0$, that $\mu_t(\theta^t) < 0$ for any $t > 0$ and any interior θ_t.■

Therefore, in a regular economy, we have that the lower the skill θ_t the lower $B_t^Q(\theta_t)$ is, and the second and third terms in (29) together reduce to a negative number. $B_t^Q(\theta_t) > 0$ can hence be thought of as a (inverse) measure of the "additional" redistributional considerations of a non-Utilitarian planner: the lower $B_t^Q(\theta_t)$ the higher the incentive to redistribute towards skill θ_t.

Thus, moving to a non-Utilitarian planner has two direct effects. First, in a regular economy, it increases the shadow price (in terms of public funds) of period$-t$ value function of type θ^t by $-\frac{1}{pq(t)}[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))]B^Q(\theta_t) < 0$, pushing towards a higher labour wedge. Note that this effect is not due to learning-by-doing. It emerges simply due to the per-period redistribution motives of a non-Utilitarian planner.

Second, it may also have an impact on the effect of higher effective labour on future social welfare, i.e. on $\frac{\partial \hat{Q}_{t+1}(\theta_t, z_t(\theta^t),.)}{\partial z}$, for $t < T$. To simplify discussion of the latter effect, we focus again on the two-period version of our model.[48] The labour wedge in period $t = 1$ is given by

$$\zeta_1^z(\theta^t) = \frac{\hat{\delta}(\theta_1, z_1(\theta_1))}{u_c(c_1(\theta_1), z_1(\theta_1)/\theta_1)} - \frac{\partial \hat{Q}_2(\theta_1, z_1(\theta_1), .)}{p\partial z}\beta + \frac{[\mu_1(\theta_1)/p]}{f_1(\theta_1)}\frac{u_e(.)\eta(e_1(\theta_1))}{\theta_1^2} \tag{30}$$

In Appendix E, we show that the learning-by-doing component of the first-period labour wedge (i.e. the first two terms in (30)) is

$$\widehat{\zeta}_1^z(\theta_1) \equiv \frac{\hat{\delta}(\theta_1, z_1(\theta_1))}{u_c(c_1(\theta_1), z_1(\theta_1)/\theta_1)} - \frac{\partial \hat{Q}_2(\theta_1, .z(\theta_1), .)}{p\partial z}\beta =$$

$$= -(\beta/p) \times \{\int_{\underline{\theta}}^{\overline{\theta}} Q(u(c_2(\theta_2, \theta_1), z_2(\theta_2, \theta_1)/\theta_2)))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))$$

$$-Q'(u(c_1(\theta_1)), .)\int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2, \theta_1), z_2(\theta_2, \theta_1)/\theta_2))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))\}$$

$$-\beta \times \{\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)}\int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2, \theta_1), z_2(\theta_2, \theta_1)/\theta_2)\frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1}dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))$$

[48] Recall again here footnote 43.

$$+ [R\beta]^{-1} \int_{\underline{\theta}}^{\overline{\theta}} [z_2(\theta_2, \theta_1) - c_2(\theta_2, \theta_1)] \frac{\partial}{\partial z_1(\theta_1)} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))\} \tag{31}$$

This wedge has two parts. The second part is the same as under a Utilitarian planner. The first part is novel and represents the additional intertemporal redistributional concerns of a non-Utilitarian planner. Recall, however, that this incentive is also present under full information - recall (6) - and hence our discussion at Section (3.1) is still relevant.

We leave this subsection by studying the limiting labour distortions (relative to those under full-information) when $t < T$ (the case of $t = T$ - where there is no learning-by-doing - is covered by the discussions in the received literature). We restrict attention to $t = 1 < T$, but the discussion extends easily to the case of $T > t > 1$. Using (30) and (31), the labour wedge that arises due to private information (i.e. the collection of terms that are multiplied by $\mu_1(\theta^1)$) relative to $-\frac{u_e(.)}{\theta_t u_c(.)}$ is equal to

$$\widetilde{\zeta}_1^z(\theta_1) \equiv -\frac{[\mu_1(\theta_1)/p]}{\theta_1 f_1(\theta_1)} u_c(.) \times \{\eta(e_1(\theta_1))$$

$$-\beta[\frac{\theta_1^2}{u_e(.)}] \int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2, \theta_1), z_2(\theta_2, \theta_1)/\theta_2) \frac{\partial^2}{\partial z_1(\theta_1) \partial \theta_1} dF_2(\theta_2 \mid \theta_1, z_1(\theta_1))\}$$

Suppose also that $\eta(.)$ and the integral are non-zero and bounded, and that $\frac{\theta_1^2}{u_e(c_1(\theta_1), z_1(\theta_1)/\theta_1)}$ is bounded. We note from the proofs of Lemmata 1 and 2 that $\mu_1(\theta_1) = [1 - F_1(\theta_1)][B_1(\theta_1) - p^{-1} B_1^Q(\theta_1)]$ is decreasing with $\mu_1(\underline{\theta}) = 0$. Recalling (29) (with $Q' \equiv 1$ for the case of a Utilitarian planner) and (24), we then have directly that the limiting labour distortions are equal to their full-information counterparts (i.e. $\widetilde{\zeta}_1^z(\theta_1) \to 0$ as $\theta_1 \to \underline{\theta}$ and $\theta_1 \to \overline{\theta}$) if $lim_{\theta_1 \to \underline{\theta}} \frac{u_c(c_1(\theta_1), z_1(\theta_1)/\theta_1)}{\theta_1 f_1(\theta_1)} < \infty$, for the bottom end, and if $lim_{\theta_t \to \overline{\theta}} \frac{[1-F_1(\theta_1)] u_c(c_1(\theta_1), z_1(\theta_1)/\theta_1)}{\theta_1 f_1(\theta_1)} = 0$, for the top end of the skill distribution.[49]

6.2 The General Model

We now move to the case when $u_{ec} \geq 0$. Some of the details of the second-best wedges change in this case, but the basic insights are still present.

[49] This result echoes results in the received literature. In particular, the result that there are no distortions at the lowest skill level follows, in part, from our restricting attention to interior solutions, i.e. to the case when labour is positive for all interior skills (see Seade, 1977, and Ebert, 1992, for instance, for this result in a static model). Note also that under an unbounded support, and with bounded marginal utility of consumption at optimum, the limiting distortion at the top may be non-zero if $lim_{\theta_t \to \overline{\theta}} \frac{[1-F_1(\theta_1)]}{\theta_1 f_1(\theta_1)} > 0$; this is the case, for instance, when the skill distribution has a Pareto tail, where $\frac{[1-F_1(\theta_1)]}{\theta_1 f_1(\theta_1)}$ is constant (see also the discussions in Diamond, 1998, and Saez, 2001).

6.2.1 The savings tax wedge

(18) gives the *generalised* inverse Euler condition in the general model. To understand this condition, let us consider again the utility perturbation we have considered in discussing (23): an increase in period-t utility by one unit alongside a decrease in period-$t-1$ utility by β units for all workers with period-$t-1$ type θ^{t-1}. Moving to a general utility has two effects on the inverse Euler condition. These effects will be present even if there is no learning-by-doing.

First, (18) takes into account that the utility perturbation in question will have an effect on period-$t-1$ information rents associated with type θ_{t-1}. The decrease in period-$t-1$ utility by β units has a *positive* impact on period-$t-1$ current information rents associated with skill θ_{t-1} equal to $\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}{u_c(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}$ (in terms of period-t utils per worker of skill-history θ^{t-1}). The increase in period-t utility by one unit has an impact on period-$t-1$ future rents associated with skill θ_{t-1} conditional on period-t skill being x that equals $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ (in terms of period-t utils per worker of period-t skill-history (x,θ^{t-1})). The per-unit social welfare cost of period-$t-1$ utility associated with skill-history θ^{t-1} is $-\frac{\mu(\theta^{t-1})}{pq(t)f_{t-1}(\theta^{t-1}|.)}$ (in terms of period-t funds per worker of skill-history θ^{t-1}) and hence

$$-\frac{1}{pq(t)}\frac{\mu_{t-1}(\theta^{t-1})}{f_{t-1}(\theta^{t-1}\mid .)}\times\{\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}$$

$$+\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}{u_c(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}\}$$

is the social evaluation of the above net effect.

The second effect on the inverse Euler condition from moving to a non-separable per-period utility is that the conditional expectation (of the various effects of the utility rearrangement in question) over the period-t skill are taken by using the density $\widehat{f}_t(\theta_t\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))$ instead of the true density $f_t(\theta_t\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))$. In fact, note that taking the expectation of the above expression by using the density $\widehat{f}_t(\theta_t\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))$ does give the last term in (18) (which reduces to zero when $u_{ec}=0$). Moreover, the difference in the remaining terms between (18) and (23) is the use of density $\widehat{f}$ instead of f. To understand this "change of measure", note that, with non-separable utility, an increase (for given effective labour) of one unit in period-t utility of some period$-t$ type θ_t leads also to a reduction of the period$-t$ information rents associated with this type that equals $\frac{z_t(\theta^t)}{\theta_t^2}\frac{u_{ec}(c_t(\theta^t),z_t(\theta^t)/\theta_t)}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)}\equiv\delta(\theta_t)$ (recall the local IC (15)). Therefore, the necessary increase (for given earnings) in the value function of skill $\theta_t+\Delta$, with Δ very small - in order to ensure local incentive compatibility - is equal to $1-\delta(\theta_t)\Delta$.[50] Consequently, and similarly, the necessary increase in the value function of skill $\theta_t+2\Delta$ - in order to ensure local incentive compatibility - is equal to $[1-\delta(\theta_t)\Delta][1-\delta(\theta_t+\Delta)\Delta]$, and so on. Therefore, the non-separability of utility implies a gain in the necessary information rents the planner has to leave with all skills *above* θ_t in order to redistribute one util to skill θ_t. To determine this gain,

[50] This uses that for very small Δ we have $v(\theta_t+\Delta,\theta^{t-1})=v(\theta^t)+\frac{\partial}{\partial\theta_t}v(\theta_t,\theta^{t-1})\Delta$, and that, by the local IC (15) and the definition of $C(.)$, $\frac{\partial}{\partial\theta_t}v(\theta_t,\theta^{t-1})=-\frac{z_t(\theta^t)}{\theta_t^2}u_{ec}(v_t(\theta^t)-\beta\widehat{U}(\theta^t),z_t(\theta^t),\theta_t),z_t(\theta^t)/\theta_t)+\beta\widehat{V}(\theta^t)$.

discretise the interval $[\theta_t, x]$ according to $\vartheta_i = \theta_t + i\Delta$ with $i = 0, 1, ..., \frac{x-\theta_t}{\Delta} \equiv N$, and note that the necessary increase in the value function of skill $x \geq \theta_t$ - in order to ensure local incentive compatibility - is equal to $\prod_{i=0}^{N-1}[1 - \delta(\theta_t + i\Delta)\Delta]$. Using then the approximation $1 - \widetilde{\delta} = e^{-\widetilde{\delta}}$,[51] and taking the limit as $\Delta \to 0$, we have that an increase by one unit of utility of skill θ_t requires a corresponding increase of $\Gamma(x, \theta^t)$ units for any skill $x \geq \theta_t$ (recall (20)). Thus, under non-separable utility, the period-t effects we have discussed under an additively-separable utility of increasing the utility associated with skill θ_t by one util have to be "weighed" now by $\Gamma(x, \theta^t)$. Equivalently, the relative weight of any period-t type x in terms of the planner's information problem (as measured by the conditional likelihood) has to be "weighed" by $\Gamma(x, \theta^t)$, which can be thought of as a measure of the reduction in the information problem under non-separable utility. $\widehat{f}_t(. \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$ is then the "information-adjusted" density that is used by the planner to evaluate the expected effects of increasing the period-t utility of the lowest skill by one unit. [We note that the "information-adjusted" weight $\Gamma(x, \underline{\theta}, \theta^{t-1}) f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$ is divided by the total weight $\int_{\underline{\theta}}^{\overline{\theta}} \Gamma(x, \underline{\theta}, \theta^{t-1}) f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))dx$ to normalise the new weight into a density.]

A direct implication of the above, after comparing (23) and (18), is that the intertemporal wedge in period-t will be higher than the one under an additively-separable utility if and only if

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{1}{u_c(c_t(x, \theta^{t-1}), z_t(x, \theta^{t-1})/x)} [\widehat{f}_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}) - f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})] dx$$

$$-\frac{1}{pq(t)} \int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_t(x, \theta^{t-1}), z_t(x, \theta^{t-1})/x)) [\widehat{f}_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}) - f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})] dx$$

$$-\frac{1}{pq(t)} \frac{\mu(\theta^{t-1})}{f_{t-1}(\theta^{t-1} \mid .)} \times$$

$$\{ \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))/\partial \theta_{t-1}}{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} [\widehat{f}_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}) - f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})] dx$$

$$+\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2} \frac{u_{ec}(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} \} > 0$$

and vice versa. (The second term is dropped in the case of a Utilitarian planner.) This condition states that, in terms of period-t revenues, and given skill-history θ^{t-1}, the "average" net social welfare loss from marginally increasing period$-t$ private utility, which takes into account the associated effects on current revenues and social welfare and period-$t-1$ information rents, is greater than the "average" net social welfare gain from reducing period-$t-1$ utility and arises

[51] Note that the approximation error converges to zero as $\widetilde{\delta} \to 0$.

from the associated effect on period-$t-1$ information rents, where the "averages" are calculated by using the "net information-adjusted" density $[\widehat{f}_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}) - f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})]$. In other words, it states that the aforementioned "information-adjusted" average net loss is greater than the unadjusted average net loss. To the best of our knowledge this is a novel characterisation of the intertemporal wedge for general per-period private and social utility.[52]

6.2.2 The labour tax wedge

The intuition for (22) follows that of (24) and our discussion above for the use of the "information-adjusted" density $\widehat{f}$ in place of f.

To study the labour tax wedge we thus proceed in understanding (19) first. We can obtain the intuition of this equation in a similar manner to the one behind (27) and (29) in Section 6.1.2. Suppose thus, given skill history θ^{t-1}, that the planner increases marginally period-t utility of all taxpayers with period-t skill (weakly) above θ_t, while, if $t > 1$, decreasing by $[1 - F_t(\theta_t \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})]\beta$ units period-$t-1$ utility of all workers with period-$t-1$ type θ^{t-1}. The effects discussed in Section 6.1.2 of such a perturbation apply here as well, *if information rents are kept constant*. However, now, as we have already emphasised, the utility perturbation in question *does* affect information rents because of the effect of consumption on labour decisions.

In more detail, the decrease in period-$t-1$ utility, if $t > 1$, implies also an increase in period-$t-1$ information rents associated with skill θ^{t-1} by $\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}),z_t(\theta^{t-1})/\theta_{t-1})}{u_c(c_{t-1}(\theta^{t-1}),z_t(\theta^{t-1})/\theta_{t-1})}\beta$ (in terms of period-$t-1$ utils per worker with period-t skill-history $(x,\theta^{t-1}), x \geq \theta_t$) with a per-unit cost of $-\mu(\theta^{t-1})$. This entails a social welfare cost (in terms of period-t funds per worker of period-$t-1$ skill-history θ^{t-1}) equal to

$$\frac{\mu_{t-1}(\theta^{t-1})/pq(t)}{f_{t-1}(\theta^{t-1} \mid .)}\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}), z_t(\theta^{t-1})/\theta_{t-1})}{u_e(c_{t-1}(\theta^{t-1}), z_t(\theta^{t-1})/\theta_{t-1})}$$

Furthermore, the increase in period-t utility, $t \geq 1$, of period-t type θ_t does reduce period-t information rents associated with this type by $\frac{z_t(\theta^t)}{\theta_t^2}\frac{u_{ec}(c_t(\theta^t),z_t(\theta^t)/\theta_t)}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)}$, which implies, as we have already discussed, the use of "information-adjusted" weights $\Gamma(x,\theta^t)f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$ *for all* period-t types $x \geq \theta_t$. Therefore, recalling also our corresponding discussion in Section 6.1.2, the re-arrangement of utility under consideration implies the following effect conditional on period-t type x with $x \geq \theta_t$ (in terms of period-t funds)

$$\frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} - \frac{Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1}))}{pq(t)}$$

$$-[\frac{1}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} - \frac{Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x))}{pq(t)}]$$

$$+\frac{\mu(\theta^{t-1})/pq(t)}{f_{t-1}(\theta^{t-1} \mid .)}\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}), z_t(\theta^{t-1})/\theta_{t-1})}{u_e(c_{t-1}(\theta^{t-1}), z_t(\theta^{t-1})/\theta_{t-1})}$$

[52] For an alternative (but of course related) characterisation see Golosov et al. (2011).

$$+\frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1}\mid .)}\beta R\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}$$

for $t>1$, while the same expression but without the first, third and fourth terms is relevant for $t=1$. Weighing it by the "information-adjusted" weight $\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))$ and integrating over all period-t types $x\geq\theta_t$ gives that the social value of period-t value function (in terms of period-t public funds) is

$$\mu_t(\theta^t)/pq(t)=$$

$$\int_{\theta_t}^{\overline{\theta}}\{\frac{Q'(u(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x))}{pq(t)}-\frac{1}{u_c(c_t(x,\theta^{t-1}),z_t(x,\theta^{t-1})/x)}\}\Gamma(x,\theta^t)dF_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))$$

$$+\int_{\theta_t}^{\overline{\theta}}\{\frac{\beta R}{u_c(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1})}-\frac{Q'(u(c_{t-1}(\theta^{t-1}),z_{t-1}(\theta^{t-1})/\theta_{t-1}))}{pq(t)}\}\Gamma(x,\theta^t)dF_t(x\mid .)$$

$$+\int_{\theta_t}^{\overline{\theta}}\frac{\mu(\theta^{t-1})/pq(t)}{f_{t-1}(\theta^{t-1}\mid .)}\frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2}\frac{u_{ec}(c_{t-1}(\theta^{t-1}),z_t(\theta^{t-1})/\theta_{t-1})}{u_e(c_{t-1}(\theta^{t-1}),z_t(\theta^{t-1})/\theta_{t-1})}\Gamma(x,\theta^t)dF_t(x\mid .)$$

$$+\frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1}\mid .)}R\beta\int_{\theta_t}^{\overline{\theta}}\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\Gamma(x,\theta^t)dF_t(x\mid .)$$

for $t>1$, and without the last three terms for $t=1$. Recalling then (22) for when $t=1$, and noting (18) for when $t>1$, leads to (19).

We can show that the social value of the value function associated with skill-history θ^t is indeed negative under an empirically relevant assumption on the skills distribution (see also Golosov et al, 2011).[53]

Lemma 3: *Suppose that* $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ *is non-decreasing in* x*. Then, in a regular economy,* $\mu_t(\theta^t)/pq(t)<0$ *for any interior skill* θ_t *and any* t*.*

Proof: Let

$$B_t^H(\theta_t)\equiv$$

$$\int_{\underline{\theta}}^{\overline{\theta}}\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\widehat{f}_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$$

$$-\int_{\theta_t}^{\overline{\theta}}\frac{\partial f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}\frac{\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))}{\int_{\theta_t}^{\overline{\theta}}\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx}dx$$

Note now that $B_t^H(\theta_t)<0$ for interior θ_t, if $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ is non-decreasing with

[53] In fact, this assumption is satisfied by both the Lognormal and Pareto distributions, and hence by a Lognormal distribution fitted with a Pareto tail.

skill x (at optimum). Note also that $B_t^H(\underline{\theta})=0$ and hence $B_t^H(\theta_t)<0$ for any $\theta_t>\underline{\theta}$. The third term in (19) equals $-\frac{\mu_{t-1}(\theta^{t-1})/pq(t-1)}{f_{t-1}(\theta^{t-1}|.)}R\beta\int_{\theta_t}^{\overline{\theta}}\Gamma(x,\theta^t)f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$ times $B_t^H(\theta_t)$. These imply, by backward induction on (19) and $\mu_0(\theta^0)=0$ (and recalling the proofs of Lemmata 1 and 2), that $\mu(\theta^t)<0$ for any $t>0$ and any interior θ_t.■

Therefore, in a regular economy, if $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ is non-decreasing in x, we have that the lower the skill θ_t the higher $B_t^H(\theta_t)$ is, and the third term in (19) is negative. $B_t^H(\theta_t)$ can hence be thought of as a measure of the "additional" redistributional considerations of a planner when private utility is non-separable: the higher the skill the higher its labour wedge is.

We can now discuss the details of the labour wedge. We recall from condition (17) (and $\eta(e)\equiv 1+e\frac{u_{ee}(e)}{u_e(e)}$) that

$$-\frac{\partial\hat{Q}_{t+1}(\theta_t,z_t(\theta^t),.)}{\partial z_t}\beta+\frac{\mu_t(\theta^t)/pq(t)}{f_t(\theta_t\mid.)}\frac{1}{\theta_t^2}\frac{u_e(.)[\eta(e(\theta_1))-\frac{z_t(\theta^t)}{\theta_t}\frac{u_{ce}(.)}{u_c(.)}]}{\theta_1^2}=$$

$$[1+\frac{u_e(.)}{\theta_t u_c(.)}] \tag{32}$$

It differs from the corresponding formula in Section 6.1.2, (25), because the marginal benefit from lower effective labour takes now into account that consumption is also affected positively from a marginal increase in (effective) labour (for given per-period value function). Thus, it takes into account that reducing effective labour $z_t(\theta^t)$ marginally leads to an additional drop in information rents of $\frac{z_t(\theta^t)}{\theta_t^2}\frac{u_{ec}(.)u_e(.)}{\theta_t u_c(.)}$ units.[54]

Finally, moving to non-separable per-period utility may also have an impact on the effect of higher effective labour on future social welfare, i.e. on $\frac{\partial\hat{Q}_{t+1}(\theta_t,z_t(\theta^t),.)}{\partial z_t}$, for $t<T$, and thereby on the labour tax wedge. To simplify the discussion, we focus hereafter on the two-period version of our model.[55] The labour tax wedge in period $t=1$ is given by

$$\zeta_1^z(\theta^t)=\frac{\hat{\delta}(\theta_1,z_1(\theta_1))}{u_c(c_1(\theta_1),z_1(\theta_1)/\theta_1)}-\frac{\partial\hat{Q}_2(\theta_1,z_1(\theta_1),.)}{\partial z}\frac{\beta}{p}+\frac{[\mu_1(\theta^1)/p]}{f_1(\theta_1)}\frac{u_e(.)[\eta(e(\theta_1))-\frac{z_t(\theta^t)}{\theta_t}\frac{u_{ce}(.)}{u_c(.)}]}{\theta_1^2} \tag{33}$$

In Appendix E, we show that the learning-by-doing component of the first-period labour wedge is

$$\widehat{\zeta}_1^z(\theta_1)\equiv\frac{\hat{\delta}(\theta_1,z_1(\theta_1))}{u_c(c_1(\theta_1),z_1(\theta_1)/\theta_1)}-\frac{\partial\hat{Q}_2(\theta_1,z_1(\theta_1),.)}{\partial z}\frac{\beta}{p}=$$

$$=-\beta\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)}\frac{z_1(\theta_1)u_{ec}(c_1(\theta_1),.)}{\theta_1^2u_c(c_1(\theta_1),.)}\int_{\underline{\theta}}^{\overline{\theta}}u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))$$

[54] Of course, this difference is not due to learning-by-doing and has also been highlighted in the received literature. For expressing the term in the curly brackets in terms of labour elasticities (for given consumption and savings) see Saez (2001) and Golosov et al. (2011).

[55] For the case of more than two periods, see again footnote43.

$$-(\beta/p)\times\{\int_{\underline{\theta}}^{\overline{\theta}} Q(u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2)))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))$$

$$-Q'(u(c_1(\theta_1)),.)\int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))\}$$

$$-\beta\times\{\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)}\int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2))\frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1}dF_2(\theta_2\mid\theta_1,z_1(\theta_1)))$$

$$+[R\beta]^{-1}\int_{\underline{\theta}}^{\overline{\theta}}[z_2(\theta_2,\theta_1)-c_2(\theta_2,\theta_1)]\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))\} \tag{34}$$

This has three parts. The second and third parts are the same as under an additively-separable per-period private utility, and have been discussed in Section 6.1.2. The first part is novel. To understand this term, recall from the corresponding discussion in Section 6.1.2 that a marginal increase in period-t (effective) labour, with $t=1$ here, makes, under net skill persistence, high future skills more likely, which raises expected discounted continuation utility of consumers of first-period type θ_1. The latter allows the planner to generate more revenues in period $t=1$ by reducing first-period utility by $\beta\int_{\underline{\theta}}^{\overline{\theta}} u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))$ units. With a non-separable private utility, this reduction in first-period utility (for given earnings and promised utility) leads also to a per-unit increase of $\frac{z_1(\theta_1)u_{ec}(c_1(\theta_1),.)}{\theta_1^2u_c(c_1(\theta_1),.)}$ units in first-period information rents associated with skill θ_1. The per-unit social valuation of the latter in terms of first-period funds per worker with skill θ_1 is $\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)}$, and thus the discounted effect on future social welfare of a marginal increase in first-period (effective) labour includes now also the negative of the first term in the learning-by-doing component of the labour wedge above. With negative social value of first-period value function (recall Lemma 3), the novel first term in 34 is positive, pushing the learning-by-doing component of the first-period labour wedge upwards.[56]

7 Numerical Solution

Here we investigate numerically the case of additively-separable utility with iso-elastic labour supply to help us improve our understanding of the effects of learning-by-doing on optimal dynamic redistribution, and in particular on the path over time of (average) cross-sectional distortions. We also compute the welfare loss from resorting instead to a linear (either time-independent or time-dependent) tax and compare it with the one derived in the received literature.

TO BE COMPLETED

[56] One can also easily adapt our discussion of the limiting labour distortions (relative to full-information) in the end of the previous subsection for the case of non-separable utility.

8 Conclusions

Focusing on a Markov environment with standard time-separable preferences, and assuming pre-commitment on the part of a planner with per-period redistributional motives, we have investigated optimal dynamic redistribution under learning-by-doing. We have first derived an envelope condition that must characterise incentive-compatible allocations. This condition is similar to the one derived recently in Kapicka (2008) and Pavan et al. (2009), though proving it in our context requires some additional steps. Echoing the approaches in Kapicka (2008), Pavan et al (2009), and more recently in Farhi and Werning (2011) and Golosov et al. (2011), we then derive the optimal allocations from the perspective of the planner, if this envelope condition was fully characterising incentive-compatible allocations. We also derive various conditions under which this "first-order approach" is a valid means of characterising second-best allocations. The validity of the first-order-approach and the implied Relaxed problem faced by the planner is also confirmed in our numerical analysis. Given our discussion on sufficiency conditions, we then move to discuss the allocations, and the associated tax wedges, that are characterised by the Relaxed problem of the planner.

We have shown that the second-best intertemporal wedge will be higher (resp. lower), all other things equal, to that set by a Utilitarian planner (and analysed by the New Dynamic Public Finance Literature), if at optimum the (conditional) expected marginal social welfare is decreasing (resp. increasing) over time, for given current skill. The reason is that in this case, current consumption is valued more (resp. less) socially. Moreover, in a regular economy with net skill persistence, a non-Utilitarian planner sets a higher labour wedge, across all skills, all other things equal.

We have also identified a new intertemporal trade-off between current and future information rents. As in the static model, a positive labour wedge on some skill level reduces current information rents of higher skills (at the margin). However, under learning-by-doing, such a wedge leads also to an increase in future information rents from mimicking this skill, when skill and effective labour are substitutes in making future high skills more likely. This trade-off implies that, for all skills above a certain (small) threshold, the labour distortion (relative to full-information) is more progressive under learning-by-doing, all other things equal, when the distribution of skills is the Lognormal fitted at the modal skill level with a Pareto tail with sufficiently high Pareto index, and the mean of the Lognormal is increasing in current skill and labour with negative cross-derivative. The latter captures an environment where the positive returns from experience are higher for low skill levels, which is consistent with empirical evidence.

We, finally, identify how the above effects change when private utility is non-separable. For instance, under an empirically relevant condition on the skill distribution, and substitutability between leisure and consumption, learning-by-doing pushes, all other things equal, the labour wedge upwards for all interior skills.

There are a number of ways our analysis can be generalised, but we leave these for future work. First, one could allow for skill distributions with moving support. Second, one might want

to discuss optimal redistribution in an infinite-horizon model. More interesting extensions, in our opinion, would be to allow for private investment in human capital (along the lines of Kapicka (2010)), or for endogenous public investment in human capital. Lastly, a very interesting and challenging extension would be to allow for no pre-commitment on the part of the planner, where the planner may be tempted to exploit information that it has learned about agents in earlier periods.[57]

9 References

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

Atkeson, A. and R. E. Lucas, Jr. (1992): "On Efficient Distribution with Private Information", Review of Economic Studies, 59, 427-453.

Balle, S., and A. Spadaro, (2006): "Optimal Nonlinear Labor Income Taxation in Dynamic Economies", Mimeo, IMEDEA and PSE.

Bartle, R.G., (2000): A Modern Theory of Integration, Graduate Studies in Mathematics, American Mathematical Society.

Battaglini, M., and S. Coate, (2008): "Pareto Efficient Income Taxation with Stochastic Abilities", Journal of Public Economics, 92(3-4), 844-868.

Berliant, M., and J. Ledyard, (2004): "Dynamic Optimal Nonlinear Income Taxation with No Commitment," Washington University, mimeo.

Chamley, C., (1986): "Optimal Taxation of Capital in General Equilibrium", Econometrica, 54, 607–622.

Chari, V.V., and P.J. Kehoe, (1999): "Optimal fiscal and monetary policy", Handbook of Macroeconomics, in: J. B. Taylor J.B, and M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 26, pp. 1671-1745 Elsevier.

Diamond, P.A., and J.A. Mirrlees, (1978): "A Model of Social Insurance with Variable Retirement", Journal of Public Economics 10(3), 295–336.

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.

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

Ebert, Udo, (1992): "A Reexamination of the Optimal Nonlinear Income Tax", Journal of Public Economics, 49(1), 47-73.

Farhi, E., and I. Werning, (2010): "Insurance and Taxation over the Life Cycle", Mimeo, Harvard and MIT.

Fernandes, A. and C. Phelan, (2000): "A Recursive Formulation for Repeated Agency with History Dependence", Journal of Economic Theory, 91 (2), 223{247.

[57] For related discussions see also Berliant and Ledyard (2004), Sleet and Yeltekin (2005), Battaglini and Coate, (2008) and Krause (2009).

Golosov, M., N. Kocherlakota, and A. Tsyvinski, (2003): "Optimal Indirect and Capital Taxation", Review of Economic Studies, 70(3), 569-587.

Golosov, M., M. Troshkin, and A. Tsyvinski, (2010): "Optimal Dynamic Taxes", Mimeo, Yale University.

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

Green, E. J., (1987): "Lending and the Smoothing of Uninsurable Income" in Contractual Arrangements for Intertem- poral Trade, eds E. Prescott and N. Wallace, University of Minnesota Press, Minneapolis.

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

Judd, K. L., (1985): "Redistributive Taxation in a Simple Perfect Foresight Model", Journal of Public Economics, 28, 5983.

Kapicka, M., (2008): "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach", University of California, Santa Barbara, mimeo.

Kapicka, M., (2010): "The Dynamics of Optimal Taxation when Human Capital is Endogenous", , University of California, Santa Barbara, mimeo.

Kapicka, M., (20011): "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach", University of California, Santa Barbara, mimeo.

Kocherlakota, N. R., (2005): "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation", Econometrica, 73(5), 1587-1621.

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

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

La France, J. T., and L. D. Barney, (1991): "The Envelope Theorem in Dynamic Optimization", Journal of Economic Dynamics and Control, 15, 355-385.

Milgrom, P., and I. Segal, (2002): "Envelope Theorems for Arbitrary Choice Sets", Econometrica, 70(2), 583.601.

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

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

Piketty, T., (1995): "Social Mobility and Redistributive Politics", Quarterly Journal of Economics, 110(3), 551- 584.

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

Thomas, J. and T. Worrall, (1990): "Income Fluctuations and Asymmetric Information: An Example of the Repeated Principal Agent Problem", Journal of Economic Theory, 51, 367-390.

Topel, R., (1991): "Specific Capital, Mobility and Wages: Wages Rise with Seniority", Journal

of Political Economy, 99(t), 145-176.

Seade, J.K., (1977): "On the Shape of Optimal Tax Schedules", Journal of Public Economics, 7, pp. 203-236.

Sleet, C., and S. Yeltekin, (2004): "Credible Social Insurance, " Carnegie-Mellon, mimeo.

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.

10 Appendices

10.1 Appendix A: Necessary Condition for IC

The proof for the case of Assumption AF6(i) (and not AF6(ii)-(iv)) holding can be found in Pavan et al. (2009).

The proof for the case of Assumptions AF6(ii)-AF6(iv) (instead of AF6(i)) holding combines steps from Pavan et al. (2009) and Kapicka (2008) as follows:

We will be using the following lemmas (we present their proofs for completeness).

Lemma A1 (Pavan et al. (2009)): For any (locally) Lipschitz real-valued function $\Gamma(\theta)$ on Θ,

$$\int \Gamma(\theta) dF_t(\theta \mid \theta_{t-1}, z) - \int \Gamma(\theta) dF_t(\theta \mid \eta_{t-1}, z) = -\int \Gamma'(\theta)[F(\theta \mid \theta_{t-1}, z) - F(\theta \mid \eta_{t-1}, z)]d\theta$$

(where the integrals are over (an arbitrary open subset of) Θ.

Proof: Note first that the first two integrals exist because (by Lipschitz continuity) $\mid \Gamma(\theta) \mid \leq \Gamma_1 + \Gamma_2 \mid \theta \mid$ for some non-negative constants Γ_1 and Γ_2, and assumption AF1. Moreover, $\Gamma(\theta)[F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z)]$ is a.e. differentiable in θ on Θ (by AF2 and (local) Lipschitz continuity of $\Gamma(\theta)$). We can thus use (Henstock-Kurzweil) integration by parts (see Bartle, 2000, Theorems 4.11 and 4.7), and the desired result follows directly if $\overline{\theta}$ is finite. If, on the other hand, $\overline{\theta} \to \infty$, note that

$$\begin{aligned}
&\mid \Gamma(\theta)[F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z)] \mid \\
&\leq \Gamma_1 \mid F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z) \mid \\
&+\Gamma_2 \theta \mid F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z) \mid \\
&= \Gamma_1 \mid F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z) \mid \\
&+\Gamma_2 \theta \mid \int_{\vartheta \geq \theta} dF_t(\vartheta \mid \eta_{t-1}, z) - \int_{\vartheta \geq \theta} dF_t(\vartheta \mid \theta_{t-1}, z) \mid \\
&\leq \Gamma_1 \mid F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z) \mid
\end{aligned}$$

$$+\Gamma_2\theta[\int_{\vartheta\geq\theta} dF_t(\vartheta \mid \eta_{t-1}, z) + \int_{\vartheta\geq\theta} dF_t(\vartheta \mid \theta_{t-1}, z)]$$
$$\leq \Gamma_1 \mid F_t(\theta \mid \theta_{t-1}, z) - F_t(\theta \mid \eta_{t-1}, z) \mid$$
$$+\Gamma_2[\int_{\vartheta\geq\theta} \vartheta dF_t(\vartheta \mid \eta_{t-1}, z) + \int_{\vartheta\geq\theta} \vartheta dF_t(\vartheta \mid \theta_{t-1}, z)]$$

Clearly then (recalling also the continuity of $F_t(\theta \mid \theta_{t-1}, z)$ in θ_{t-1} by AF5) both terms to the right of the above inequality vanish as $\theta \to \infty$, and the required result follows directly.■

Lemma A2 (Kapicka (2008)): Under AF6(ii), for any (locally) Lipschitz real-valued function $\Gamma(\theta)$ on Θ

$$lim_{\eta\to 0} \int \Gamma'(\theta) \frac{F_t(\theta \mid \theta' + \eta, z) - F_t(\theta \mid \theta', z)}{\eta} d\theta = \int \Gamma'(\theta) \frac{\partial F_t(\theta \mid \theta', z)}{\partial \theta'} d\theta$$

where the integrals are over (an arbitrary open subset of) Θ.

Proof: If $\frac{\partial F_t(\theta|.,z)}{\partial \theta_{t-1}}$ is non-increasing then $F_t(\theta \mid ., z)$ is concave and hence for any scalar $\alpha \in (0,1)$ we have for $\eta > 0$

$$\alpha F_t(\theta \mid \theta' + \eta, z) + (1-\alpha) F_t(\theta \mid \theta', z) \leq F_t(\theta \mid \theta' + \alpha\eta, z)$$

This can be rewritten (after rearranging terms and dividing both sides with $\alpha\eta$) as

$$\frac{F_t(\theta \mid \theta' + \eta, z) - F_t(\theta \mid \theta', z)}{\eta} \leq \frac{F_t(\theta \mid \theta' + \alpha\eta, z) - F_t(\theta \mid \theta', z)}{a\eta}$$

Thus, $\frac{F_t(\theta|\theta'+x,z) - F_t(\theta|\theta',z)}{x}$ is non-increasing in x. Similarly, if $\frac{\partial F_t(\theta|.,z)}{\partial \theta_{t-1}}$ is non-decreasing then $F_t(\theta \mid ., z)$ is convex and hence for any scalar $\alpha \in (0,1)$ we have for $\eta < 0$

$$\alpha F_t(\theta \mid \theta', z) + (1-\alpha) F_t(\theta \mid \theta' + \eta, z) \geq F_t(\theta \mid \theta' + (1-\alpha)\eta, z)$$

This can be rewritten (after rearranging terms and dividing both sides with $(1-\alpha)\eta$) as

$$\frac{F_t(\theta \mid \theta' + \eta, z) - F_t(\theta \mid \theta', z)}{\eta} \leq \frac{F_t(\theta \mid \theta' + (1-\alpha)\eta, z) - F_t(\theta \mid \theta', z)}{(1-a)\eta}$$

Thus, $-\frac{F_t(\theta|\theta'+x,z) - F_t(\theta|\theta',z)}{x}$ is non-increasing in x.

Note also that by assumption $lim_{\eta\to 0} \frac{F_t(\theta|\theta'+\eta,z) - F_t(\theta|\theta',z)}{\eta} = \frac{\partial F_t(\theta|\theta',z)}{\partial \theta'}$. The Monotone Convergence Theorem then implies that

$$lim_{\eta\to 0} \int \Gamma'(\theta) \frac{F_t(\theta \mid \theta' + \eta, z) - F_t(\theta \mid \theta', z)}{\eta} d\theta = \int \Gamma'(\theta) lim_{\eta\to 0} \frac{F_t(\theta \mid \theta' + \eta, z) - F_t(\theta \mid \theta', z)}{\eta} d\theta$$

and the result follows from the definition of the derivative.■

We can proceed now in deriving the Necessary Conditions for IC.

Following Pavan et al. (2009), let

$$\Psi_t(\theta^t, m^t, y^{t-1}) =$$

$$u(c_t(m^t, y^{t-1}), z_t(m^t, y^{t-1})/\theta_t) + \beta \int W_{t+1}(\theta^{t+1}, m^t, y^{t-1}))dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(m^t, y^{t-1}))$$

$$W_{t+1}(\theta^{t+1}, m^t, y^{t-1})) = V_{t+1}(\theta^{t+1}, m^t, (y_t(m^t, y^{t-1}), y^{t-1}))$$

$$V_{t+1}(\theta^{t+1}, m^t, y^t) = sup_{m_{t+1}} \Psi_{t+1}(\theta^{t+1}, (m_{t+1}, m^t), y^t)$$

for any history $h_t = (\theta^t, m^{t-1}, y^{t-1})$ and any IC-h_t mechanism. To understand these definitions, $\Psi_t(\theta^t, m^t, y^{t-1})$ is the expected continuation utility (in terms of period-t utils) given history $(\theta^t, m^{t-1}, y^{t-1})$ from reporting m_t in period t and from period-$t+1$ (inclusive) and onwards choosing reports optimally. $V_{t+1}(\theta^{t+1}, m^t, y^t)$ is the expected continuation utility (in terms of period-$t+1$ utils) given history (θ^{t+1}, m^t, y^t) from submitting the optimal reports from period-$t+1$ (inclusive) and onwards.

The proof then proceeds in a number of steps.

First Step: We now show that

$$V_t(\theta^t, m^{t-1}, (y_{t-1}(m^{t-1}, y^{t-2}), y^{t-2}))$$

is (locally) equi-Lipschitz continuous in θ^t for any t, if $\Psi_t(\theta^t, m^t, y^{t-1})$ is (locally) equi-Lipschitz continuous in θ^t for any t.

To see this note that for any m^{t-1} and y^{t-1}:

$$\mid V_t(\eta^t, m^{t-1}, y^{t-1}) - V_t(\theta^t, m^{t-1}, y^{t-1}) \mid$$

$$\leq sup_{m_t} \mid \Psi_t(\eta^t, (m_t, m^{t-1}), y^{t-1}) - \Psi_t(\theta^t, (m_t, m^{t-1}), y^{t-1}) \mid \leq M \parallel \eta^t - \theta^t \parallel$$

where M is the Lipschitz constant of $\Psi_t(\theta^t, m^t, y^{t-1})$. Thus, $V_t(\theta^t, m^{t-1}, y^{t-1})$ is (locally) equi-Lipschitz continuous in θ^t. Clearly, then, this implies the required result after setting $y^{t-1} = (y_{t-1}(m^{t-1}, y^{t-2}), y^{t-2})$.

Second Step: We now show by induction that $\Psi_t(\theta^t, m^t, y^{t-1})$ is locally equi-Lipschitz continuous in θ^t for any t.

Note first that $\Psi_{T+1}(\theta^{T+1}, m^{T+1}, y^T) = 0$, which is (trivially) equi-Lipschitz continuous (and hence locally equi-Lipschitz continuous) in θ^{T+1}. So, suppose that $\Psi_{t+1}(\theta^{t+1}, m^{t+1}, y^t)$ is locally equi-Lipschitz continuous in θ^{t+1} for any t. Let $\hat{u}(y_t(m^t, y^{t-1}), \theta) = u(c_t(m^t, y^{t-1}), z_t(m^t, y^{t-1})/\theta)$. Note that, after using the definition of Ψ,

$$\mid \Psi_t(\eta^t, m^t, y^{t-1}) - \Psi_t(\theta^t, m^t, y^{t-1}) \mid =$$

$$\frac{\hat{u}(y_t(m^t, y^{t-1}), \eta_t) - \hat{u}(y_t(m^t, y^{t-1}), \theta_t)}{\eta_t - \theta_t}$$

$$+ \mid \beta \int W_{t+1}((\theta, \eta^t), m^t, y^{t-1}) dF_{t+1}(\theta \mid \eta_t, z_t(m^t, y^{t-1}))$$

$$-\beta \int W_{t+1}((\theta, \theta^t), m^t, y^{t-1}) dF_{t+1}(\theta \mid \theta_t, z_t(m^t, y^{t-1})) \mid$$

$$\leq \frac{\hat{u}(y_t(m^t, y^{t-1}), \eta_t) - \hat{u}(y_t(m^t, y^{t-1}), \theta_t)}{\eta_t - \theta_t}$$

$$+ \mid \beta \int [W_{t+1}((\theta, \eta^t), m^t, y^{t-1}) - W_{t+1}((\theta, \theta^t), m^t, y^{t-1})] dF_{t+1}(\theta \mid \eta_t, z_t(m^t, y^{t-1})) \mid$$

$$+ \mid \beta \int W_{t+1}((\theta, \theta^t), m^t, y^{t-1}) d[F_{t+1}(\theta \mid \eta_t, z_t(m^t, y^{t-1})) - F_{t+1}(\theta \mid \theta_t, z_t(m^t, y^{t-1}))] \mid$$

$$\leq N \mid \eta_t - \theta_t \mid + \beta M \parallel \eta^t - \theta^t \parallel$$

$$+\beta \int [K + M\theta] \mid d[F_{t+1}(\theta \mid \eta_t, z_t(m^t, y^{t-1})) - F_{t+1}(\theta \mid \theta_t, z_t(m^t, y^{t-1}))] \mid$$

$$\leq N \mid \eta_t - \theta_t \mid + \beta M \parallel \eta^t - \theta^t \parallel + \beta ML \mid \eta_t - \theta_t \mid$$

$$\leq [N + \beta(M + ML)] \parallel \eta^t - \theta^t \parallel$$

where the second inequality follows from (a) AU1 with N being the corresponding Lipschitz constant and (b) the First Step (recalling $W_{t+1}(\theta^{t+1}, m^t, y^{t-1})) = V_{t+1}(\theta^{t+1}, m^t, (y_t(m^t, y^{t-1}), y^{t-1}))$), with M being the Lipschitz constant, which, in turn, implies that $\mid W_{t+1}((\theta_{t+1}, \theta^t), m^t, y^{t-1})) \mid \leq K + M\theta_{t+1}$ for some non-negative constant K. Moreover, the third inequality follows from AF6(iii) where L is the relevant Lipschitz constant. Thus, $\Psi_t(\theta^t, m^t, y^{t-1})$ is also locally equi-Lipschitz continuous in θ^t for any t, and the proof is complete.

Third Step: Here we show that

$$\frac{\partial \Psi_t(\theta^t, m^t, y^{t-1})}{\partial \theta_t} = -\frac{z_t(m^t, y^{t-1})}{\theta_t^2} u_e(c_t(m^t, y^{t-1}), z_t(m^t, y^{t-1})/\theta_t)$$

$$-\beta \int \frac{\partial W_{t+1}(\theta^{t+1}, m^t, y^{t-1})}{\partial \theta_{t+1}} [\frac{\partial F_{t+1}(\theta_{t+1} \mid \theta_t, z_t(m^t, y^{t-1}))}{\partial \theta_t}] d\theta_{t+1}$$

with $\frac{\partial W_{T+1}(\theta^{T+1}, m^T, y^{T-1})}{\partial \theta_{T+1}} \equiv 0$.

We have by definition of Ψ, and the structure of our environment,

$$\frac{\Psi_t((\eta_t, \theta^{t-1}), m^t, y^{t-1}) - \Psi_t((\theta_t, \theta^{t-1}), m^t, y^{t-1})}{\eta_t - \theta_t} =$$

$$\frac{\hat{u}(y_t(m^t, y^{t-1}), \eta_t) - \hat{u}(y_t(m^t, y^{t-1}), \theta_t)}{\eta_t - \theta_t}$$

$$+\beta \int \frac{W_{t+1}((\theta,\theta^t),m^t,y^{t-1})}{\eta_t-\theta_t} d[F_{t+1}(\theta \mid \eta_t, z_t(m^t,y^{t-1})) - F_{t+1}(\theta \mid \theta_t, z_t(m^t,y^{t-1}))]$$

Recall now that $V_{t+1}(\theta^{t+1},m^t,(y_t(m^t,y^{t-1}),y^{t-1})) = W_{t+1}(\theta^{t+1},m^t,y^{t-1})$. Fix ϑ,ϑ' in Θ with $\vartheta<\vartheta'$ and take $\theta_t,\eta_t \in (\vartheta,\vartheta')$. By the Second Step we have that $\Psi_t(\theta^t,m^t,y^{t-1})$ is locally equi-Lipschitz continuous in θ^t for any t, and hence almost everywhere differentiable with bounded derivative in θ_t on (ϑ,ϑ') for any $(\theta^{t-1},m^t,y^{t-1})$. Recalling the First and Second Steps, we have that $W_{t+1}(\theta^{t+1},m^t,y^{t-1})$ is locally equi-Lipschitz continuous in θ^{t+1} for any t, and hence almost everywhere differentiable with bounded derivative in θ_{t+1} for all (θ^t,m^t,y^{t-1}). Using Lemma A1, taking the limits at both sides as $\eta_t \to \theta_t$, using Lemma A2 and AU1, and then taking the limits as $\vartheta \to \underline{\theta}$ and $\vartheta' \to \overline{\theta}$ we have the desired result.

Fourth Step: In an IC mechanism Ω,

$$\frac{\partial \Psi_t(\theta^t,\theta^t,y^{t-1})}{\partial \theta_t} = \frac{\partial W_t(\theta^t,\theta^{t-1},y^{t-2})}{\partial \theta_t}$$

By IC we have $W_t((\theta_t,\theta^{t-1}),\theta^{t-1},y^{t-2}) = \Psi_t((\theta_t,\theta^{t-1}),\theta^t,y^{t-1})$. By the definitions of W and Ψ we also have $W_t((\eta_t,\theta^{t-1}),\theta^{t-1},y^{t-2}) \geq \Psi_t((\eta_t,\theta^{t-1}),\theta^t,y^{t-1})$. Thus, we have

$$W_t((\eta_t,\theta^{t-1}),\theta^{t-1},y^{t-2}) - W_t(\theta^t,\theta^{t-1},y^{t-2}) \geq$$

$$\Psi_t((\eta_t,\theta^{t-1}),\theta^t,y^{t-1}) - \Psi_t(\theta^t,\theta^t,y^{t-1})$$

Fix ϑ,ϑ' in Θ such that $\vartheta<\vartheta'$ and take $\theta_t \in (\vartheta,\vartheta')$. Recall $V_t(\theta^t,\theta^{t-1},(y_{t-1}(\theta^{t-1},y^{t-2}),y^{t-2})) = W_t(\theta_t,\theta^{t-1},y^{t-2})$. By the First and Second Steps we have that $\Psi_t(\theta^t,\theta^t,y^{t-1})$ and $W_t(\theta^t,\theta^{t-1},y^{t-2})$ are almost everywhere differentiable with bounded derivative in θ_t on (ϑ,ϑ'). Dividing thus the above inequality by $\eta_t-\theta_t$ (with $\eta_t \in (\vartheta,\vartheta')$), and taking the limits at both sides as $\eta_t \to \theta_t$, and then taking the limits as $\vartheta \to \underline{\theta}$ and $\vartheta' \to \overline{\theta}$ we have the desired result.

The Necessary conditions follow from (a) noting, by IC, that $v(\theta^t,\theta^{t-1},y^{t-1}) = W_t(\theta^t,\theta^{t-1},y^{t-2})$ for any history truthful history $(\theta^t,\theta^{t-1},y^{t-1})$ and any t, (b) the First Step, and (c) the Third (after setting $m^t=\theta^t$) and Fourth steps, and using forward iteration along a truth-telling path.

We then have the following result:

Lemma A3: Under AF3, AF6(iii) and AF6(iv), the local incentive-compatibility constraint is

$$\frac{\partial v(\theta^t,\theta^{t-1},y^{t-1})}{\partial \theta_t} = -\frac{z_t(\theta^t)}{\theta_t^2} u_e(c_t(\theta^t),z_t(\theta^t)/\theta_t)$$

$$+\beta \int \frac{\partial^2 F_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial \theta_t \partial \theta_{t+1}} v(\theta^{t+1},\theta^t,y^t) d\theta_{t+1}$$

Proof: Note first that, due to $\mid v((\theta,\theta^t),\theta^t,y^t) \mid \leq V_1 + V_2 \mid \theta \mid$ for some non-negative constants V_1 and V_2 (which follows from Lipschitz continuity of $v(\theta^{t+1},\theta^t,y^t)$) and Assumption AF6(iii) (which implies that the conditional average skill has an essentially bounded partial derivative

with respect to past skill on every open compact subset of Θ), the integral

$$\int \frac{\partial^2 F_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))}{\partial \theta_t \partial \theta_{t+1}} v(\theta^{t+1}, \theta^t, y^t) d\theta_{t+1}$$

exists for every θ_t in every open compact subset of Θ. Moreover, $\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, z_t(\theta^t))}{\partial \theta_t} v(\theta^{t+1}, \theta^t, y^t)$ is differentiable almost everywhere in θ_{t+1} (by AF3 and Lipschitz continuity of $v(\theta^{t+1}, \theta^t, y^t)$ in θ_{t+1}). We can thus use (Henstock-Kurzweil) integration by parts (see Bartle, 2000, Theorems 4.11 and 4.7) over an arbitrary open compact subset of Θ, implying that the above integral is equal to

$$-\int \frac{\partial v((\theta, \theta^t), \theta^t, y^t)}{\partial \theta} \frac{\partial F_{t+1}(\theta \mid \theta_t, z_t(\theta^t)))}{\partial \theta_t} d\theta$$

if $\overline{\theta}$ is finite. For the case of $\overline{\theta} \to \infty$, we have

$$\mid v((\theta, \theta^t), \theta^t, y^t) \frac{\partial F_{t+1}(\theta \mid \theta_t, z_t(\theta^t))}{\partial \theta_t} \mid$$

$$\leq V_1 \mid \frac{\partial F_{t+1}(\theta \mid \theta_t, z_t(\theta^t))}{\partial \theta_t} \mid$$

$$+V_2 \theta \mid \frac{\partial F_{t+1}(\theta \mid \theta_t, z_t(\theta^t))}{\partial \theta_t} \mid$$

Thus, recalling AF6(iv), $lim_{\theta \to \infty} \mid v((\theta, \theta^t), \theta^t, y^t) \frac{\partial F_{t+1}(\theta|\theta_t, z_t(\theta^t))}{\partial \theta_t} \mid = 0$, and the result follows directly after recalling the local incentive-compatibility constraint (9).■

10.2 Appendix B: Sufficiency Conditions for IC

Recall that $J_t^\tau(\theta^\tau, z^{\tau-1}) = -\prod_{k=t+1}^{\tau} \frac{\partial F_k(\theta_k|\theta_{k-1}, z_{k-1})/\partial \theta_{k-1}}{f_k(\theta_k|\theta_{k-1}, z_{k-1})}$.

Let $\widehat{u}(y_\tau, \theta_\tau) \equiv u(c_\tau, z_\tau/\theta_\tau)$. Define $D^\Omega(h) \equiv E^{\mu[\Omega]|h}\left[\sum_{\tau=t}^{T} J_t^\tau(\tilde{\theta}^\tau, \tilde{z}^{\tau-1}) \beta^{\tau-t} \frac{\partial \widehat{u}(\tilde{y}_\tau, \tilde{\theta}_\tau)}{\partial \theta_\tau}\right]$ for **any** period$-t$ history h, and note that $D(m_t, (\theta^t, \theta^{t-1}, y^{t-1})) = D^\Omega(\theta^t, (m_t, \theta^{t-1}), y^{t-1})$.

For any history $((\theta_t, \theta^{t-1}), m^{t-1}, y^{t-1})$,

$$E^{\mu[\Omega]|((\theta_t, \theta^{t-1}), m^{t-1}, y^{t-1})}\left[X(\tilde{\theta}, \tilde{y})\right] = E^{\mu[\Omega]|((\theta_t, \theta^{t-1}), (\theta_t, m^{t-1}), y^{t-1})}\left[X(\tilde{\theta}, \tilde{y})\right]$$

and so $D^\Omega(\theta^t, \theta^{t-1}, y^{t-1}) = D^\Omega(\theta^t, (\theta_t, \theta^{t-1}), y^{t-1})$.

Denote with $\widehat{\Psi}_t(\theta, m)$ the period-t expected discounted utility in Ω from one-stage deviation m, i.e. the period-t expected discounted utility in Ω given history $((\theta, \theta^{t-1}), (m, \theta^{t-1}), y^{t-1})$.

We then have:

Proposition A1 (Pavan et al., 2009): *Ω is IC at any truthful period–t history, if, for any truthful history $(\theta^t, \theta^{t-1}, y^{t-1})$ (referred simply as history θ^t for the remainder of this Appendix) and any t,*

1. *$D^{\Omega}(\theta^t, (m_t, \theta^{t-1}), y^{t-1})$ is non-decreasing in the report m_t*
2. *$\widehat{\Psi}_t(\theta_t, \theta_t)$ is Lipschitz continuous in θ_t, and $\frac{\partial \widehat{\Psi}_t(\theta_t, \theta_t)}{\partial \theta_t} = D^{\Omega}(\theta^t, \theta^t, y^{t-1})$ for almost all θ_t.*

Proof: Fix truthful history θ^t. Note first that in a Markov environment, Ω is IC at any period-$t+1$ history $(\theta^{t+1}, (m, \theta^{t-1}), (y_t(m, \theta^{t-1}), y^{t-1}))$. The reason is that allocations depend only on current and past reports, and so incentives depend only on current skill (which affects utility), and current and past reports, but not directly on past skills. Formally, if Ω is not IC at history $(\theta^{t+1}, (m, \theta^{t-1}), (y_t(m, \theta^{t-1}), y^{t-1}))$, then Ω is also not IC at truthful history $\hat{\theta}^{t+1} = (\theta_{t+1}, m_t, \theta^{t-1})$.

Thus, **suffices** to check **one-stage** deviations. Let Ω' be the mechanism that differs from Ω only in that ignores any report in period t and has $y^{\Omega'}(., \theta^{t-1}) = y^{\Omega}(m, \theta^{t-1})$. Since Ω is IC at any period-$t+1$ history $(\theta^{t+1}, (m, \theta^{t-1}), (y_t(m, \theta^{t-1}), y^{t-1}))$, the same is true for Ω'. This implies that Ω' is IC at the truthful period-t history θ^t for any m (recall that the report in period t is ignored). Thus $\widehat{\Psi}_t(\theta, m)$, for any m, is **also** the period-t value function under Ω' given truthful history θ^t, and hence, due to the necessary condition for IC, we have that $\widehat{\Psi}_t(\theta, m)$ is Lipschitz continuous, for any m, in θ and such that a.e. $\frac{\partial \widehat{\Psi}_t(\theta, m)}{\partial \theta} = D^{\Omega'}((\theta, \theta^{t-1}), \theta^{t-1}, y^{t-1})$.

But note that $D^{\Omega'}((\theta, \theta^{t-1}), \theta^{t-1}, y_{t-1}) = D^{\Omega}((\theta, \theta^{t-1}), (m, \theta^{t-1}), y_{t-1})$ by construction. Note next that $\widehat{\Psi}_t(\theta_t, \theta_t)$ is (by 2.) Lipschitz continuous in θ_t and satisfies the local IC (the equality in 2.), and so

$$\widehat{\Psi}_t(\theta, \theta) - \widehat{\Psi}_t(\theta, m) =$$

$$\int_m^{\theta} [D^{\Omega}((s, \theta^{t-1}), (s, \theta^{t-1}), y^{t-1}) - D^{\Omega}((s, \theta^{t-1}), (m, \theta^{t-1}), y^{t-1})] ds$$

This is non-negative by 1. above. Therefore, (one-stage) deviations are not profitable and the mechanism Ω is incentive compatible.■

We leave this Appendix by noting that $\widehat{\Psi}_t(\theta_t, \theta_t) = v(\theta^t, \theta^{t-1}, y^{t-1})$ for any truthful history $(\theta^t, \theta^{t-1}, y^{t-1})$, and hence, recalling (10), that

$$\frac{\partial \widehat{\Psi}_t(\theta_t, \theta_t)}{\partial \theta_t} = \frac{\partial v(\theta^t, \theta^{t-1}, y^{t-1})}{\partial \theta_t} = D^{\Omega}(\theta^t, \theta^t, y^{t-1})$$

10.3 Appendix C: Recursive Optimal Control

The extended Hamiltonian (for interior allocations) of the period-t optimal control problem given (past) skill-history θ^{t-1} and phase variable $(z_{t-1}(\theta^{t-1}), \hat{U}_t(\theta^{t-1}), \hat{V}_t(\theta^{t-1}), \hat{G}_t(\theta^{t-1}))$ is

$$H \equiv [Q(v(\theta) - \beta\hat{U}(\theta)) + \beta\hat{Q}_{t+1}(\theta, z(\theta), \hat{U}(\theta), \hat{V}(\theta), \hat{G}(\theta))]f_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))-$$

$$\mu_t(\theta, \theta^{t-1})\{\frac{z(\theta)}{\theta^2}u_e(C(v(\theta) - \beta\hat{U}(\theta), z(\theta), \theta), z(\theta)/\theta) - \beta\hat{V}(\theta)\}+$$

$$\lambda_t(\theta^{t-1})(\hat{U}_t(\theta^{t-1}) - v(\theta)f_t(\theta \mid .)) + \gamma(\theta^{t-1})(\hat{V}_t(\theta^{t-1}) - v(\theta)\frac{\partial^2 F_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))}{\partial\theta_{t-1}\partial\theta})+$$

$$p_t(\theta^{t-1})\{z(\theta) - C(v(\theta) - \beta\hat{U}(\theta), z(\theta), \theta) + R^{-1}\hat{G}(\theta) - \hat{G}_t(\theta^{t-1})\}f_t(\theta \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))-$$

$$k_t(\theta^{t-1})\hat{G}(\theta)$$

where

$$k_t(\theta^{t-1}) \equiv 0 \text{ for all } 0 < t < T \text{ and } \lambda_0(\theta^0) = \gamma_0(\theta^0) = \hat{G}_1(\theta^0) \equiv 0$$

and, when $t = T$, $\hat{U}(\theta) = \hat{V}(\theta) = 0$.

Here, $\mu_t(\theta, \theta^{t-1})$ is the co-state variable (as a continuous function of θ), $p_t(\theta^{t-1})$ and $k_t(\theta^{T-1})$ are the Kuhn-Tucker multipliers of the period-t fiscal constraint and the inequality constraint on the last-period net liabilities, respectively, and $\lambda_t(\theta^{t-1})$ and $\gamma_t(\theta^{t-1})$ are the Lagrange multipliers of the two period-t constraints on promised utility and marginal promised utility. To facilitate exposition, we will use hereafter the notation $u^t(.)$ for $u(c_t(\theta^t), z_t(\theta^t)/\theta_t)$ and similarly for the derivatives with respect to consumption and effective labour.

To lighten up notation we drop for the rest of this Appendix the subscript t, whenever there is no risk of confusion.

The first-order conditions (FOCs) from the maximisation of $H + \dot{\mu}(\theta, \theta^{t-1})v(\theta)$ with respect to $v(\theta), z(\theta), \hat{G}(\theta)$ are

$$\dot{\mu}(\theta, \theta^{t-1}) =$$

$$-Q'(u^t(.))f_t(\theta \mid .)+[\mu(\theta, \theta^{t-1})\frac{z(\theta)}{\theta^2}u^t_{ec}(.)+p(\theta^{t-1})f_t(\theta \mid .)]\frac{1}{u^t_c(.)}+\lambda(\theta^{t-1})f_t(\theta \mid .)+\gamma(\theta^{t-1})\frac{\partial^2 F_t(\theta \mid .)}{\partial\theta_{t-1}\partial\theta_t}$$

$$-\frac{\partial\hat{Q}_{t+1}(\theta, z(\theta), .)}{\partial z}\beta f_t(\theta \mid .) + \mu(\theta, \theta^{t-1})\frac{1}{\theta^2}\{u^t_e(.) + \frac{z(\theta)}{\theta}u^t_{ee}(.)\} =$$

$$[\mu(\theta, \theta^{t-1})\frac{z(\theta)}{\theta^2}u^t_{ec}(.) + p(\theta^{t-1})f_t(\theta \mid .)]\frac{u^t_e(.)}{\theta u^t_c(.)} + p(\theta^{t-1})f_t(\theta \mid .)$$

$$p(\theta^{t-1})R^{-1} = k(\theta^{t-1}) + \beta p(\theta^t)$$

for all t with

$$p(\theta^T) \equiv 0$$

Moreover, **only for** $0 < t < T$, the FOCs with respect to $\hat{U}(\theta), \hat{V}(\theta)$ are

$$[\lambda(\theta^t) - Q'(u^t(.))]f_t(\theta \mid .) + [\mu(\theta, \theta^{t-1})\frac{z(\theta)}{\theta^2}u^t_{ec}(.) + p(\theta^{t-1})f_t(\theta \mid .)]\frac{1}{u^t_c(.)} = 0$$

$$\gamma(\theta^t)f_t(\theta \mid .) + \mu(\theta, \theta^{t-1}) = 0$$

Note that here we have used the envelope conditions $\partial\hat{Q}_{t+1}(.)/\partial\hat{U}(\theta) = \lambda(\theta^t)$, $\partial\hat{Q}_{t+1}(.)/\partial\hat{V}(\theta) = \gamma(\theta^t)$, and $\partial\hat{Q}_{t+1}(.)/\partial\hat{G}(\theta) = -p(\theta^t)$. (See, for instance, La France and Barney, 1991). Moreover, as we show in Appendix D, $\frac{\partial\hat{Q}_{t+1}(\theta,z(\theta),.)}{\partial z}$ exists.

We also have, for all t, the terminal conditions

$$lim_{\theta\to\underline{\theta}}\mu(\theta,\theta^{t-1}) = lim_{\theta\to\bar{\theta}}\mu(\theta,\theta^{t-1}) = 0 \tag{35}$$

In Appendix F we show that the sufficiency conditions of this optimal control problem are satisfied (and hence the FOCs above characterise the optimal solution), if (a) $C(u,z,\theta)$ is (weakly) convex, (b) the economy is "regular" and $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ is non-decreasing in x (so that Lemma 3 holds), (c) $zu_e(C(u,z,\theta),z/\theta)$ is (weakly) concave,[58] and (d) $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ is (weakly) concave in $(z,\hat{U},\hat{V},\hat{G})$.[59]

Denote hereafter $p \equiv p(\theta^0)$.

By backward iteration of $p(\theta^{t-1})R^{-1} = k(\theta^{t-1}) + \beta p(\theta^t)$ (recalling that $p(\theta^T) = 0$ and $k(\theta^{t-1}) = 0, t = 1,...,T-1$) we thus have for $t = 1,...,T$ that $p(\theta^{t-1}) = k(\theta^{T-1})\frac{R^{T-t+1}\prod_{\tau=t+1}^{T+1}\beta^{T+1-\tau}}{\prod_{\tau=t}^{T-2}\beta^{T-1-\tau}}$. We therefore have that

$$p = k(\theta^{T-1})\frac{R^T\prod_{\tau=2}^{T+1}\beta^{T+1-\tau}}{\prod_{\tau=1}^{T-2}\beta^{T-1-\tau}} > 0$$

and hence for all $t = 1,...,T$

$$p(\theta^{t-1}) = p(\beta R)^{1-t}$$

and for $t = T$

$$\hat{G}(\theta) = 0$$

Note then that the FOC's with respect to $v(\theta), z(\theta)$ can be re-written (after using the remaining FOCs) as follows, for any $t = 1,...,T$:

$$\dot{\mu}(\theta,\theta^{t-1}) = -Q'(u^t(.))f_t(\theta \mid .)+[\mu(\theta,\theta^{t-1})\frac{z(\theta)}{\theta^2}u^t_{ec}(.)+pq(t)f_t(\theta \mid .)]\frac{1}{u^t_c(.)}-\frac{\mu(\theta^{t-1})}{f_{t-1}(\theta_{t-1} \mid .)}\frac{\partial^2 F_t(\theta \mid .)}{\partial\theta_{t-1}\partial\theta}$$

$$+ Q'(u^{t-1}(.))f_t(\theta \mid .) - [\mu(\theta^{t-1})\frac{z_{t-1}(\theta^{t-1})}{\theta^2_{t-1}}u^{t-1}_{ec}(.) + pq(t-1)f_{t-1}(\theta \mid .)]\frac{1}{u^{t-1}_c(.)}\frac{f_t(\theta \mid .)}{f_{t-1}(\theta \mid .)} \tag{36}$$

$$-\frac{\partial\hat{Q}_{t+1}(\theta,.,z(\theta),.)}{\partial z}\beta f_t(\theta \mid .) + \mu(\theta,\theta^{t-1})\frac{1}{\theta^2}\{u^t_e(.) + \frac{z(\theta)}{\theta}u^t_{ee}(.)\} =$$

$$[\mu(\theta,\theta^{t-1})\frac{z(\theta)}{\theta^2}u^t_{ec}(.) + pq(t)f_t(\theta \mid .)]\frac{u^t_e(.)}{\theta u^t_c(.)} + pq(t)f_t(\theta \mid .) \tag{37}$$

[58] Note that this is true when utility is additively-separable with constant Frisch elasticity $1/(\eta-1) = [u_{ee}e/u_e]^{-1}$. Note also that (a) and (c) are common in static screening problems.

[59] Assumption (d) trivially holds for $t = T$. The assumptions (a)-(c) would also be sufficient for the concavity of $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ in $(z,\hat{U},\hat{V},\hat{G})$, $0 < t < T$, if there was no learning-by-doing (this follows standard steps from dynamic programming). Unfortunately, we have not been able to push assumption (d) further back into the fundamentals for $1 \leq t \leq T-1$.

where

$$\mu(\theta^0) \equiv q(0) \equiv Q'(u^0(.)) \equiv 0$$

and

$$q(t) \equiv (R\beta)^{1-t}$$

FOC (37) is (17) in the main text.

The FOC (36), alongside the terminal conditions (35), imply that

$$\int_{\underline{\theta}}^{\overline{\theta}} \frac{pq(t)}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)} \Gamma(x, \underline{\theta}, \theta^{t-1}) dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1})) =$$

$$\int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)) \Gamma(x, \underline{\theta}, \theta^{t-1}) dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

$$- \int_{\underline{\theta}}^{\overline{\theta}} [Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})) - \frac{pq(t-1)}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}] \Gamma(x, \underline{\theta}, \theta^{t-1}) dF_t(x \mid .)$$

$$+ \int_{\underline{\theta}}^{\overline{\theta}} \frac{\mu(\theta^{t-1})}{f_{t-1}(\theta^{t-1} \mid .)} \frac{\partial f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))/\partial \theta_{t-1}}{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} \Gamma(x, \underline{\theta}, \theta^{t-1}) dF_t(x \mid .)$$

$$+ \int_{\underline{\theta}}^{\overline{\theta}} \mu(\theta^{t-1}) \frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2} \frac{u_{ec}(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}{f_{t-1}(\theta_{t-1} \mid .) u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} \Gamma(x, \underline{\theta}, \theta^{t-1}) dF_t(x \mid .) \quad (38)$$

and

$$\mu(\theta^t) =$$

$$\int_{\theta_t}^{\overline{\theta}} [Q'(u(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)) - \frac{pq(t)}{u_c(c_t(x,\theta^{t-1}), z_t(x,\theta^{t-1})/x)}] \Gamma(x, \theta^t) dF_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))$$

$$- \int_{\theta_t}^{\overline{\theta}} [Q'(u(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})) - \frac{pq(t-1)}{u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}] \Gamma(x, \theta^t) dF_t(x \mid .)$$

$$+ \int_{\theta_t}^{\overline{\theta}} \frac{\mu(\theta^{t-1})}{f_{t-1}(\theta^{t-1} \mid .)} \frac{\partial f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))/\partial \theta_{t-1}}{f_t(x \mid \theta_{t-1}, z_{t-1}(\theta^{t-1}))} \Gamma(x, \theta^t) dF_t(x \mid .)$$

$$+ \int_{\theta_t}^{\overline{\theta}} \mu(\theta^{t-1}) \frac{z_{t-1}(\theta^{t-1})}{\theta_{t-1}^2} \frac{u_{ec}(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})}{f_{t-1}(\theta_{t-1} \mid .) u_c(c_{t-1}(\theta^{t-1}), z_{t-1}(\theta^{t-1})/\theta_{t-1})} \Gamma(x, \theta^t) dF_t(x \mid .) \quad (39)$$

where

$$\Gamma(x,\theta^t) \equiv exp[-\int_{\theta_t}^{x} \frac{z_t(\vartheta,\theta^{t-1})}{\vartheta^2}\frac{u_{ec}(c_t(\vartheta,\theta^{t-1}),z_t(\vartheta,\theta^{t-1})/\vartheta)}{u_c(c_t(\vartheta,\theta^{t-1}),z_t(\vartheta,\theta^{t-1})/\vartheta)}d\vartheta]$$

Given the allocation, condition (38) for $t=1$ determines the "shadow price of public funds", p $p(\theta^0)$:

$$p = \frac{\int_{\underline{\theta}}^{\overline{\theta}} Q'(u(c_1(x),z_1(x)/x))exp[-\int_{\underline{\theta}}^{x}\frac{z_1(\vartheta)}{\vartheta^2}\frac{u_{ec}(c_1(\vartheta),z_1(\vartheta)/\vartheta)d\vartheta}{u_c(c_1(\vartheta),z_1(\vartheta)/\vartheta)d\vartheta}]f_1(x)dx}{\int_{\underline{\theta}}^{\overline{\theta}}[u_c(c_1(x),z_1(x)/x)]^{-1}exp[-\int_{\underline{\theta}}^{x}\frac{z(\vartheta)}{\vartheta^2}\frac{u_{ec}(c_1(\vartheta),z_1(\vartheta)/\vartheta)d\vartheta}{u_c(c_1(\vartheta),z_1(\vartheta)/\vartheta)d\vartheta}]f_1(x)dx} > 0 \tag{40}$$

which implies that $\widehat{G}_{T+1}(\theta^T)=0$ and hence the intertemporal resource constraint is binding.

Rearranging (38) and dividing both sides with $\int_{\underline{\theta}}^{\overline{\theta}}\Gamma(x,\underline{\theta},\theta^{t-1})f_t(x\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))dx$ gives (18) in the main text. Dividing both sides of (40) with $\int_{\underline{\theta}}^{\overline{\theta}}\Gamma(x,\underline{\theta},\theta^0)f_1(x)dx$ gives (22) in the main text. Using (18) and (40) in (39) gives (19) in the main text.

10.4 Appendix D: Differentiability of Value Function

We prove here that $\hat{Q}_t(\theta_{t-1},z,\hat{U},\hat{V},\hat{G})$ is twice continuously differentiable in $z,\hat{U},\hat{V},\hat{G}$, for any t,θ_{t-1}. The proof proceeds in two steps.

First Step: We first note, given AU2 and AF4, and the uniqueness assumption of the optimal path, that standard steps from optimal control theory imply that $\hat{Q}_t(\theta_{t-1},z,\hat{U},\hat{V},\hat{G})$ is twice continuously differentiable in $z,\hat{U},\hat{V},\hat{G}$, if $\hat{Q}_{t+1}(\theta_t,z',\hat{U}',\hat{V}',\hat{G}')$ is twice continuously differentiable in $z',\hat{U}',\hat{V}',\hat{G}'$; see, for instance, La France and Barney (1991).

Second Step: The result follows from an inductive proof that uses $\hat{Q}_{T+1}(\theta_T,z,\hat{U},\hat{V},\hat{G})\equiv 0$.

10.5 Appendix E: The Learning-by-doing Component of Labour Wedge

Note: setting $\mu_t(\theta^t)=\gamma_t(\theta^t)=0$ in what follows gives the period-t full-information labour wedge.

Recall that the labour wedge associated with skill-history θ^t in any period t is given by

$$\zeta_t^z(\theta^t) \equiv 1 + \frac{u_e(c_t(\theta^t),z_t(\theta^t)/\theta_t)}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)\theta_t} + \frac{\hat{\delta}(\theta^t,z^t(\theta^t))}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)}$$

Recalling (37) (or (17) in the main text), the above becomes

$$\zeta_t^z(\theta^t) \equiv \frac{\hat{\delta}(\theta^t,z^t(\theta^t))}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)} - \frac{\partial\hat{Q}_{t+1}(\theta_t,z_t(\theta_t),.)}{pq(t)\partial z}\beta$$

$$+\frac{\mu_t(\theta^t)/pq(t)}{f_t(\theta_t\mid\theta_{t-1},z_{t-1}(\theta^{t-1}))\theta_t^2}u_e(c_t(\theta^t),.)\left[1+\frac{z_t(\theta^t)}{\theta_t}[\frac{u_{ee}(c_t(\theta^t),.)}{u_e(c_t(\theta^t),.)}-\frac{u_{ec}(c_t(\theta^t),.)}{u_c(c_t(\theta^t),.)}]\right]$$

Let

$$\widehat{u}(\theta^{t+1}) \equiv u(c_{t+1}(\theta_{t+1},\theta^t),z_{t+1}(\theta_{t+1},\theta^t)/\theta_{t+1})$$

Note now that, after using the envelope theorem vis-a-vis the planner (see, for instance, La France and Barney, 1991), the definition of $\hat{Q}_{t+2}(\theta_{t+1}, z_{t+1}(\theta^{t+1}), .)$ and the FOCs of the planner's period$-t$ problem (see Appendix C), that

$$\frac{\partial \hat{Q}_{t+1}(\theta_t, z_t(\theta^t), .)}{\partial z_t(\theta^t)} \equiv$$

$$\frac{\partial}{\partial z_t(\theta^t)} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{\tau=t+1}^{T} \beta^{\tau-t-1} Q(\widehat{u}(\theta^\tau)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$-\lambda_t(\theta^t) \int_{\underline{\theta}}^{\overline{\theta}} [\widehat{u}(\theta^{t+1}) + \beta \widehat{U}_{t+2}(\theta^{t+1})] \frac{\partial}{\partial z_t(\theta^t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$-\gamma_t(\theta^t) \int_{\underline{\theta}}^{\overline{\theta}} [\widehat{u}(\theta^{t+1}) + \beta \widehat{U}_{t+2}(\theta^{t+1})] \frac{\partial^2}{\partial z_t(\theta^t) \partial \theta_t} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$+p_t(\theta^t) \int_{\underline{\theta}}^{\overline{\theta}} [z_{t+1}(\theta^{t+1}) - c_{t+1}(\theta^{t+1}) + \frac{\widehat{G}_{t+2}(\theta^{t+1})}{R}] \frac{\partial}{\partial z_t(\theta^t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$= \frac{\partial}{\partial z_t(\theta^t)} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{\tau=t+1}^{T} \beta^{\tau-t-1} Q(\widehat{u}(\theta^\tau)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$+\frac{\mu_t(\theta^t)}{f_t(\theta_t|\theta_{t-1}, z_{t-1}(\theta^{t-1}))} \frac{z_t(\theta^t) u_{ec}(c_t(\theta^t), .)}{\theta_t^2 u_c(c_t(\theta^t), .)} \int_{\underline{\theta}}^{\overline{\theta}} [\widehat{u}(\theta^{t+1}) + \beta \widehat{U}_{t+2}(\theta^{t+1})] \frac{\partial}{\partial z_t(\theta^t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$+[\frac{pq(t)}{u_c(c_t(\theta^t), .)} - Q'(\widehat{u}(\theta^t))] \int_{\underline{\theta}}^{\overline{\theta}} [\widehat{u}(\theta^{t+1}) + \beta \widehat{U}_{t+2}(\theta^{t+1})] \frac{\partial}{\partial z_t(\theta^t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$+\frac{\mu_t(\theta^t)}{f_t(\theta_t|\theta_{t-1}, z_{t-1}(\theta^{t-1}))} \int_{\underline{\theta}}^{\overline{\theta}} [\widehat{u}(\theta^{t+1}) + \beta \widehat{U}_{t+2}(\theta^{t+1})] \frac{\partial^2}{\partial z_t(\theta^t) \partial \theta_t} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$+pq(t+1) \int_{\underline{\theta}}^{\overline{\theta}} [z_{t+1}(\theta^{t+1}) - c_{t+1}(\theta^{t+1}) + \frac{\widehat{G}_{t+2}(\theta^{t+1})}{R}] \frac{\partial}{\partial z_t(\theta^t)} dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$= \frac{\partial}{\partial z_t(\theta^t)} \int_{\underline{\theta}}^{\overline{\theta}} \sum_{\tau=t+1}^{T} \beta^{\tau-t-1} Q(\widehat{u}(\theta^\tau)) dF_T(\theta_T \mid \theta_{T-1}, z_{T-1}(\theta^{T-1})) ... dF_{t+1}(\theta_{t+1} \mid \theta_t, z_t(\theta^t))$$

$$-Q'(\widehat{u}(\theta^t))\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+\frac{\mu_t(\theta^t)}{f_t(\theta_t|\theta_{t-1},z_{t-1}(\theta^{t-1}))}\frac{z_t(\theta^t)u_{ec}(c_t(\theta^t),.)}{\theta_t^2u_c(c_t(\theta^t),.)}\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+pq(t)\times$$

$$\{\frac{\hat{\delta}(\theta_t,z_t(\theta^t))}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)\beta}$$

$$+\frac{\mu_t(\theta^t)/pq(t)}{f_t(\theta_t|\theta_{t-1},z_{t-1}(\theta^{t-1}))}\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial^2}{\partial z_t(\theta^t)\partial\theta_t}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+[R\beta]^{-1}\int_{\underline{\theta}}^{\overline{\theta}}[z_{t+1}(\theta^{t+1})-c_{t+1}(\theta^{t+1})+\frac{\widehat{G}_{t+2}(\theta^{t+1})}{R}]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))\}$$

Therefore, we have that the learning-by-doing component of the period-t labour wedge is

$$\widehat{\zeta}_t^z(\theta^t)\equiv$$

$$\frac{\hat{\delta}(\theta_t,z_t(\theta^t))}{u_c(c_t(\theta^t),z_t(\theta^t)/\theta_t)}-\frac{\partial\hat{Q}_{t+1}(\theta_t,z_t(\theta^t),.)}{pq(t)\partial z}\beta$$

$$=-\beta\times$$

$$\{\frac{\partial}{\partial z_t(\theta^t)}\int_{\underline{\theta}}^{\overline{\theta}}\sum_{\tau=t+1}^{T}\beta^{\tau-t-1}\frac{Q(\widehat{u}(\theta^\tau))}{pq(t)}dF_T(\theta_T\mid\theta_{T-1},z_{T-1}(\theta^{T-1}))...dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$-\frac{Q'(\widehat{u}(\theta^t))}{pq(t)}\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+\frac{\mu_t(\theta^t)/pq(t)}{f_t(\theta_t|\theta_{t-1},z_{t-1}(\theta^{t-1}))}\frac{z_t(\theta^t)u_{ec}(c_t(\theta^t),.)}{\theta_t^2u_c(c_t(\theta^t),.)}\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+\frac{\mu_t(\theta^t)/pq(t)}{f_t(\theta_t|\theta_{t-1},z_{t-1}(\theta^{t-1}))}\int_{\underline{\theta}}^{\overline{\theta}}[\widehat{u}(\theta^{t+1})+\beta\widehat{U}_{t+2}(\theta^{t+1})]\frac{\partial^2}{\partial z_t(\theta^t)\partial\theta_t}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))$$

$$+[R\beta]^{-1}\int_{\underline{\theta}}^{\overline{\theta}}[z_{t+1}(\theta^{t+1})-c_{t+1}(\theta^{t+1})+\frac{\widehat{G}_{t+2}(\theta^{t+1})}{R}]\frac{\partial}{\partial z_t(\theta^t)}dF_{t+1}(\theta_{t+1}\mid\theta_t,z_t(\theta^t))\}$$

For the two-period model one needs to set $T-1=t=1$ above, and that $\hat{G}_3(\theta^2)/R_3=0$ and $\hat{U}_3(\theta^2)=0$. For the case of non-Utilitarian planner and additively-separable utility, we thus have (recalling that $q(1)=1$)

$$\widehat{\zeta}_1^z(\theta_1)\equiv\frac{\hat{\delta}(\theta_1,z_1(\theta_1))}{u_c(c_1(\theta_1),z_1(\theta_1)/\theta_1)}-\frac{\partial\hat{Q}_2(\theta_1,z_1(\theta_1),.)}{p\partial z}\beta=$$

$$=-(\beta/p)\times\{\int_{\underline{\theta}}^{\overline{\theta}}Q(u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2))\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))$$

$$-Q'(u(c_1(\theta_1)),.)\int_{\underline{\theta}}^{\overline{\theta}}u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2)\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))\}$$

$$-\beta\times\{\frac{\mu_1(\theta_1)/p}{f_1(\theta_1)}\int_{\underline{\theta}}^{\overline{\theta}}u(c_2(\theta_2,\theta_1),z_2(\theta_2,\theta_1)/\theta_2)\frac{\partial^2}{\partial z_1(\theta_1)\partial\theta_1}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))$$

$$+[R\beta]^{-1}\int_{\underline{\theta}}^{\overline{\theta}}[z_2(\theta_2,\theta_1)-c_2(\theta_2,\theta_1)]\frac{\partial}{\partial z_1(\theta_1)}dF_2(\theta_2\mid\theta_1,z_1(\theta_1))\}$$

while for the case of a Utilitarian planner the first two terms cancel each other out.

10.6 Appendix F: Second Order Sufficiency Conditions for Optimal Control Problem

We prove next that the second order sufficiency conditions for the Relaxed Program are satisfied (namely,[60] that the Hamiltonian H in Appendix C is concave in the phase and control variables), if (a) $C(u,z,\theta)$ is (weakly) convex in u and z, (b) the economy is "regular" and $\frac{\partial f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))/\partial\theta_{t-1}}{f_t(x|\theta_{t-1},z_{t-1}(\theta^{t-1}))}$ is non-decreasing in x, (c) $-zu_e(C(u,z,\theta),z/\theta)$ is (weakly) convex in u and z, and (d) $\hat{Q}_{t+1}(\theta,z,\hat{U},\hat{V},\hat{G})$ is (weakly) concave in $(z,\hat{U},\hat{V},\hat{G})$.

First Step: We first note that the FOCs in Appendix C imply that $p(\theta^{t-1})\geq 0$ for any t, and that under (b) we have that $\mu(\theta^t)\leq 0$ (see Lemma 3).

[60] See, for instance, Kamien and Schwartz, 1991, pp. 133-135, 228-229, and references therein.

The proof then proceeds in three steps.

Second Step: We show, under (a) and (c), that $C(v(\theta)-\beta\hat{U}(\theta),z(\theta),\theta)$ and $-z(\theta)u_e(C(v(\theta)-\beta\hat{U}(\theta),z(\theta),\theta),z(\theta)/\theta)$ are (weakly) convex functions.

Consider first a Hessian of the form

$$\begin{bmatrix} A_1 & A_2 & -A_1\beta \\ B_1 & B_2 & -B_1\beta \\ -A_1\beta & -A_2\beta & A_1\beta^2 \end{bmatrix}$$

with $A_1 \geq 0$, $B_2 \geq 0$ and $A_1B_2 \geq A_2B_1$. Note, by using the expansion with respect to the last row, that the determinant of this Hessian is zero. Note also that the determinant of the principal minor that emerges from deleting the second row and column is also zero. Thus, all principal minors of the above Hessian are non-negative.

Note now that both the above functions have a Hessian of the above form, with $v(\theta), z(\theta), \hat{U}(\theta)$ being the first, second and third arguments respectively, when (a) and (c) hold.

Third Step: We prove, under (a)-(d), that for any t, the following Hamiltonian is concave in $(v(\theta),z(\theta),\hat{U}(\theta),\hat{V}(\theta),\hat{G}(\theta))$:

$$[Q(v(\theta)-\beta\hat{U}(\theta))+\beta\hat{Q}_{t+1}(\theta,z(\theta),\hat{U}(\theta),\hat{V}(\theta),\hat{G}(\theta))]f_t(\theta \mid \theta_{t-1},z_{t-1}(\theta^{t-1}))-$$

$$\mu(\theta^t)\{\frac{z(\theta)}{\theta^2}u_e(C(v(\theta)-\beta\hat{U}(\theta),z(\theta),\theta),z(\theta)/\theta)-\beta\hat{V}(\theta)\}-$$

$$\lambda(\theta^{t-1})v(\theta)f_t(\theta \mid \theta_{t-1},z_{t-1}(\theta^{t-1}))-\gamma(\theta^{t-1})v(\theta)\frac{\partial f_t(\theta_t \mid \theta_{t-1},z_{t-1}(\theta^{t-1}))}{\partial\theta_{t-1}})+$$

$$p(\theta^{t-1})\{z(\theta)-C(v(\theta)-\beta\hat{U}(\theta),z(\theta),\theta)+R^{-1}\hat{G}(\theta)\}f_t(\theta \mid \theta_{t-1},z_{t-1}(\theta^{t-1}))$$

The proof follows directly after recalling that $\beta \geq 0$, $Q(.)$ is concave and increasing, and using the First and Second Steps.

10.7 Appendix G

Let Φ denote the cdf of the standard normal. Let $F^{LOG}(x|\Xi) = \Phi(\frac{logx-log\Xi+(\sigma^2/2)}{\sigma\sqrt{2}})$ for any $0 < x \leq \Xi e^{-\frac{3}{2}\sigma^2} \equiv \theta^m$ and $F^{PAR}(x|\Xi) = 1-(\frac{\Xi e^{-\frac{3}{2}\sigma^2}}{x})^\alpha$ for any $x \geq \Xi e^{-\frac{3}{2}\sigma^2}$, $\alpha > 1$. Let also $\overline{F}(x|\Xi) = F^{LOG}(\theta^m|\Xi)+(1-F^{LOG}(\theta^m|\Xi))F^{PAR}(x|\Xi)$ for any $x \geq \theta^m$.

Let $F(x|\theta,z) \equiv F^{LOG}(x|\Xi(\theta,z))$ for $0 < x \leq \theta^m$ and $F(x|\theta,z) \equiv \overline{F}(x|\Xi(\theta,z))$ for $x \geq \theta^m$. Note that $F(x|\Xi)$ is a cdf that is constructed by fitting a Lognormal distribution with mean $\Xi = \Xi(\theta,z) > 0$ with a Pareto tail at the modal skill level $x = \theta^m$.

Let $\Xi(\theta,z) \equiv \xi+\rho^{-1}\theta^\rho+log(\theta+\frac{2}{1+e^{-(z/\theta)}})$, $\xi \geq 0$,$\rho \in (0,1)$, $\rho \to 0$. We have that $\Xi(\theta,z)$ is strictly increasing, and it is also concave in θ. Let θ_L be the skill that ensures $\frac{\partial^2\Xi}{\partial\theta\partial z} \leq 0$ for any $\theta \geq \theta_L$. For our case here $\theta_L < 1/2$. We note also that $\theta \geq \theta_L$ ensures that $\frac{\partial^2\Xi}{\partial\theta\partial z}$ is bounded.

Assume that $\theta \geq \theta_L$.

We have that $\frac{\partial}{\partial \Xi} F^{LOG}(x|\Xi) \leq 0$ and $\frac{\partial^2}{\partial \Xi^2} F^{LOG}(x|\Xi) > 0$ (note here that $logx < log\Xi - (\sigma^2/2)$ and so $erf''(\frac{logx - log\Xi + (\sigma^2/2)}{\sigma\sqrt{2}}) > 0$).

Moreover, $\frac{\partial}{\partial \Xi} F^{PAR}(x|\Xi) = -(\frac{e^{-\frac{3}{2}\sigma^2}}{x})^{\alpha} \alpha \Xi^{\alpha-1} < 0$, $\frac{\partial^2}{\partial \Xi^2} F^{PAR}(x|\Xi) = \frac{a-1}{\Xi} \frac{\partial}{\partial \Xi} F^{PAR}(x|\Xi) < 0$ and $1 - F^{PAR}(x|\Xi) = -\frac{\Xi}{\alpha} \frac{\partial}{\partial \Xi} F^{PAR}(x|\Xi)$ and hence

$$\frac{\partial^2 F^{PAR}(x|\Xi(\theta,z))}{\partial\theta\partial z} = \frac{\partial^2 F^{PAR}}{\partial \Xi^2} \frac{\partial \Xi}{\partial \theta} \frac{\partial \Xi}{\partial z} + \frac{\partial F^{PAR}}{\partial \Xi} \frac{\partial^2 \Xi}{\partial \theta \partial z}$$

$$= \frac{\partial F^{PAR}}{\partial \Xi} [\frac{a-1}{\Xi} \frac{\partial \Xi}{\partial \theta} \frac{\partial \Xi}{\partial z} + \frac{\partial^2 \Xi}{\partial \theta \partial z}]$$

Clearly, the latter is non-positive for α sufficiently high.

In addition,

$$\frac{\partial \overline{F}(x|\Xi(\theta,z))}{\partial \theta} =$$

$$\frac{\partial}{\partial \theta} F^{LOG}(\theta^m|\Xi(\theta,z))[1 - F^{PAR}(x|\Xi(\theta,z))] + (1 - F^{LOG}(\theta^m|\Xi(\theta,z))) \frac{\partial}{\partial \theta} F^{PAR}(x|\Xi(\theta,z))\}$$

and

$$\frac{\partial^2 \overline{F}(x|\Xi(\theta,z))}{\partial\theta\partial z} = -\frac{\partial}{\partial \theta} F^{LOG}(\theta^m|\Xi(\theta,z)) \frac{\partial}{\partial z} F^{PAR}(x|\Xi(\theta,z))$$

$$-\frac{\partial}{\partial z} F^{LOG}(\theta^m|\Xi(\theta,z)) \frac{\partial}{\partial \theta} F^{PAR}(x|\Xi(\theta,z))$$

$$+\frac{\partial^2}{\partial\theta\partial z} F^{LOG}(\theta^m|\Xi(\theta,z))[1 - F^{PAR}(x|\Xi(\theta,z))] + (1 - F^{LOG}(\theta^m|\Xi(\theta,z))) \frac{\partial^2}{\partial\theta\partial z} F^{PAR}(x|\Xi(\theta,z))\}$$

Note then that for any $0 < x \leq \theta^m$:

$$\frac{\partial^2 F}{\partial\theta\partial z} = \frac{\partial^2}{\partial\theta\partial z} F^{LOG}(\theta^m|\Xi(\theta,z)) = \frac{\partial^2 F^{LOG}}{\partial \Xi^2} \frac{\partial \Xi}{\partial \theta} \frac{\partial \Xi}{\partial z} + \frac{\partial F^{LOG}}{\partial \Xi} \frac{\partial^2 \Xi}{\partial \theta \partial z} > 0$$

for any θ, z.

For any $x \geq \theta^m$:

$$\frac{\partial^2 F}{\partial\theta\partial z} = \frac{\partial^2 \overline{F}(x|\Xi(\theta,z))}{\partial\theta\partial z} = \frac{\partial}{\partial \Xi} F^{PAR}(x|\Xi(\theta,z))$$

$$\times\{[-2] \frac{\partial \Xi}{\partial \theta} \frac{\partial \Xi}{\partial z} \frac{\partial}{\partial \Xi} F^{LOG}(\theta^m|\Xi(\theta,z)) + [-\frac{\Xi}{\alpha}] \frac{\partial^2}{\partial\theta\partial z} F^{LOG}(\theta^m|\Xi(\theta,z))$$

$$+[\frac{a-1}{\Xi} \frac{\partial \Xi}{\partial \theta} \frac{\partial \Xi}{\partial z} + \frac{\partial^2 \Xi}{\partial\theta\partial z}][1 - F^{LOG}(\theta^m|\Xi(\theta,z))]\}$$

Note that, in the curly brackets, the first term is non-negative, the second term is negative and the last term is positive for sufficiently high α. Clearly, then, there is a sufficiently large α such that the term in the curly brackets is positive and hence $\frac{\partial^2 F}{\partial\theta\partial z} < 0$.