NBER WORKING PAPER SERIES

LEARNING AND (OR) DOING:
HUMAN CAPITAL INVESTMENTS AND OPTIMAL TAXATION

Stefanie Stantcheva

Working Paper 21381
http://www.nber.org/papers/w21381

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2015

I want to thank James Poterba, Ivan Werning, and Emmanuel Saez for insightful comments. I also thank participants at the MIT Public Finance Lunch and seminar participants at Berkeley, Michigan, Princeton, Stanford, UPenn, Yale, Wharton, and Santa Barbara for their useful feedback and comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Learning and (or) Doing: Human Capital Investments and Optimal Taxation
Stefanie Stantcheva
NBER Working Paper No. 21381
July 2015
JEL No. H21,H23,H24,H53,J24

ABSTRACT

This paper considers a dynamic taxation problem when agents can allocate their time between working and investing in their human capital. Time investment in human capital, or "training," increases the wage and can interact with an agent's intrinsic, exogenous, and stochastic earnings ability. It also interacts with both current and future labor supply and there can be either "learning-and-doing" (when labor and training are complements, like for on-the-job training) or "learning-or-doing" (when labor and training are substitutes, like for college). Agents' abilities and labor supply are private information to them, which leads to a dynamic mechanism design problem with incentive compatibility constraints. At the optimum, the subsidy on training time is set so as to balance the total labor supply effect of the subsidy and its distributional consequences. In a one-period version of the model, particularly simple relations arise at the optimum between the labor wedge and the training wedge that can also be used to test for the Pareto efficiency of existing tax and subsidy systems. In the limit case of learning-by-doing (when training is a direct by-product of labor) or in the case in which agents who are more able at work are also more able at training, there are important modifications to the labor wedge.

Stefanie Stantcheva
Department of Economics
Littauer Center 232
Harvard University
Cambridge, MA 02138
and NBER
sstantcheva@fas.harvard.edu

1 Introduction

Most people spend a lot of time acquiring human capital throughout their life. These time investments can take many different forms: formal college, continuing education, online degrees, on-the-job training, or vocational training.[1] Mirroring the variety of possible training types, there is also a plethora of different policies throughout the world that aim to encourage people to invest time in improving their skills. They include subsidies, grants or loans for time spent in college, firm incentives for job training, regulations regarding on-the-job training, or mixed arrangements such as apprenticeships.[2] At the same time, other important and far-reaching policies, such as income taxation or social insurance, could have an ambiguous effect on time investments in human capital. Indeed, they not only affect the returns to time investments, but also the allocation of time between training and working.

Standard models of social insurance and taxation typically focus only on time spent working, abstracting from the fact that time can also be allocated to improving one's productivity.[3] Therefore, this paper studies the following issues: How should a comprehensive social insurance and tax system take into account time spent acquiring human capital, called "training"? What makes time investments in human capital different from other investments, such as investments in physical capital or material investments in human capital? What parameters do we need to know to determine the optimal policies regarding training time and the optimal tax system in the presence of training?

The first goal of this paper is to provide a flexible life cycle framework that can capture the unique characteristics of training time in order to answer these questions. The second goal is to solve for the optimal taxation and training policies over the life cycle in this general model, and to point out how they depend on some important parameters.

The peculiarity of time investments in human capital is that they are intrinsically linked to a given agent: their returns depend on an agent's type and their costs depend on the agent's labor supply. Accordingly, the wage of agents in the model depends on their age, their endogenous human capital stock, acquired through training, and on an exogenous, stochastic, and persistent earnings ability. The stochastic ability causes earnings fluctuations against which agents try to insure themselves by adjusting their labor supply, training choices, and risk-free savings. Ability and training can interact in any arbitrary way in the wage function and the wage function can change with age, allowing for

[1] Heckman (1976a) indirectly infers that 30% of time in early working years is spent training. 40% of adults participate in formal and/or non-formal education in a given year in OECD countries, and receive on average 988 hours of instruction in non-formal education during life, out of which 715 are job-related (OECD, 2013).

[2] The most notable example of which is probably the German apprenticeship system that combines formal training with hands-on experience.

[3] See Mirrlees (1971), Saez (2001), Golosov *et al.* (2006), Albanesi and Sleet (2006), Farhi and Werning (2013). The literature review will present some recent papers that have tried to relax this assumption.

different impacts of training at different points in life.

In the model, agents can invest time in training in every year of their working lives. They incur a disutility cost from working and training, and the disutility specification allows for any age-dependent interaction between work and training effort. To characterize this interaction, I introduce two concepts "Learning-or-doing" and "Learning-and-doing." The former designates the case in which labor and training time are subsitutes: time spent working cannot be spent training. In the limit, this corresponds to the standard opportunity cost of time model as in Ben-Porath (1976). "Learning-and-doing" is the case in which labor and training are complements. For instance, training could reduce the marginal disutility of working by making agents more aware of how to optimize their work effort. The limit case here is the canonical "learning-by-doing" model in which training is a direct by-product of labor.

The government wants to minimize the resource cost of providing any level of redistribution and insurance, but faces asymmetric information about agents' abilities and labor efforts. Because of this informational asymmetry, I solve a dynamic mechanism design problem with incentive compatibility constraints.

In order to build the intuition for the results, I first present a simple one-period version of the dynamic model, which features a single realization of uncertainty and a one-shot investment in human capital. I introduce the notion of a net wedge, as a measure of the net incentive that is provided for training, after one has compensated the agent for the existence of other distortions, such as the labor tax (and, in the dynamic model, the savings tax). A positive net wedge means that the government wants to encourage training acquisition on net. The net wedge has three effects: The first is to increase the wage by stimulating investments in training, which in turn increases labor supply. At the same time, the presence of learning-or-doing or learning-and-doing directly affects labor supply positively or negatively. The third effect is the differential impact of training on people with different abilities that can increase or reduce inequality, depending on the wage function. The optimal subsidy balances these three effects.

At the optimum, whether training needs to be encouraged or discouraged on net depends on two crucial parameters: the Hicksian coefficient of complementarity between training and ability in the wage function and the Hicksian coefficient of complementarity between training and labor in the disutility function. The net wedge decreases the more complementary human capital and ability are in the wage function, decreases the more learning-or-doing there is, and increases the more learning-and-doing there is. I derive a simple relation between the labor wedge and the training wedge at the optimum. In some special cases, this relation has to remain constant over life and can be used to easily test for the Pareto efficiency of any tax and subsidy system.

In the dynamic model, the sign of the net wedge also depends on the interaction of contemporaneous training with future labor supply, and the sign of this effect is exactly the opposite of the one on the interaction with current labor supply. The amplitude of the net wedge is determined by an insurance motive and the persistence of the ability shocks – factors which usually affect the dynamic labor tax in models without human capital.[4]

The labor wedge is not fundamentally affected by the presence of training, except in two cases. The first is the limit case of learning-by-doing, in which training is a direct by product of (unobservable) labor supply. In this case, the labor wedge indirectly provides incentives for human capital acquisition as well as for labor supply. The second case is when more able agents are not only more productive at working, but also at training.

Related Literature: This paper contributes to the long-standing optimal taxation literature as developed by Mirrlees (1971) and Saez (2001), and especially to the dynamic taxation literature, as in Kocherlakota (2005), Albanesi and Sleet (2006), Golosov *et al.* (2006), Kapicka (2013b), Werning (2007a,b,c), Battaglini and Coate (2008a,b). Recent taxation papers seek to expand and improve the canonical optimal taxation framework (see Rothschild and Scheuer (2014a,b), Scheuer (2013a,b), and Shourideh (2012)). The life cycle results directly speak to age-dependent taxation as in Weinzierl (2011). Particularly relevant are the dynamic taxation models of Golosov, Tsyvinski and Troshkin (2013), Shourideh and Troshkin (2012), Golosov, Tsyvinski and Werquin (2014), Werquin (2014), and, most importantly, the life cycle taxation model in Farhi and Werning (2013).

This paper is also related to the large literature on human capital acquisition, starting with Becker (1964), Ben-Porath (1967), and Heckman (1976a,b). The findings in Heckman, Todd and Lochner (2005) can justify the assumption of a disutility cost of human capital investment, above and beyond a pure opportunity cost of time.

Some papers have considered taxation in the presence of human capital investments, but have focused exclusively on monetary human capital expenses in static models (see Bovenberg and Jacobs (2005), Maldonado, 2007, Bovenberg and Jacobs, 2011, DaCosta and Maestri, 2007.) Gelber and Weinzierl (2012) consider parental investments in children's future abilities.

Dynamic papers that examine the impact of taxation on human capital bear important differences to the current paper. Bohacek and Kapicka (2008) or Kapicka (2013) focus on heterogeneity, without uncertainty in human capital returns. Findeisen and Sachs (2013) focus on a one-shot monetary investment during "college," different from the life cycle time investments with progressive realization of uncertainty in this paper. Complementary papers are those by Kapicka and Neira (2014) with

[4]See for instance, Farhi and Werning (2013).

unobservable human capital, and Krueger and Ludwig (2013) with an overlapping generations model and Ramsey taxation.

The closest paper is Stantcheva (2014), which considers monetary investments in human capital. Time investments, as already emphasized earlier in the introduction, are very different under asymmetric information.[5] Their costs interact with the unobservable labor supply, which affects the incentive provision mechanism in a way that monetary expenses on human capital or physical capital investments do not. Accordingly, the results are substantially different and introduce new concepts such as the learning-or-doing or learning-and-doing, which turn out to be the major drivers of the optimal training policies.

The rest of the paper is organized as follows. Section 2 presents a model of work and training over the life cycle. Section 3 solves for the optimal policies in a simple one-period version of the model which helps build the intuitions. Section 4 sets up the dynamic planning problem and explains the solution method. Section 5 presents the results on the optimal dynamic training policies. Section 6 digs deeper into how the optimal labor wedge depends on the specification of the model with training. Section 7 numerically illustrates some of the results. Section 8 concludes.

2 A model of work and training

In this Section, I present a dynamic model with work and training. I then solve for the first-best allocation if there is no asymmetric information between the Planner and agents.

2.1 Environment

The economy consists of agents who live for T years, during which they work and spend time investing in human capital. Agents who work $l_t \geq 0$ hours in period t at a wage rate w_t earn a gross income $y_t = w_t l_t$. Each period, agents can build their stock of human capital by spending time. These time investments in human capital will be called "training." They could represent time spent in college, in formal classes or training courses, time spent reading, or, importantly, time spent for on-the-job training. Let z_t denote the stock of training time, or the stock of human capital of an agent, at time t and i_t the incremental training time acquired in period t. Human capital z_t, evolves according to

$$z_t = z_{t-1} + i_t.$$

The disutility cost to an agent who provides $l_t \geq 0$ units of work and spends $i_t \geq 0$ units in training

[5]Note that under perfect information, in the first best, and with no other frictions, any time cost could be converted into an equivalent monetary cost and the distinction between time and money investments would be blurred. This is no longer true under asymmetric information, unless training does not interact with labor supply, a special case which will be explored in the paper.

is $\phi_t(l_t, i_t)$. ϕ_t is strictly increasing and convex in each of its arguments ($\frac{\partial\phi(l,i)}{\partial i}, \frac{\partial\phi(l,i)}{\partial l} > 0, \forall\ l, i > 0$; $\frac{\partial^2\phi(l,i)}{\partial l^2}, \frac{\partial^2\phi(l,i)}{\partial i^2} \geq 0, \forall l, i$); however, no assumption is yet made on the cross partial $\frac{\partial^2\phi(l,i)}{\partial l \partial i}$.[6]

The wage rate w_t is determined by the training acquired as of time t and by a stochastic ability θ_t:

$$w_t = w_t(\theta_t, z_t)$$

w_t is strictly increasing and concave in each of its arguments ($\frac{\partial w}{\partial m} > 0$, $\frac{\partial^2 w}{\partial m^2} \leq 0$ for $m = \theta, z$). Again, there are no restrictions placed on the cross-partials yet. Note that because the wage depends on time t, which is the equivalent of age for the given cohort, human capital and ability are allowed to have different wage impacts at different times in life.

Agents start their life at time $t = 1$ with a heterogeneous type θ_1, distributed according to a density function $f^1(\theta_1)$. Earning ability in each period is private information, and evolves according to a Markov process with a time-varying transition function $f^t(\theta_t|\theta_{t-1})$, over a fixed support $\Theta \equiv [\underline{\theta}, \overline{\theta}]$. The ability shock θ can be interpreted in several different ways. It could be a health shock, an idiosyncratic labor market shock, or a shock to the return to an agent's skills. To follow the taxation literature, θ_t will simply be called "ability" throughout.

The agent's per period utility is separable in consumption and in the effort required for labor and training:

$$\tilde{u}_t(c_t, y_t, z_t; \theta_t, z_{t-1}) = u_t(c_t) - \phi_t\left(\frac{y_t}{w_t(\theta_t, z_t)}, z_t - z_{t-1}\right)$$

u_t is increasing, twice continuously differentiable, and concave.

Denote by θ^t the history of ability shocks up to period t, by Θ^t the set of possible histories at t, and by $P(\theta^t)$ the probability of a history θ^t, $P(\theta^t) \equiv f^t(\theta_t|\theta_{t-1}) ... f^2(\theta_2|\theta_1) f^1(\theta_1)$. An allocation $\{x_t\}_t$ specifies consumption, output, and the training stock for each period t, conditional on the history θ^t, i.e., $x_t = \{x(\theta^t)\}_{\Theta^t} = \{c(\theta^t), y(\theta^t), z(\theta^t)\}_{\Theta^t}$. The expected lifetime utility from an allocation, discounted by a factor β, is given by:

$$\begin{aligned} & U(\{c(\theta^t), y(\theta^t), z(\theta^t)\}) \\ & = \sum_{t=1}^{T} \int \beta^{t-1} \left[u_t(c(\theta^t)) - \phi_t\left(\frac{y(\theta^t)}{w_t(\theta_t, z(\theta^t))}, z(\theta^t) - z(\theta^{t-1})\right)\right] P(\theta^t) d\theta^t \end{aligned} \quad (1)$$

where, with some abuse of notation, $d\theta^t \equiv d\theta_t ... d\theta_1$.

Let $w_{m,t}$ denote the partial of the wage function with respect to argument m ($m \in \{\theta, z\}$), and $w_{mn,t}$ the second order partial with respect to arguments m,$n \in \{\theta, z\} \times \{\theta, z\}$.

There are two important parameters that drive the optimal policies. The first is the Hicksian coefficient of complementarity between ability and training in the wage function at time t, denoted by

[6]The Appendix links this model to the standard Ben-Porath model.

$\rho_{\theta z}$:[7]

$$\rho_{\theta z} \equiv \frac{w_{\theta z} w}{w_z w_\theta} \tag{2}$$

A positive Hicksian complementarity between human capital z and ability θ means that the effect of ability on the wage is amplified by human capital z, i.e, that human capital and ability are complements ($w_{\theta z} \geq 0$). Put differently, human capital compounds the exposure of the agent to stochastic ability and to risk. A Hicksian complementarity greater than 1 means that the wage elasticity with respect to ability is increasing in human capital, i.e., $\frac{\partial}{\partial z}\left(\frac{\partial w}{\partial \theta}\frac{\theta}{w}\right) \geq 0$.[8] Some simple cases can illustrate the complementarity coefficient: A separable wage such as $w_t = \theta_t + h_t(z_t)$, where h_t is an increasing, concave function, has $\rho_{\theta z,t} = 0$. A multiplicative wage $w_t = \theta_t h_t(z)$ has $\rho_{\theta z,t} = 1$. A CES wage such as:

$$w_t = \left[\alpha_{1t}\theta^{1-\rho_t} + \alpha_{2t} z^{1-\rho_t}\right]^{\frac{1}{1-\rho_t}} \tag{3}$$

has $\rho_{\theta z,t} = \rho_t$.

2.1.1 Learning-and-doing and Learning-or-doing

For training time, the most important parameter will be the one that characterizes the interaction between training and labor. Define ρ_{lz}^{ϕ} to be the Hicksian complementarity coefficient between labor and training in the disutility function ϕ:

$$\rho_{lz}^{\phi} \equiv \frac{\phi_{lz}\phi}{\phi_l \phi_z} \tag{4}$$

where the partials and second order partials of ϕ are denoted by: $\phi_{z,t} \equiv \frac{\partial \phi_t}{\partial z_t}$, $\phi_{l,t} \equiv \frac{\partial \phi_t}{\partial l_t}$, $\phi_{lz,t} \equiv \frac{\partial^2 \phi}{\partial l \partial z}$.

How can we interpret this disutility function? It is really just a general way of capturing everything in between two polar cases we often consider in the labor literature, namely an "opportunity cost of time" specification, as in the canonical Ben-Porath (1967) model, and the "learning-by-doing" model of Arrow (1962). In the former, training time and work time are perfect substitutes and time spent training cannot be spent working. This can be captured by a disutility function linear in i and l, with a total time endowment constraint. In the learning-by-doing case, the disutility is a Leontieff function.[9]

But in between these two polar cases, there are many other possibilities. A less extreme case of opportunity cost of time is when training and labor are substitutes, but not perfect ones. Time spent training takes away effort from working. Fatigue effects may set in, at a convex rate. In this case, the

[7] On studies of the Hicksian complementarity as see Hicks (1970), Samuelson (1973), Bovenberg and Jacobs (2011).

[8] Equivalently, the wage elasticity with respect to human capital is increasing in ability.

[9] This equivalence between training and labor in the learning-by-doing case ignores the fact that, in the planning problem, labor is unobservable while training is not. Section 6.2 will consider this limit case in more detail.

Hicksian complementarity between training and labor in the disutility function is positive ($\rho^{\phi}_{lz,t} > 0$), a case I will label "Learning-or-doing." Similarly, a less extreme case of learning-by-doing is when labor supply and training are complements over at least some range ($\rho^{\phi}_{lz,t} < 0$). The reasoning could be the same as the standard learning-by-doing justification: skills or experience acquired by working might make training less costly at the margin. Or, less commonly thought, the guarantee of regular training, which gives workers a prospective for progress, might boost motivation and make labor effort seem less painful. I will call this case "Learning-and-doing."

In addition, at different levels of training and labor supply, different effects could dominate: for instance, labor and training could be complementary at low levels, but substitutable at high levels of labor supply. The formulation is general enough to accommodate all these cases. Of course, the results on optimal training policies in both the static and dynamic cases (Sections 3 and 5) will strongly depend on this interaction between training and labor.

2.2 First best planning problem

Before turning to the planning problem under asymmetric information, let us first consider the first best solution that would arise if the Planner had the same information as the agent. In the first best, the Planner seeks to maximize the lifetime expected utility in (1) subject to the intertemporal resource constraint, where R is the gross, exogenously given, borrowing rate:

$$\sum_{t=1}^{T}\left(\frac{1}{R}\right)^{t-1}\int_{\Theta^t}\left(c_t\left(\theta^t\right)-y_t\left(\theta^t\right)\right)P\left(\theta^t\right)d\theta^t\leq 0$$

Suppose that $\beta = \frac{1}{R}$. There is perfect insurance against earnings risk, as consumption is equalized across all histories: $u'\left(c_t\left(\theta^t\right)\right) = \lambda$, with λ the multiplier on the resource constraint. Output and training for each θ^t, given any level z_{t-1}, solve simultaneously:

$$\phi_{l,t}\left(\frac{y_t}{w_t\left(\theta_t,z_t\right)},z_t-z_{t-1}\right)=w_t\left(\theta_t,z_t\right)\lambda$$

$$\phi_{z,t}\left(\phi^{-1,l}_{l,t}\left(w_t\lambda,z_t-z_{t-1}\right),z_t-z_{t-1}\right)=\sum_{\tau=t}^{T}\frac{1}{R^{\tau-t}}E_t\left(w_{z,\tau}\phi^{-1,l}_{l,\tau}\left(w_\tau\lambda,z_\tau-z_{\tau-1}\right)\right)\lambda$$

Hence, higher ability agents produce more output, for a given level of training. The optimal training equates the marginal disutility of training to its marginal benefit. The marginal expected benefit of training is increasing in ability, as long as training and ability are complements in the wage function $\left(\frac{\partial^2 w}{\partial\theta\partial z}\geq 0\right)$ and assumption 2 i) holds. The cost of training, however, can also rise or fall with ability depending on whether there is learning-or-doing or learning-and-doing. If the latter is prevalent, or if there is no interaction between labor supply and training (ϕ is separable), then unambiguously higher types both work more and train more. However, if learning-or-doing is prevalent, higher types face both a higher marginal benefit, and a higher marginal disutility from training. The net effect is ambiguous.

When there is asymmetric information about the types θ between the Planner and the agents, this first-best allocation is no longer achievable. The second-best allocation will have to take into account the agents' incentive compatibility constraints, an issue to which I turn now.

3 Illustrating the Optimal Policies in a Simple One-Period Model

Before solving the full-fledged dynamic planning problem, it is informative to draw out all the main effects in a one-period model. The dynamic solution in Section 5 will extend these results, and it is very helpful to build up the intuition from the static case.

3.1 Gross wedges versus net wedges

In the second best, ability θ is private information to the agent. Consumption and output are observable. Because of the asymmetric information, the allocations will be distorted. The marginal distortions in agents' choices can be described using "wedges," which represent the implicit local marginal tax and subsidy rates (see Golosov *et al.* (2006) for an introduction to the wedges in dynamic public finance). The wedges capture the magnitudes and signs of distortions in the allocation relative to the laissez-faire (or, autarky) allocation. They are just definitions of variables that, evaluated at the optimal second-best allocation, will characterize that allocation in a way analogous to taxes. Expressing the properties of the optimal allocation in terms of wedges allows us to appeal to intuitions related to a standard tax system.

In the one-period model, suppose that types θ are distributed according to a distribution $f(\theta)$ on Θ and that the starting stock of training is zero. Hence, $z = i$.

For a given type θ, define the labor wedge $\tau_L(\theta)$ and the training wedge $\tau_Z(\theta)$ as follows:

$$\tau_L(\theta) \equiv 1 - \frac{\phi_l(l,z)}{w(\theta,z)u'(c)} \tag{5}$$

$$\tau_Z(\theta) \equiv \frac{\phi_z}{u'(c)} - (1-\tau_L(\theta))w_z l \tag{6}$$

The labor wedge τ_L is defined as the gap between the marginal rate of substitution and the marginal rate of transformation between consumption and labor. It is akin to a locally linear marginal tax because, without any government intervention, it would be zero. On the other hand, if the government would set the marginal labor tax for agents of type θ equal to $\tau_L(\theta)$ and let these agents optimize their labor supply in a small neighborhood around their optimal allocation, and conditional on the optimal training choice, they would choose their labor supply so that this relation would hold.

The implicit subsidy on training τ_Z can be thought of as the incremental pay received by an agent for training for one more unit of time. If the implicit subsidy is positive, human capital is distorted upwards. Like any subsidy, it is defined as the gap between marginal cost and marginal benefit. The only subtlety is that, since training entails a disutility cost, the right-hand side is scaled by marginal utility in order to convert it to a dollar amount.

Without the presence of training, the wedges just defined would be sufficient to characterize the allocation. However, there are two simultaneous distortions: saying that the wegde on human capital, τ_Z, is zero does not actually mean human capital is undistorted – if there is a positive labor wegde $\tau_L > 0$, then human capital is still distorted downwards, as part of its return is taxed away.

It would thus be useful to find a measure of the *net* distortion on human capital, the one that goes beyond just compensating for the presence of a labor distortion. We can ask: what is the subsidy on human capital needed to exactly undo the labor distortion on human capital?

To answer this, let's consider a thought-experiment. Suppose that we set the training subsidy to be $\tau_Z = \tau_L(\frac{\phi_z}{u'(c)})$. Then, locally, for a given labor supply l, the agent solves:

$$\max_z u(w(z,\theta)l(1-\tau_L) + \tau_Z z) - \phi(z,l)$$

The first-order condition of the agent with respect to training then yields:

$$w_z l u'(c)(1-\tau_L) + u'(c)\tau_L \frac{\phi_z}{u'(c)} - \phi_z(z,l) = 0$$

where we have used the supposed relation between τ_L and τ_Z. This implies that: $(1-\tau_L)(w_z l u'(c) - \phi_z) = 0$, i.e., conditional on the labor choice, human capital is set optimally.

This relation between τ_L and τ_Z is the equivalent of making training costs fully tax deductible from the income tax base. More precisely, the agent is allowed to convert his disutility costs into monetary equivalents, and then deduct this monetary equivalent from his income tax base.[10]

Inspired by this example, let us define the net wedge on training as the gross wedge minus the amount needed to compensate the agent for the distortion on labor.

Definition 1 *The (static) net wedge on training time t_z is defined as:*

$$t_z \equiv \frac{\tau_Z - \tau_L(\frac{\phi_z}{u'(c)})}{((\frac{\phi_z}{u'(c)}) - \tau_Z)(1-\tau_L)} \tag{7}$$

As a matter of definition, the net wedge is scaled in the denominator by the net (monetary) cost of training and the net of tax rate. This will serve to make the optimal formulas simpler.

[10] This idea of full deductibility holds perfectly locally and conditional on the labor choice. Of course, given that the wedges are nonlinear and conditional on the type, the analogy does not hold globally.

It will be useful to denote by ε_{xy} the elasticity of any variable x with respect to variable y, $\varepsilon_{xy} \equiv d\log(x)/d\log(y)$. Let ε^u and ε^c be the uncompensated and compensated labor supply elasticities to the net wage, holding training fixed.[11]

3.2 The one-period planning problem

I briefly set up the planning problem in the static case: the full-fledged dynamic problem is rigorously explained in Section 5 and nests this simple static case. Hence, many technical details (in particular, about how the first-order approach works) are omitted here. Hence, the reader interested in the derivations can refer to Section 5 and to the proofs in the Appendix.

Let $\omega(\theta)$ denote the utility of an agent of type θ:

$$\omega(\theta) = u(c(\theta)) - \phi(l(\theta), z(\theta)).$$

I use a first-order approach that consists in replacing the incentive compatibility constraints by envelope conditions of the agents. The envelope condition tells us how the utility for each type has to vary at the optimum and how the informational rent received by each type evolves. In this one-period case, the envelope condition is simply:[12]

$$\dot{\omega}(\theta) := \frac{\partial \omega(\theta)}{\partial \theta} = \frac{w_\theta}{w} l(\theta)\phi_l(l(\theta), z(\theta))$$

The one-period planning problem is then to minimize resources:

$$\min \int (c(\theta) - l(\theta)w(z(\theta), \theta))f(\theta)d\theta$$

subject to the envelope condition and the utility constraint for each type:

$$\omega(\theta) \geq \underline{\omega}(\theta).$$

where the vector of lower bounds for utility is arbitrarily given.

3.3 The optimal training wedge

The next proposition (a special case of the full fledged, dynamic optimum in Proposition 2) shows how the net wedge is set at the optimum.

Proposition 1 *At the optimum, the net subsidy for training and the labor wedge are set according to:*

$$t_z^*(\theta) = \frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} \frac{\varepsilon^c}{1 + \varepsilon^u} (1 - \rho_{\theta z} - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}} \rho_{lz}^{\phi}) \tag{8}$$

[11] With utility separable in consumption and labor, $\frac{\varepsilon^c}{1+\varepsilon^u-\varepsilon^c}$ is the Frisch elasticity of labor.
[12] This is just a special case of the dynamic envelope condition in (16).

where: $\varepsilon_{\phi z} \equiv d\log(\phi)/d\log(z)$ *is the elasticity of disutility with respect to training time.* ρ^{ϕ}_{lz} *is the Hicksian coefficient of complementarity between* l *and* z *in the disutility* ϕ *as defined in (4).*

In this one-period model, there are three effects from subsidizing training that are balanced at the optimum. First, a higher net wedge increases investments in training, and, hence, increases the wage and labor supply. This is the so-called "labor supply effect." However, there is also a differential effect on different ability types. If $\rho_{\theta z} > 0$, training benefits able agents more and further spreads out earnings inequality. On the other hand, if $\rho_{\theta s} < 0$, earnings inequality is compressed by training. This is the so-called "inequality effect." Given that ability is stochastic here, a complementarity to ability also increases risk exposure.

In mechanism design language, the inequality effect is just the result of the informational rent that needs to be forfeited to higher ability agents. When $\rho_{\theta z} > 0$, the informational rent increases in training.

The third effect is the direct interaction with labor supply through the disutility function, i.e., either learning-or-doing or learning-and-doing. In the formula, $\left(1 - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho^{\phi}_{lz,t}\right)$ is the *total* effect of the bonus on labor, i.e., the sum of the labor supply effect (the increase in labor due to the higher wage), and either one of the learning-and-doing or learning-or-doing effects, which can increase or decrease work.

Rearranging this, we can see that the total effect of a training subsidy on labor is positive if and only if

$$\varepsilon_{wz} > \varepsilon_{\phi_l z} \quad \text{with } \varepsilon_{\phi_l z} \equiv \partial \log\left(\phi_l\right) / \partial \log\left(z\right)$$

i.e., if and only if the wage is more sensitive to training than the marginal disutility of work is.[13]

The question then is, whether the increase in total resources from the total labor effect of training more than makes up for the increased rent transfers (the inequality effect). The net subsidy on training will be positive if and only if the answer to this question is yes.

Corollary 1 *The static net subsidy on training is positive if and only if:*

$$1 - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho^{\phi}_{lz} > \rho_{\theta z} \tag{9}$$

With learning-and-doing ($\rho^{\phi}_{lz} < 0$), as long as ability and human capital are not too complementary (say, $\rho_{\theta z} < 1$), the net subsidy on training is positive. Intuitively, training does not distract from labor effort, and so it's good to encourage it as long as high types do not disproportionately benefit from it. However, if there is learning-or-doing, training makes labor supply more costly. In this case, even

[13] Note that: $\left(\varepsilon_{\phi z,t}/\varepsilon_{wz,t}\right)\rho^{\phi}_{lz,t} = \varepsilon_{\phi_l z,t}/\varepsilon_{wz,t}$.

if the coefficient of complementarity $\rho_{\theta z}$ is small, it might not be sufficiently small to compensate for the lost work effort.

Testing for Pareto efficiency: Note that the simple relation derived between the labor wedge and the training wedge in (8) can be used to test for the optimality of an allocation.

In the special cases highlighted in the next two subsections, that relation is constant over time and does not depend on the exact allocation. In this case, we can use the ratio of the labor wedge and the training wedge to test for Pareto efficiency of any tax schedule.

As a special case, note that if the wage is multiplicatively separable and ϕ is additively separable, we have $\rho_{\theta z} = 1$ and $\rho^{\phi}_{lz} = 0$, and, accordingly, $t_z = 0$. This case turns out to be of special interest, as explained next.

3.4 A special case: an application of Atkinson-Stiglitz

The special case in which the wage is multiplicatively separable and ϕ is additively separable is an application of the Atkinson-Stiglitz theorem (Atkinson and Stiglitz, 1976). Let the separable disutility function $\phi(l, z)$ be given by: $\phi(l, z) = \phi^1(l) + \phi^2(z)$. The separable utility then satisfies all assumptions of the Atkinson-Stiglitz theorem, where the two commodities are c and z.

We can directly show that, in this case, the tax schedule should be neutral with respect to human capital by using a simple variational argument.

Recall that the objective is to maximize net resources (or minimize net cost):

$$c(\theta, \hat{z}(\theta)) - l(\theta, \hat{z}(\theta))w(\theta, \hat{z}(\theta)) \tag{10}$$

where $\hat{z}(\theta)$ is a perturbation from $z(\theta)$ and $c(\theta, \hat{z}(\theta)), l(\theta, \hat{z}(\theta))$ are the solutions to the two constraints below, as a function of $\hat{z}(\theta)$. Taking the first-order condition with respect to $\hat{z}$ yields:

$$\frac{\partial c}{\partial \hat{z}} - \frac{\partial l}{\partial \hat{z}} w(\theta, \hat{z}) - l(\theta, \hat{z}) \frac{\partial w}{\partial \hat{z}} \tag{11}$$

There are, as previously explained, two constraints on this problem. The first one is the envelope condition for an agent of type θ, given here by:

$$\dot{\omega}(\theta) = \frac{w_\theta}{w} l(\theta) \phi^{1\prime}(l(\theta)) \tag{12}$$

The case with $\rho_{\theta z} = 1$ is the case where $\frac{w_\theta}{w}$ does not change as z changes, so that at the optimum given these parameters, this constraint yields $dl/dz = 0$.

The second constraint is that utility for each type is unchanged. Using the implicit function theorem gives:

$$u'(c(\theta))dc - \phi^{1\prime}(l(\theta))dl - \phi^{2\prime}(z(\theta))dz = 0 \tag{13}$$

Substituting this into the first-order condition, yields:

$$\frac{\phi^{2\prime}(z(\theta))}{u'(c(\theta))} - l(\theta)\frac{\partial w}{\partial z} = 0 \tag{14}$$

Hence, conditional on $l(\theta)$ and $c(\theta)$, $z(\theta)$ should not be distorted. If we replace the definitions of τ_Z and τ_L from above, and rearrange to make t_z appear, this will be equivalent to $t_z = 0$.

3.5 Other Special cases

There are three other special cases that are illustrative because they yield particularly simple relations at the optimum between the training wedge and the labor wedge. Note that each of these special relations will also hold in the dynamic model, if human capital fully depreciates between periods.

Multiplicatively separable wage and Cobb-Douglas disutility: Suppose that the wage is multiplicatively separable

$$w(\theta, z) = \theta z$$

and the disutility is Cobb-Douglas:

$$\phi(l, z) = \frac{1}{\gamma\alpha} l^{\gamma} z^{\alpha}.$$

In this case, $\rho^{\phi}_{lz} = 1$, $\varepsilon_{\phi z} = \alpha$, and $\varepsilon_{wz} = 1$, so that the formula for the optimal wedge yields a very simple negative relation between the optimal labor wedge and the optimal training wedge at any point in the skill distribution:

$$t^*_z(\theta) = -\frac{\tau^*_L(\theta)}{1 - \tau^*_L(\theta)} \frac{\alpha}{\gamma}$$

Based on the fact that at every interior type $\tau_L > 0$, we can sign the wedges.[14] In this case, $t_z(\theta) < 0$ at any interior type.

CES wage and separable disutility: Suppose that the wage takes a CES form

$$w(\theta, z) = (\theta^{1-\rho} + z^{1-\rho})^{\frac{1}{1-\rho}}$$

and the disutility is separable with $\phi(l, z) = \phi^1(l) + \phi^2(z)$. Then, we have that $\rho_{\theta z} = \rho$ and $\rho^{\phi}_{lz} = 0$ so that at the optimum:

$$t^*_z(\theta) = \frac{\tau^*_L(\theta)}{1 - \tau^*_L(\theta)} \frac{\varepsilon^c}{1 + \varepsilon^u}(1 - \rho)$$

[14] It is shown that the optimal labor wedge is positive for all interior types in the proof of proposition 2. There, it is also shown that τ_L is optimally zero for the lowest and highest types.

If the disutility is in addition isoelastic in labor with $\phi^1(l) = \frac{1}{\gamma} l^\gamma$, then we have:

$$t_z^*(\theta) = \frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} \frac{(1 - \rho)}{\gamma}$$

Hence, the net wedge on training is positive if and only if $\rho_{\theta z} < 1$, i.e., if ability and training are not too complementary in generating earnings.

Separable wage and Cobb-Douglas disutility: Suppose that the wage is separable with

$$w(\theta, z) = \theta + z$$

and the disutility is again Cobb-Douglas as above. Then the optimal training wedge is set according to:

$$t_z^*(\theta) = \frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} \frac{(1 - \alpha)}{\gamma}$$

In this case, the optimal wedge is again negative since $\alpha > 1$. See also Section 5.3 for a version of this special case in the dynamic model.

Equipped with the intuitions and results from this simple one-period model, I now turn to solving the full dynamic life cycle model.

4 The dynamic planning problem

In the dynamic problem, the planner cannot see the ability realization θ_t in any period. Thus, the wage realization $w_t(\theta_t, z_t)$, or labor supply $l_t = y_t / w_t$ are also unknown. Output y_t, consumption c_t, and training z_t on the other hand are still observable in every period.

The solution method of such a dynamic mechanism design problem is based on Farhi and Werning (2013) and Stantcheva (2014). It is hence presented only briefly here. The reader can refer to Stantcheva (2014) for a more detailed presentation of this approach.

To solve for the constrained efficient allocations, I suppose that the planner designs a direct revelation mechanism, in which agents report their types each period and the planner assigns allocations as a function of the history of reports made. Each agent has a history of reports as a function of his true history of shocks, and this history of reports is denoted by $r^t(\theta^t)$. For any reporting strategy r, define the continuation value after history θ^t to be:

$$\omega^r(\theta^t) = u_t(c(r^t(\theta^t))) - \phi_t\left(\frac{y\left(r^t\left(\theta^t\right)\right)}{w_t(\theta_t, z(r^t(\theta^t)))}, z\left(r^t\left(\theta^t\right)\right) - z\left(r^{t-1}\left(\theta^{t-1}\right)\right)\right) + \beta \int \omega^r\left(\theta^{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1}$$

The continuation value under truthful revelation, $\omega\left(\theta^t\right)$, is the unique solution to:

$$\omega\left(\theta^t\right)=u_t\left(c\left(\theta^t\right)\right)-\phi_t\left(\frac{y\left(\theta^t\right)}{w_t(\theta_t,z(\theta^t))},z\left(\theta^t\right)-z\left(\theta^{t-1}\right)\right)+\beta\int\omega\left(\theta^{t+1}\right)f^{t+1}\left(\theta_{t+1}|\theta_t\right)d\theta_{t+1}$$

Incentive compatibility requires that truth-telling yields a weakly higher continuation utility than any reporting strategy r, after all histories θ^t:

$$(IC):\omega\left(\theta^t\right)\geq\omega^r\left(\theta^t\right)\quad\forall\theta^t,\forall r \tag{15}$$

Assumption 1 *i)* $\tilde{u}_t\left(c,y,z;\theta,z_-\right)$ *and* $\frac{\partial\phi(l,i)}{\partial l}\frac{\partial w(\theta,z)}{\partial\theta}\frac{l}{w}$ *are bounded. ii)* $\frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial\theta_{t-1}}$ *exists and is bounded. iii)* $f^t\left(\theta_t|\theta_{t-1}\right)$ *has full support on* Θ.

The first order approach consists in replacing the incentive constraint by the envelope condition of the agent. As in Section 3, it constrains the informational rent of each type of agent. In this dynamic model, there is an additional future term that captures the fact that the agent has some advance information about his future type because of the Markov persistence.

$$\dot{\omega}\left(\theta^t\right):=\frac{\partial\omega\left(\theta^t\right)}{\partial\theta_t}=\frac{w_{\theta,t}}{w_t}l\left(\theta^t\right)\phi_{l,t}\left(l\left(\theta^t\right),z_t\left(\theta^t\right)-z_{t-1}\left(\theta^{t-1}\right)\right)+\beta\int\omega\left(\theta^{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_t\right)}{\partial\theta_t}d\theta_{t+1} \tag{16}$$

The envelope condition is only a necessary condition for incentive compatibility.[15] I use a numerical ex post verification procedure to ensure that the solution found is indeed incentive compatible (see Farhi and Werning (2013) and Stantcheva (2014) for more details).

The next step is to rewrite the problem recursively. Denote the gradient of the expected future continuation utility with respect to the type by:

$$\Delta\left(\theta^t\right)\equiv\int\omega\left(\theta^{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_t\right)}{\partial\theta_t}d\theta_{t+1} \tag{17}$$

The envelope condition can then be rewritten as:

$$\dot{\omega}\left(\theta^t\right)=\frac{w_{\theta,t}}{w_t}l\left(\theta^t\right)\phi_{l,t}\left(l\left(\theta^t\right),z_t\left(\theta^t\right)-z_{t-1}\left(\theta^{t-1}\right)\right)+\beta\Delta\left(\theta^t\right) \tag{18}$$

Let $v\left(\theta^t\right)$ be the expected future continuation utility:

$$v\left(\theta^t\right)\equiv\int\omega\left(\theta^{t+1}\right)f^{t+1}\left(\theta_{t+1}|\theta_t\right)d\theta_{t+1} \tag{19}$$

Continuation utility $\omega\left(\theta^t\right)$ can hence be rewritten as:

$$\omega\left(\theta^t\right)=u_t\left(c\left(\theta^t\right)\right)-\phi_t\left(\frac{y\left(\theta^t\right)}{w_t\left(\theta_t,z\left(\theta^t\right)\right)},z\left(\theta^t\right)-z\left(\theta^{t-1}\right)\right)+\beta v\left(\theta^t\right) \tag{20}$$

[15] An application of Theorem 2 in Milgrom and Segal (2002), under assumption 1.

Define the expected continuation cost of the planner at time t, given v_{t-1}, Δ_{t-1}, θ_{t-1}, and z_{t-1}:

$$K(v, \Delta, \theta_{t-1}, z_{t-1}, t) = \min[\sum_{\tau=t}^{T}(\frac{1}{R})^{\tau-t}\int(c_\tau(\theta^\tau) - y_\tau(\theta^\tau)P(\theta^{\tau-t})d\theta^{\tau-t}]$$

where, with some abuse of notation, $d\theta^{\tau-t} = d\theta_\tau d\theta_{\tau-1}...d\theta_t$, and $P\left(\theta^{\tau-t}\right) = f^\tau\left(\theta_\tau|\theta_{\tau-1}\right)...f^t\left(\theta_t|\theta_{t-1}\right)$.

A recursive formulation of the relaxed program is then for $t \geq 2$:

$$K\left(v, \Delta, \theta_-, z_-, t\right) = \min\int\left(c\left(\theta\right) - w_t\left(\theta, z\left(\theta\right)\right)l\left(\theta\right) + \frac{1}{R}K\left(v\left(\theta\right), \Delta\left(\theta\right), \theta, z\left(\theta\right), t+1\right)\right)f^t\left(\theta|\theta_-\right)d\theta \quad (21)$$

subject to:

$$\begin{aligned}
\omega\left(\theta\right) &= u_t\left(c\left(\theta\right)\right) - \phi_t\left(l\left(\theta\right), z\left(\theta\right) - z_-\right) + \beta v\left(\theta\right)\\
\dot{\omega}\left(\theta\right) &= \frac{w_{\theta,t}}{w_t}l\left(\theta\right)\phi_{l,t}\left(l\left(\theta\right), z\left(\theta\right) - z_-\right) + \beta\Delta\left(\theta\right)\\
v &= \int\omega\left(\theta\right)f^t\left(\theta|\theta_-\right)d\theta\\
\Delta &= \int\omega\left(\theta\right)\frac{\partial f^t\left(\theta|\theta_-\right)}{\partial\theta_-}d\theta
\end{aligned}$$

where the maximization is over $\left(c\left(\theta\right), l\left(\theta\right), z\left(\theta\right), \omega\left(\theta\right), v\left(\theta\right), \Delta\left(\theta\right)\right)$.

For period $t = 1$, suppose all agents have identical initial human capital level z_0.[16] The problem for $t = 1$ is then indexed by $\left(\underline{U}\left(\theta\right)\right)_\Theta$, the set of target lifetime utilities $\underline{U}\left(\theta\right)$ for each type θ:

$$K\left(\left(\underline{U}\left(\theta\right)\right)_\Theta, 1\right) = \min\int\left(c\left(\theta\right) - w_1\left(\theta, z\left(\theta\right)\right)l\left(\theta\right) + \frac{1}{R}K\left(v\left(\theta\right), \Delta\left(\theta\right), \theta, z\left(\theta\right), 2\right)\right)f^1\left(\theta\right)d\theta$$

$$\begin{aligned}
\text{s.t.} &: \omega\left(\theta\right) = u_1\left(c\left(\theta\right)\right) - \phi_1\left(l\left(\theta\right), z\left(\theta\right) - z_0\right) + \beta v\left(\theta\right)\\
\dot{\omega}\left(\theta\right) &= \frac{w_{\theta,1}}{w_1}l\left(\theta\right)\phi_{l,1}\left(l\left(\theta\right), z\left(\theta\right)\right) + \beta\Delta\left(\theta\right)\\
\omega\left(\theta\right) &\geq \underline{U}\left(\theta\right)
\end{aligned}$$

Technical assumptions: The following assumptions are sufficient conditions to determine the sign of the optimal wedges and are used only for that purpose in some proofs.[17]

Assumption 2 *i)* $\int_{\underline{\theta}}^{\theta'} f^t\left(\theta|\theta_b\right)d\theta \leq \int_{\underline{\theta}}^{\theta'} f^t\left(\theta|\theta_s\right)d\theta$, $\forall t, \theta'$, *and* $\theta_b > \theta_s$;
ii) $\frac{\partial}{\partial\theta_t}\left(\frac{\partial f^t\left(\theta_t|\theta_{t-1}\right)}{\partial\theta_{t-1}}\frac{1}{f^t\left(\theta_t|\theta_{t-1}\right)}\right) \geq 0$, $\forall t$, $\forall\theta_{t-1}$;
iii) $\frac{\partial v(\theta)}{\partial\theta} > 0$ *for all* θ,
iv) $\frac{\partial}{\partial v}K \geq 0$ *and* $\frac{\partial^2}{\partial v^2}K \geq 0$.

They guarantee (in this order) that that there is first-order stochastic dominance, that a form of monotone likelihood ratio property is satisfied, that the expected continuation utility is increasing in the type, and that the resource cost is increasing and convex in promised utility.

[16] Otherwise, given that human capital is observable, we would need to condition on the initial levels.
[17] In addition, all theoretical results on the signs of the wedges are indeed satisfied in the simulations (section 7).

5 Optimal dynamic training policies

This section characterizes the allocations, obtained as solutions to the relaxed recursive program above.

5.1 The dynamic gross and net wedges

As already explained in the one-period model in Section 3, marginal distortions in agents' choices can be described using wedges, which represent the implicit local marginal tax and subsidy rates. In the dynamic program, wedges are defined after each history. They hence represent implicit taxes and subsidies, locally, and conditional on a given history.

Define the intratemporal wedge on labor $\tau_L\left(\theta^t\right)$, the intertemporal wedge on savings or capital $\tau_K\left(\theta^t\right)$, and the human capital wedge $\tau_Z\left(\theta^t\right)$ as follows:

$$\tau_L\left(\theta^t\right) \equiv 1-\frac{\phi_{l,t}\left(l_t, z_t-z_{t-1}\right)}{w_t\left(\theta_t, z_t\right) u_t'\left(c_t\right)} \tag{22}$$

$$\tau_K\left(\theta^t\right) \equiv 1-\frac{1}{R\beta}\frac{u_t'\left(c_t\right)}{E_t\left(u_t'\left(c_{t+1}\right)\right)} \tag{23}$$

$$\tau_Z\left(\theta^t\right) \equiv -\left(1-\tau_L\left(\theta^t\right)\right) w_{z,t} l_t+\left[\frac{\phi_{z,t}}{u_t'\left(c_t\right)}-\beta E_t\left(\frac{u_{t+1}'\left(c_{t+1}\right)}{u_t'\left(c_t\right)}\frac{\phi_{z,t+1}}{u_{t+1}'\left(c_{t+1}\right)}\right)\right] \tag{24}$$

The intratemporal labor wedge was already explained in Section 3. The savings or capital wedge τ_K is defined as the difference between the expected marginal rate of intertemporal substitution and the return on savings and is standard in the dynamic public finance literature (see Golosov *et al.* (2006)).

As in the static case, the training subsidy τ_Z can be viewed as the incremental pay received by an agent for one more unit of training time. It is again equal to the difference between marginal cost and marginal benefit. But the marginal benefit now lasts into all future periods. To write this recursively, I express it as a function of the future marginal cost. Solved forward, (24) would yield the full future stream of marginal benefits. Recall that the marginal disutility cost is scaled by marginal utility in order to be expressed as an equivalent monetary cost.

The following definitions will be needed for the dynamic formulas below. For any variable x, define the "insurance factor" of x, $\xi_{x,t+1}$:

$$\xi_{x,t+1}\equiv Cov\left(-\beta\frac{u_{t+1}'}{u_t'}, x_{t+1}\right)/\left(E_t\left(\beta\frac{u_{t+1}'}{u_t'}\right)E_t\left(x_{t+1}\right)\right)$$

with $\xi_{x,t+1}\in[-1,1]$. If x is a flow to the agent, it is a good hedge if $\xi<0$, and a bad hedge otherwise. With some abuse of notation, define also:

$$\xi_{x,t+1}'\equiv -Cov\left(\frac{\beta u_{t+1}'}{u_t'}-1, x_{t+1}\right)/\left(E_t\left(\frac{\beta u_{t+1}'}{u_t'}-1\right)E_t\left(x_{t+1}\right)\right)$$

which, up to an additive constant, captures the same risk properties as $\xi_{x,t+1}$.

Denote by $\varepsilon_{xy,t}$ the elasticity of x_t to y_t, $\varepsilon_{xyt} \equiv d\log(x_t)/d\log(y_t)$. Let ε_t^u and ε_t^c be the uncompensated and compensated labor supply elasticities to the net wage, holding both savings and training fixed.[18]

As in the simple one-period model of Section 3, it would be useful to have a measure of the net distortion and real incentive effect on training, that takes into account the need to filter out whatever goes towards compensating for other distortions. By analogy to the one-period case, we can define a net wedge on training, t_{zt}, that looks more complicated, but follows the same logic.

Definition 2 *The net wedge or net bonus on training time t_{zt} is defined as:*

$$t_{zt} \equiv \frac{\tau_{Zt} - \tau_{Lt}\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d + P_{Zt}}{\left(\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d - \tau_{Zt}\right)(1-\tau_{Lt})} \tag{25}$$

$\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d \equiv \frac{\phi_{z,t}}{u_t'(c_t)} - \frac{1}{R(1-\tau_K)}\left(1-\xi_{\phi'/u',t+1}\right)E_t\left(\frac{\phi_{z,t+1}}{u_{t+1}'(c_{t+1})}\right)$ *is the dynamic risk-adjusted disutility cost, converted into monetary units.*

$P_{Zt} \equiv \frac{\tau_K}{R(1-\tau_K)}\left(1-\xi'_{\phi'/u',t+1}\right)(1-\tau_{Lt})E_t\left(\frac{\phi'_{z,t+1}}{u'_{t+1}(c_{t+1})}\right)$ *is the risk-adjusted savings distortion.*

Setting the net wedge to zero can be thought of as making training costs locally fully tax deductible from taxable income *in a dynamic risk-adjusted way.* Recall that in the static case, a zero net wedge corresponded to contemporaneous full deductibility. In the dynamic case, we need to take into account three more effects. First, the cost is dynamic, due to the fact that training acquired today persists into the future. Second, marginal utility is not smooth across states and time and, hence, the agent does not value one dollar of transfer today the same as one dollar of transfer tomorrow. Thus, we need to reweight any monetary amount by the ratio of marginal utilities. Finally, there is a savings wedge that distorts intertemporal transfers and which also needs to be filtered out.

The net wedge is positive if the government wants to encourage training on net, i.e., above and beyond simply compensating for the distortive effects of the other wedges.

5.2 The optimal dynamic training subsidy

The following proposition characterizes the net training subsidy at the optimum.

[18] I.e., ε^c and ε^u are defined as in the static framework (e.g., Saez, 2001), at constant savings and training time:

$$\varepsilon^u = \frac{\phi_l(l,i)/l + \frac{\phi_l(l,i)^2}{u'(c)^2}u''(c)}{\phi_{ll}(l,i) - \frac{\phi_l(l)^2}{u'(c)^2}u''(c)} \qquad \varepsilon^c = \frac{\phi_l(l,i)/l}{\phi_{ll}(l,i) - \frac{\phi_l(l)^2}{u'(c)^2}u''(c)}$$

Again, with per-period utility separable in consumption and labor, $\frac{\varepsilon_t^c}{1+\varepsilon_t^u-\varepsilon_t^c}$ is the Frisch elasticity of labor.

Proposition 2 *i) At the optimum, the net training subsidy is given by:*

$$t_{zt}^{*}\left(\theta^{t}\right)=\frac{\tau_{Lt}^{*}\left(\theta^{t}\right)}{1-\tau_{Lt}^{*}\left(\theta^{t}\right)}\frac{\varepsilon_{t}^{c}}{1+\varepsilon_{t}^{u}}\left(1-\rho_{\theta z,t}-\frac{\varepsilon_{\phi z,t}}{\varepsilon_{wz,t}}\rho_{lz,t}^{\phi}\right)+\frac{1}{R}E_{t}\left(\frac{\tau_{Lt+1}^{*}\left(\theta^{t+1}\right)}{1-\tau_{Lt}^{*}\left(\theta^{t}\right)}\frac{\varepsilon_{t+1}^{c}}{1+\varepsilon_{t+1}^{u}}\frac{w_{z,t+1}}{w_{z,t}}\frac{l_{t+1}}{l_{t}}\frac{\varepsilon_{\phi z,t+1}}{\varepsilon_{wz,t+1}}\rho_{lz,t+1}^{\phi}\right) \tag{26}$$

with $\varepsilon_{\phi z,t}=d\log\left(\phi_{t}\right)/d\log\left(z_{t}\right)$ *the elasticity of disutility with respect to training time.* $\rho_{lz,t}^{\phi}$ *is the Hicksian coefficient of complementarity between* l *and* z *in the disutility* ϕ_t*. The labor wedges* $\tau_{Lt}^{*}\left(\theta^{t}\right)$ *and* $\tau_{Lt+1}^{*}\left(\theta^{t+1}\right)$ *are at their optimal levels.*

ii) The net training wedges are zero at the bottom and the top of the ability distribution:

$$t_{zt}^{*}(\theta^{t-1},\underline{\theta})=t_{zt}^{*}(\theta^{t-1},\overline{\theta})=0$$

As in the one-period case, we see the labor supply effect, the inequality effect, and the contemporaneous learning-or-doing or learning-and-doing effects. However, in this dynamic model, the subsidy on training time has an additional direct interaction with future labor supply through the disutility function.

Even if training diverts time away from contemporaneous labor supply, the effects on future labor supply can motivate a positive net subsidy. Whatever the pattern of complementarity or substitutability between contemporaneous labor supply and training, the relation between current training and future labor supply is its mirror image. If the former are complements, the latter are substitutes, and vice versa, because investing in human capital today means having to invest less tomorrow to reach any given level.[19]

If there is learning-or-doing, the net wedge co-moves positively with the future income tax rate τ_{Lt+1}, but negatively with the current tax rate τ_{Lt}. Intuitively, if there is a higher contemporaneous wedge on labor, there already is an indirect stimulus to training because training is a substitute for labor; hence the need for an additional stimulus through a net subsidy is reduced. However, a higher future labor wedge is a stimulus for training in the future, which is a substitute for training today. Hence, to stimulate training today, there is a need for a higher net subsidy.

Corollary 2 illustrates two cases, among other possible ones.

Corollary 2 *Under assumption 2:*

i) The net subsidy is positive if $\left(1-\frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho_{lz,t}^{\phi}\right)\geq\rho_{\theta z,t}$ *and* $\rho_{lz,t}^{\phi}\geq 0$.

ii) The net subsidy is negative if $\left(1-\frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho_{lz,t}^{\phi}\right)\leq\rho_{\theta z,t}$ *and* $\rho_{lz,t}^{\phi}\leq 0$.

The sufficient conditions in i) guarantee that, although training diverts time away from contemporaneous labor effort, the total labor supply effect $\left(1-\frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho_{lz,t}^{\phi}\right)$ (the direct labor supply effect plus the

[19] This is because only the flow $i_t = z_t - z_{t-1}$, and not the stock z_{t-1} enters the disutility. This reasoning is also only valid at interior solutions for training.

interaction of labor and training) outweighs the inequality effect $\rho_{\theta z}$. On the other hand, the sufficient conditions in ii) imply that, despite stimulating contemporaneous labor supply through learning-and-doing, training disproportionately benefits high productivity workers.

One can revert back to the one-period special cases in Section 3.5 if it is assumed that the training stock depreciates entirely each period. In this case, all the simple relations derived in that section between the labor wedge and the training wedge continue to hold.

Life cycle evolution of the optimal net training wedge: It is natural to think that the parameters affecting the optimal net wedge in (26) are age-dependent and would be evolving over the life cycle. In particular, the interaction between labor and training could be very different at different ages in life – something that is easily accommodated by the general, age-dependent disutility function ϕ_t. It is an interesting empirical question how labor supply and the acquisition of training evolve over life, and the answer turns out to be important for optimal tax considerations as well.

5.3 A special case: separability between training and labor

The peculiarity of time spent investing in human capital, unlike monetary human capital investments or physical capital, is that it interacts with labor supply. With perfect information this would not matter, as any given disutility cost can be converted into an equivalent monetary cost. This is no longer true under asymmetric information, unless the disutility function is separable in labor and human capital $\left(\frac{\partial^2 \phi}{\partial z \partial l} = 0\right)$. Then, as corollary (3) shows, the training subsidy is set only as a function of its complementarity with ability.

Corollary 3 *If* $\frac{\partial^2 \phi_t}{\partial z \partial l} = 0$ *and assumption* 2 *holds:*

i) The net wedge satisfies:

$$t^*_{zt} = \frac{\tau^*_{Lt}}{1 - \tau^*_{Lt}} \frac{\varepsilon^c_t}{1 + \varepsilon^u_t} (1 - \rho_{\theta z,t}) \tag{27}$$

ii)

$$t^*_{zt}\left(\theta^t\right) \geq 0 \Leftrightarrow \rho_{\theta z,t} \leq 1$$

ii) If the wage function is CES with parameter ρ *as in* (3) *and the disutility is separable and isoelastic*

$$\phi(l, i) = \frac{1}{\gamma} l^\gamma + \frac{1}{\eta} i^\eta \quad (\gamma > 1, \eta > 1), \tag{28}$$

then the net wedge satisfies:

$$t^*_{zt} / \left(\frac{\tau^*_{Lt}}{1 - \tau^*_{Lt}}\right) = \frac{(1 - \rho)}{\gamma}$$

First, when there no longer is a direct interaction with labor supply, the net stimulus to training is positive if and only if training is not too complementary to ability ($\rho_{\theta z} \leq 1$). In addition, because the future interaction with labor supply no longer matters, the training wedge and labor wedge are set according to a contemporaneous inverse elasticity rule, in which the training wedge is higher whenever its complementarity with ability is lower. With a CES wage function, this relation is particularly simple: in any period, the relation between the training wedge and the labor wedge needs to stay constant.

6 The optimal labor wedge in the presence of training

What happens to the labor wedge in the presence of training? It is natural to think that the presence of training, given its strong interaction with labor, would affect the optimal labor wedge. However, it turns out that in the case just analyzed, the presence of learning-or-doing or learning-and-doing does not really change the labor wedge. The specification of the model matters a lot for this result. I therefore consider two modifications to the model, which have important implications for the labor wedge. The first is the limit case of learning-by-doing, in which all training comes directly as a by-product of the (unobservable) labor supply. The second is the case in which higher ability agents are not only more productive at producing output, but are also more able at training.

6.1 Does learning-or-doing versus learning-and-doing matter for the optimal labor distortion?

In the model analyzed up to here, the labor wage is only slightly affected by the presence of human capital, but not by either learning-or-doing or learning-and-doing. The full formula for the labor wedge is given in the next proposition.

Proposition 3 *i) At the optimum, the labor wedge is equal to:*

$$\frac{\tau_{Lt}^{*}\left(\theta^{t}\right)}{1-\tau_{L}^{*}\left(\theta^{t}\right)}=\frac{\mu\left(\theta^{t}\right)u_{t}'\left(c\left(\theta^{t}\right)\right)}{f^{t}\left(\theta_{t}|\theta_{t-1}\right)}\frac{\varepsilon_{w\theta,t}}{\theta_{t}}\frac{1+\varepsilon_{t}^{u}}{\varepsilon_{t}^{c}} \tag{29}$$

with $\mu\left(\theta^{t}\right)=\eta\left(\theta^{t}\right)+\kappa\left(\theta^{t}\right)$, *where* $\eta\left(\theta^{t}\right)$ *can be rewritten recursively as a function of the past labor wedge,* $\tau_{L,t-1}^{*}$*:*

$$\eta\left(\theta^{t}\right)=\frac{\tau_{Lt-1}^{*}\left(\theta^{t-1}\right)}{1-\tau_{L,t-1}^{*}\left(\theta^{t-1}\right)}\left[\frac{R\beta}{u_{t-1}'\left(c\left(\theta^{t-1}\right)\right)}\frac{\varepsilon_{t-1}^{c}}{1+\varepsilon_{t-1}^{u}}\frac{\theta_{t-1}}{\varepsilon_{w\theta,t-1}}\int_{\theta_{t}}^{\bar{\theta}}\frac{\partial f\left(\theta_{s}|\theta_{t-1}\right)}{\partial\theta_{t-1}}d\theta_{s}\right]$$

and

$$\kappa\left(\theta^{t}\right)=\int_{\theta_{t}}^{\bar{\theta}}\frac{1}{u'\left(c\left(\theta\right)\right)}\left(1-u'\left(c\left(\theta\right)\right)\int_{\underline{\theta}}^{\bar{\theta}}\frac{1}{u'\left(c\left(m\right)\right)}f\left(m|\theta_{-}\right)dm\right)f^{t}\left(\theta|\theta_{-}\right)$$

ii) $\tau_{Lt}^{*}\left(\theta^{t-1},\bar{\theta}\right)=\tau_{Lt}^{*}\left(\theta^{t-1},\underline{\theta}\right)=0,\ \forall t.$

Hence, the only effect of training is to introduce the wage elasticity with respect to ability $\varepsilon_{w\theta,t}$, which is usually constant and equal to one in a model without training (in which $w_t = \theta_t$). With training, a higher elasticity amplifies the labor wedge as it increases the value of insurance and redistribution. However, it does not directly matter whether there is learning-or-doing or learning-and-doing. The standard zero distortion at the bottom and the top results from the static Mirrlees model continue to apply.

The multiplier $\mu\left(\theta^t\right)$ on the envelope condition captures two effects that amplify or dampen the labor wedge over the lifecycle. Because the training wedge in (26) was derived in relation to the labor wedge, they also influence the magnitude of the net training wedge. The first effect is an insurance motive, captured by $\kappa\left(\theta^t\right)$, familiar from the static taxation literature. It would be zero with linear utility. The term $\eta\left(\theta^t\right)$ captures the previous period's labor wegde, hence indirectly the previous period's insurance motive, weighted by a measure of ability persistence. In the limit, if θ_t is identically and independently distributed (iid), only the contemporaneous insurance motive $\kappa\left(\theta_t\right)$ matters.

6.2 The limit case of Learning-by-Doing

A limit case of learning-and-doing would be the learning-by-doing model introduced by Arrow (1962), which posits that agents acquire human capital through work itself. With asymmetric information, this limit case actually leads to very different implications for the labor wedge. This is because human capital with learning-by-doing is a direct by-product of labor, but labor is unobserved by the government. There is no separate instrument for training and all the incentives have to be set indirectly through the labor wedge. Indeed, this case is much closer to a case with unobservable training. As soon as training is a direct by-product of labor, either the government will be able to perfectly infer labor effort from observing training (if we maintain the assumption that training is observable) or it has to be that training like labor is unobservable.

Learning-by-doing can be modelled by letting the investment in training be the direct result of labor, not of a separate effort, so that training time as previously defined, i_t is equal to $i_t = i\left(l_t\right)$ where $i\left(\right)$ is an increasing, concave function of l_t. The stock of human capital z_t is $z_t = \sum_{\tau=1}^{t} i\left(l_\tau\right)$. For simplicity, I will consider a linear function with $i_t = l_t$, but the more general formulation gives similar results.

The next proposition highlights how the optimal labor wedge needs to be modified in order to take into account the presence of learning-by-doing.

Proposition 4 *With learning-by-doing, such that $i_t = l_t$ and $z_t = \sum_{\tau=1}^{t} l_\tau$:*

i) The optimal labor wedge is given by:

$$\frac{\tau_{Lt}}{(1-\tau_{Lt})}\left(1+\frac{w_{z,t}}{\tau_{Lt}w_t}l_t\right) = \frac{\mu(\theta_t)}{f(\theta_t|\theta_{t-1})}u'(c_t)\frac{\varepsilon_{w\theta}}{\theta_t}\left[\frac{1+\varepsilon_t^u}{\varepsilon_t^c}-\frac{w_{zt}}{w_t}l_t\left[1-\rho_{\theta z}\right]\right]$$
$$+\frac{1}{R}E\left(\left[\frac{w_{t+1}}{w_t}\frac{\tau_{L,t+1}}{(1-\tau_{Lt})}-\frac{(1-\tau_{L,t+1})}{(1-\tau_{Lt})}\frac{\mu(\theta_{t+1})}{f(\theta_{t+1}|\theta_t)}\frac{\varepsilon_{w\theta,t+1}}{\theta_{t+1}}\frac{1+\varepsilon_{t+1}^u}{\varepsilon_{t+1}^c}u'(c_{t+1})\right]\right)$$

ii) Except in the last period $t = T$, the zero tax results at the bottom and the top no longer hold: $\tau_{Lt}^*\left(\theta^{t-1},\bar{\theta}\right) \neq 0$ and $\tau_{Lt}^*\left(\theta^{t-1},\underline{\theta}\right) \neq 0$.

Three elements are noteworthy in this formula relative to (29). First, the relevant efficiency cost of taxation is no longer just the own-price elasticity, as captured by $\frac{1+\varepsilon_t^u}{\varepsilon_t^c}$. Instead, it also includes the incentive effect of training time, as captured by the Hicksian coefficient of complementarity $\rho_{\theta z}$. All else equal, a lower Hicksian coefficient of complementarity of training and ability tends to decrease the labor wedge because it is attractive to encourage training acquisition (which corresponds to supplying more labor when there is learning-by-doing). The reason for this was explained in Section 3: when the Hicksian coefficient of complementarity is below 1, training has a redistributive and insurance effect.

Secondly, the optimal labor wedge, which is normally an intratemporal wedge, now becomes forward-looking and dynamic. Labor supply now has an intertemporal effect and lasting impacts because it indirectly builds human capital. There is a need to take into account the next period's incentive effects of having a higher human capital stock and the next period's tax.

Third, and related to the dynamic implications of the labor wedge, the zero tax result at the bottom and top of the ability distribution no longer holds (except in the last period, which is akin to a static model). The intertemporal effects of the labor wedge lead to different incentive implications.

6.3 Higher ability for working and training

There is another natural modification of the model which also has different implications for the labor wedge. Suppose that higher ability agents are not only more productive at producing output, but are also more productive at training, i.e., they have a lower marginal cost for any training achieved.

In this case, the per-period utility of an agent of current type θ_t takes the form:

$$u(c_t)-\phi_t\left(\frac{y_t}{w(\theta_t,z_t)},\frac{z_t-z_{t-1}}{\theta_t}\right) \tag{30}$$

The envelope condition needs to be modified accordingly (see the Appendix) and is now given by:

$$\dot{\omega}(\theta)=\frac{w_{\theta t}}{w_t}l(\theta)\phi_{lt}+\frac{z(\theta)-z_-}{\theta^2}\phi_{it}+\beta\Delta(\theta)$$

where ϕ_{it} is the derivative of ϕ with respect to its full second argument.

Solving the problem in exactly the same way as before, with the augmented envelope condition, yields a modified formula for the optimal labor wedge.

Proposition 5 *If higher ability types are also more able at acquiring training, according to ϕ as in (30), the optimal labor wedge takes the form:*

$$\frac{\tau_L(\theta)}{1-\tau_L(\theta)} = \frac{\mu(\theta)}{f(\theta|\theta_-)} u'(c_t) \left[\frac{w_{\theta t}}{w_t} \left(\frac{1+\varepsilon^u}{\varepsilon^c} \right) + \left(\frac{z_t - z_{t-1}}{z_t} \right) \rho^{\phi}_{lzt} \varepsilon_{\phi z} \right] \tag{31}$$

where $\mu(\theta)$ still takes the same form as in proposition 3.

Note that the first term on the right-hand side of (31) is the same as the optimal tax in formula (29). The second term tells us that, when agents who have high ability for work also have high ability at training, the labor tax needs to be adjusted. In particular, whenever there is learning-or-doing ($\rho^{\phi}_{lz} > 0$), the labor wedge is adjusted upwards. The intuition for this result is that, if labor and training are substitutes, then discouraging labor encourages training and relaxes the incentive compatibility constraint. The opposite happens when there is learning-and-doing ($\rho^{\phi}_{lz} < 0$).

This additional effect can be first-order. Typically, whenever the horizon is short the labor wedges are small because the compensated elasticity in the data is small. However, the crowding out or crowding in of training could be a much stronger consideration.

7 A Simple Numerical Illustration

In this section, I illustrate numerically some of the theoretical results from Sections 5 and 6. The goal is not to provide a full-fledged quantitative analysis of the results, but rather to point out how the interaction between labor and training can influence optimal policies.[20]

7.1 Calibration

Functional specifications: Agents work for 20 periods, corresponding to 40 years in the data and are retired for 10 periods, corresponding to 20 years in the data.

The focus is on the complementarity between training and labor in the disutility function, i.e., on ρ^{ϕ}_{lz}. Therefore, the rest of the specifications are intentionally kept very simple. Ability follows a log-normal iid process, with:

$$\log \theta_t = \psi_t$$

[20] An estimation of the model is complicated and might maybe not be that informative. Indeed, while this life cycle model is more general than other models of dynamic taxation, the quest for analytical results limits the complexity we can introduce into the model. Hence, many features of the real-world are highly stylized and simplified. A separate analysis could give up on theoretical and analytical results and instead perform a more rigorous quantitative analysis.

with $\psi_t \sim N\left(-\frac{1}{2}\sigma_\psi^2, \sigma_\psi^2\right)$. The wage function is also taken to be simply additively separable, with $\rho_{\theta z} = 0$, namely:

$$w_t(\theta_t, z_t) = \theta_t + z_t \tag{32}$$

The focus is hence on the preferences. The first simulation considers the baseline separable preferences between training and labor, with $\rho_{lz}^{\phi} = 0$:

$$\tilde{u}(c_t, y_t, z_t, \theta_t) = \log(c_t) - \frac{\kappa}{\gamma}\left(\frac{y_t}{w_t(\theta_t, z_t)}\right)^{\gamma} - \frac{1}{\alpha}(z_t - z_{t-1})^{\alpha}, \kappa > 0, \gamma > 1, \alpha > 1 \tag{33}$$

The second simulation considers a specification with learning-or-doing in which training and labor are substitutes:

$$\tilde{u}(c_t, y_t, z_t, \theta_t) = \log(c_t) - \frac{\kappa}{\gamma\alpha}\left(\frac{y_t}{w_t(\theta_t, z_t)}\right)^{\gamma}(z_t - z_{t-1})^{\alpha}, \kappa > 0, \gamma > 1, \alpha > 1 \tag{34}$$

In this case, $\rho_{lz}^{\phi} = 1$.

During retirement, utility is in both cases simply $\tilde{u}^R(c_t) = \log(c_t)$.

Parameter values: The parameter values are calibrated exogenously from the existing literature. I set $\gamma = 3$ and $\kappa = 1$, in order to obtain a Frisch elasticity of 0.5 as in Chetty (2012). Given the lack of evidence on the elasticity of training to the net return, I select $\alpha = 2$. The discount factor is set to $\beta = 0.95$, and the gross borrowing rate to $R = \frac{1}{\beta} = 1.053$. The variance of ability is taken to be, to a first order, the variance of productivity, for which there are many different estimates in the literature. I take a medium-range value of 0.0095 (Heathcote *et al.*, 2005). [21]

Simulations: After using the recursive formulation of the problem to compute the policy functions, I simulate 100,000 draws through a Monte Carlo simulation for each of the two preference specifications (i.e., for the values $\rho_{lz}^{\phi} = 0$ and $\rho_{lz}^{\phi} = 1$). The initial states are set to yield a zero present value resource cost for the allocation. This ensures comparability across simulations, and gives a sense of outcomes achievable without outside government revenue.

7.2 Discussion of the numerical illustration

Figure 1 shows the simulation results. Each panel depicts the cross-sectional average (taken across all 100,000 simulations) of the corresponding wedge at the given time period. It is most natural to interpret time $t = 0$ as approximately age 20 –when the agent is adult and makes labor and training choices on his own – and time $t = 20$ as the retirement age, or age 60.

[21] It naturally is very challenging to assess what is uncertainty versus what is heterogeneity from the point of view of the agent (as opposed to the point of view of the econometrician). Ideally, we would like estimates of the initial hetereogeneity, and of the transition function between different ability states, which is also uncertain for the agent.

As is usual, the labor wedge in panel (a) grows over the life cycle. Because the ability shock is iid over time. the labor wedge is relatively small, even at age 60. More persistent or more volatile shocks would make the labor wedge higher and steeper over the life cycle.[22] The gross training wedge in panel (b) closely tracks the labor wedge. The training wedge τ_{Zt} is lower and grows slower when there is learning-or-doing. The intuition for this is that when there is learning-or-doing ($\rho_{lz}^{\phi} = 1$), encouraging training discourages labor and is less attractive than when there is no learning-or-doing ($\rho_{lz}^{\phi} = 0$).

Figure 1: Average Wedges over the Life Cycle

(a) Labor Wedge τ_{Lt}

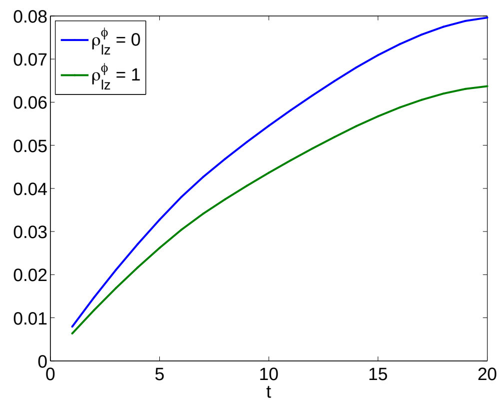

(b) Gross Wedge τ_{Zt}

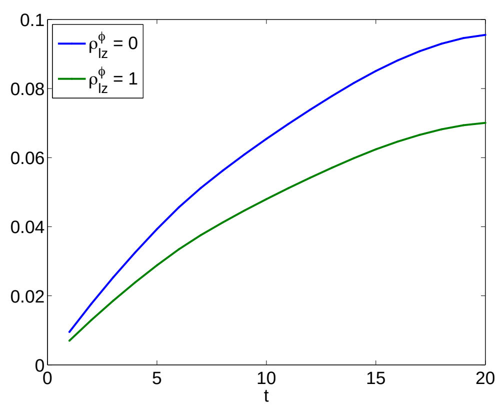

(c) Net Training Wedge t_{zt}

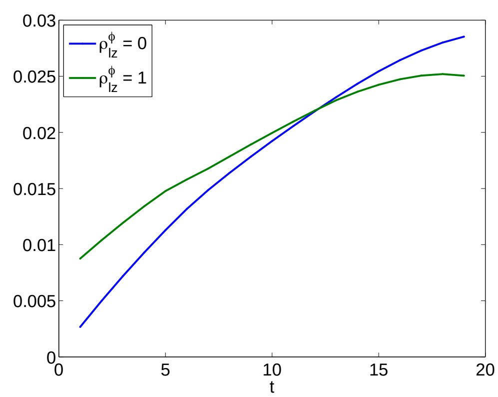

Each of the three panels depicts the average wedge at the given time period across the 100,000 simulations. Panel (a) depicts the labor wedge τ_{LT}, panel (b) depicts the gross training wedge, τ_{Zt}, and panel (c) depicts the net training wedge t_{zt}. The blue line represents the benchmark separable case with preferences as in (33) and $\rho_{lz}^{\phi} = 0$. The green line represents the case with learning-or-doing and nonseparable preferences as in (34) for which $\rho_{lz}^{\phi} = 1$. Each time period in the simulations represents two years in the data. It is most natural to view time $t = 0$ as approximately age 20 and time $t = 20$ as the retirement age, or age 60.

Given the close parallel between the labor wedge and the gross wedge, the net wedge in panel (c) is very small. Since training is a dynamic decision, it is very costly to distort it for redistributive and incentive reasons. Hence, the optimal allocation remains very close to neutrality with respect to

[22]As compared, for instance, to the numbers in Farhi and Werning (2013) and Stantcheva (2014).

training. The net wedge is less steep when there is learning-or-doing. In that case, it is especially attractive to encourage human capital relatively more in the earlier periods than in later periods. Inutitively, training will always distract effort away from labor when there is learning-or-doing, but the benefits will be recouped over more periods when the agent is younger.

8 Conclusion

This paper considered optimal taxation and training policies when agents allocate their time between working and acquiring human capital. Time investments in human capital are peculiar because they cannot be dissociated from the agent: their returns interact with an agent's intrinsic and stochastic type, and their costs interact with an agent's labor effort. The model allows for a general wage function and a general interaction between labor and training that ranges from learning-and-doing (training and labor are complements) to learning-or-doing (training and labor are substitutes), and which can be different at different ages in life.

I first considered a simple one-period model to highlight the main intuitions related to human capital acquired though training. The notion of a net wedge was introduced, as a measure of the net incentive that the government wants to provide for training, above and beyond simply compensating for the labor distortion. At the optimum, the ratio of the net training wedge to the labor wedge needs to satisfy a simple relation that depends on two parameters: negatively on the Hicksian coefficient of complementarity between training and ability in the wage function, which measures how much more high ability agents benefit from training, and negatively on the Hicksian coefficient of complementarity between training and labor in the disutility function. The latter is positive if there is "learning-or-doing" and negative in the presence of "learning-and-doing." The net wedge is positive if and only if the total labor supply effect outweights the inequality effect. Some special cases yield particularly simple, constant ratios between the labor wedge and training wedge and can be used to test for the Pareto efficiency of any tax and subsidy system.

In the dynamic life cycle model, there is an additional interaction between contemporaneous training and future labor supply, which is the mirror image of the interaction with the contemporaneous labor supply. This makes the sign of the net wedge more complex and there are many possible configurations.

The labor wedge is not affected by learning-or-doing or learning-and-doing. I considered two modifications of the model which do induce important changes for the labor wedge: the limit case of "learning-by-doing," in which training is a direct by-product of labor effort, and the case in which agents who are more able at work are also more able at training.

In future work, it would be interesting to pin down empirically how the interaction between labor supply and time investments in human capital evolves over the life cycle. At what ages and at what levels of labor supply and training is learning-or-doing more prevalent than learning-and-doing, and vice-versa? Then, a more complete quantitative analysis of the optimal policies would be feasible and insightful.

9 References

Albanesi, S. and C. Sleet. 2006. "Dynamic Optimal Taxation with Private Information," *Review of Economic Studies* 73, 1-30.

Arrow, Kenneth 1962. "The Economic Implications of Learning by Doing," *The Review of Economic Studies* 29(3), 155-173.

Atkinson, A. and J. Stiglitz 1976. "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics* 6, 55-75.

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

Battaglini, M. and S. Coate. 2008b. "A Dynamic Theory of Public Spending, Taxation, and Debt," *American Economic Review* 99(1), 201-236.

Becker, G. 1964. *A Theoretical and Empirical Analysis, with Special Reference to Education*, The University of Chicago Press.

Ben-Porath, Y. 1967. "The Production of Human Capital and the Life Cycle of Earnings", *Journal of Political Economy*, Vol. 75, No. 4, pp. 352-365.

Bohacek, R. and M. Kapicka. 2008. "Optimal Human Capital Policies," *Journal of Monetary Economics* 55(1), 1-16.

Bovenberg, L. and B. Jacobs. 2005. "Redistribution and Education Subsidies are Siamese Twins," *Journal of Public Economics*, 89(11-12), 2005–2035.

Bovenberg, L. and B. Jacobs. 2011. "Optimal Taxation of Human Capital and the Earnings Function," *Journal of Public Economic Theory*, 13, (6), 957-971.

Da Costa, C. E., and L. J. Maestri. 2007. "The risk Properties of Human Capital and the Design of Government Policies," *European Economic Review* 51(3), 695-713.

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

Findeisen, S. and D. Sachs 2012. "Education and Optimal Dynamic Taxation: The Role of Income-contingent Student Loans," *Working Paper, University of Konstanz.*

Gelber, A. and M. Weinzierl 2012. "Equalizing Outcomes and Equalizing Opportunities: Optimal Taxation When Children's Abilities Depend on Parents' Resources," *NBER Working Paper, 18332.*

Golosov, M, A. Tsyvinski and M. Troshkin. 2013. "Redistribution and Social Insurance," *Working paper, Princeton University.*

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

Golosov, M, and A. Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122(2), 487-534.

Golosov, Mikhail, Aleh Tsyvinski, and Nicolas Werquin 2014. "A Variational Approach to the Analysis of Tax Systems," *Working Paper, Yale.*

Heathcote, J., K. Storesletten and G. Violante. 2005. "Two Views of Inequality Over the Life Cycle," *Journal of the European Economic Association,* 3, 765-775.

Heckman, James. 1976a. "Estimates of a Human Capital Production Function Embedded in a Life-Cycle Model of Labor Supply," in N. Terleckyj (ed.), *Household Production and Consumption*, NBER, p. 225-264.

Heckman, James. 1976b. "A Life-Cycle Model of Earnings, Learning, and Consumption," *Journal of Political Economy*, 84(4), S11-S44.

Heckman, J., L. Lochner, and P. Todd. 2005. "Earnings Functions, Rates of Return, and Treatment Effects: The Mincer Equation and Beyond," *NBER Working Paper 11544.*

Hicks, J. 1970. "Elasticity of Substitution Again: Substitutes and Complements," *Oxford Economic Papers,* 25, 289-296.

Kapicka, M. and J. Neira 2014. "Optimal Taxation in a Life-Cycle Economy with Endogenous Human Capital Formation," *Working Paper, University of California, Santa Barbara.*

Kapicka, M. 2013a. "Optimal Mirrleesean Taxation with Unobservable Human Capital Formation," *Working Paper, University of California, Santa Barbara.*

Kapicka, M. 2013b. "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies,* 80(3), 1027-1054.

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

Krueger, D., and A. Ludwig. 2013. "Optimal Capital and Progressive Labor Income Taxation with Endogenous Schooling Decisions and Intergenerational Transfers," *University of Pennsylvania Working Paper.*

Maldonado, D. 2007. "Education Policies and Optimal Taxation," *International Tax and Public Finance*, 15(2), 131-143.

Milgrom, R. 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.

Williams, Timothy 2014. "Seeking New Start, Finding Steep Cost Workforce Investment Act Leaves

Many Jobless and in Debt," *New York Times* , August 17, 2014.

OECD. 2004. "Improving Skills for More and Better Jobs: Does Training Make a Difference?," *OECD Employment Outlook.*

OECD. 2013. "Education at a Glance, 2013," *OECD, Paris.*

Rothschild, Casey and Florian Scheuer. 2014a. "A Theory of Income Taxation under Multidimensional Skill Heterogeneity," *NBER Working Paper* 19822.

Rothschild, Casey and Florian Scheuer. 2014b. "Optimal Taxation with Rent-Seeking," *NBER Working Paper* 17035.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies* 68(1), 205-229

Samuelson, P. 1973. "Complementarity – An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory", *Journal of Economic Literature,* 12, (4), 1255-1289.

Shourideh, Ali 2012. "Optimal Taxation of Wealthy Individuals," *Working Paper University of Pennsylvania.*

Shourideh, Ali, and Maxim Troshkin 2012. "Providing Efficient Incentives to Work: Retirement Ages and the Pension System," *Working Paper University of Pennsylvania.*

Scheuer, Florian 2014. "Entrepreneurial Taxation with Endogenous Entry," *American Economic Journal: Economic Policy,* 6, 126-163.

Stantcheva, Stefanie 2014. "Optimal Taxation and Human Capital Policies over the Life Cycle," *Working Paper, Harvard University.*

Weinzierl, M. 2011. "The Surprising Power of Age-Dependent Taxes," *Review of Economic Studies,* 78, 1490-1518.

Werning, Ivan 2002 "Optimal Unemployment Insurance with Unobservable Savings," *Working Paper MIT.*

Werning, Ivan 2007a "Nonlinear Capital Taxation," *Working Paper MIT.*

Werning, Ivan 2007b "Optimal Fiscal Policy with Redistribution," Quarterly Journal of Economics, 122, 925-967.

Werning, Ivan 2007c "Pareto Efficient Income Taxation," *Working Paper MIT.*

Werquin, Nicolas 2014 "Income Taxation with Frictional Labor Supply," *Working Paper, Yale.*

A Appendix

Notation:

$w_{\theta,t}$, $w_{z,t}$ are the partials of the wage with respect to ability, human capital z respectively. $w_{\theta z,t}$, $w_{zz,t}$ are the second order derivatives. $\phi_{l,t}$ and $\phi_{z,t}$ are the partials of the disutility with respect to labor and training respectively, and $\phi_{zz,t}, \phi_{lz,t}, \phi_{ll,t}$ the corresponding second order derivatives. ε_{xyt} is the elasticity of variable x with respect to variable y, $\varepsilon_{xyt} = d\log(x_t)/d\log(y_t)$, hence for instance, $\varepsilon_{w\theta,t}$ is the elasticity of the wage with respect to ability θ. When clear, the history dependence of allocations is omitted, e.g., c_t denotes $c(\theta^t)$.

Link to the Ben-Porath model in discrete time.

In the Ben-Porath model, the general accumulation process for human capital is $z_{t+1} = H(z_t, \alpha_t) + (1-\delta) z_t$ with α_t the fraction of human capital reinvested into human capital acquisition and δ the depreciation rate of human capital. Labor supply is equal to $l_t = 1 - \alpha_t$, the residual time from human capital acquisition. The rental rate of human capital w_t is constant and earnings are $y_t = w_t z_t l_t$. The most general functional form used is $z_{t+1} = \beta_0 \alpha_t^{\beta_1} (\alpha_t z_t)^{\beta_2} + (1-\delta) z_t$, so that $w_t \alpha_t z_t$ is the opportunity cost of human capital acquisition.

In this paper, a variation of this model is used, which is better adapted to a Mirrleesian analysis. The accumulation process is simplified along some dimensions and made more complex along others. First, instead of α_t, the input into human capital is effort or time i_t. i_t plays the same role of α_t but has a disutility cost, potentially nonseparable from the disutility cost of labor $\phi_t(l, i)$, and not just an opportunity cost of time. This is because, first, disutility costs from learning are empirically relevant (see Heckman et al. 2005) and second, with observable human capital, labor would also be observable if we only had: $l_t = 1 - i_t$. Second, the parameters are fixed so that only time, but not the previous stock of human capital matters for the accumulation, and there is no depreciation: $\beta_1 = 1 = \beta_0$ and $\beta_2 = \delta = 0$. This leads to a linear accumulation process $z_{t+1} = z_t + i_t$. The diminishing returns to human capital are instead captured in the more complex rental rate for human capital, i.e., the wage, which is heterogeneous, nonlinear, and stochastic $w_t = w_t(\theta_t, z_t)$.

Generalized model a la Ben-Porath: It is possible to generalize the model to feature an accumulation process close to the one in Ben-Porath. The key change is to make the accumulation process depend on the human capital stock. Hence, let $z_t = Z_t(z_{t-1}, i_t) = z_{t-1} + H_t(z_{t-1}, i_t)$ be a more general production function without depreciation (which is without loss of generality). Define the inverse function with respect to i, $I_t(z_t, z_{t-1}) \equiv Z_t^{-1,i}(z_{t-1}, z_t)$. The accumulation process can be time-dependent. Note now that with $\frac{\partial \phi_t}{\partial z_t} = \frac{\partial \phi_t}{\partial i_t}\frac{\partial i_t}{\partial z_t} = \frac{\partial \phi_t}{\partial i_t}\frac{\partial I_t}{\partial z_t}$, the definition for the net bonus t_{zt} is unchanged. The same formula for the optimal bonus on training time as in (26) applies.

Derivations and proofs for Propositions (2) and (3) **and Corollary** (2):

The expenditure function: $\tilde{c}(l, \omega - \beta v, i, \theta)$ defines consumption indirectly as a function of labor l, current period utility ($\tilde{u} = \omega - \beta v$), training, and the current realization of the type. Then, $\omega(\theta) = u_t(c(\theta)) - \phi_t(l(\theta), z(\theta) - z_-) + \beta v(\theta)$ becomes redundant as a constraint, and the choice variables are $(l(\theta), z(\theta), \omega(\theta), v(\theta), \Delta(\theta))$. Let the multipliers in program (21) be (in the order of the constraints there) $\mu(\theta)$, λ_-, and γ_-. The problem is solved using the optimal control approach where the "types" play the role of the running variable, $\omega(\theta)$ is the state (and $\dot{\omega}(\theta)$ its law of motion), and the controls are $l(\theta), v(\theta)$ and $\Delta(\theta)$. The Hamiltonian is:

$$\begin{aligned}
&(\tilde{c}(l(\theta), \omega(\theta) - \beta v(\theta), z(\theta) - z_-, \theta) - w_t(\theta, z(\theta))\, l(\theta))\, f^t(\theta|\theta_-) \\
&+\frac{1}{R} K(v(\theta), \Delta(\theta), \theta, z(\theta), t+1)\, f^t(\theta|\theta_-) \\
&+\lambda_- \left[v - \omega(\theta) f^t(\theta|\theta_-)\right] + \gamma_- \left[\Delta - \omega(\theta) \frac{\partial f^t(\theta|\theta_-)}{\partial \theta_-}\right] + \mu(\theta) \left[\frac{w_{\theta,t}}{w_t} l(\theta)\, \phi_{l,t}(l(\theta), z(\theta) - z_-) + \beta \Delta(\theta)\right]
\end{aligned}$$

with boundary conditions:

$$\lim_{\theta \to \bar{\theta}} \mu(\theta) = \lim_{\theta \to \underline{\theta}} \mu(\theta) = 0$$

First, we solve for the optimal labor wedge, which is then used in the formula for the optimal training wedge.

Taking the first order condition of the recursive planning problem with respect to the labor choice:

$$[l(\theta)] : \frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \frac{w_{\theta,t}}{w_t} u_t'(c(\theta)) \left(1 + \frac{l(\theta)\, \phi_{ll,t}(l(\theta), i(\theta))}{\phi_{l,t}(l(\theta), i(\theta))}\right)$$

using the definitions of ε^c, ε^u and ε in the text:

$$\frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \frac{\varepsilon_{w\theta}}{\theta} u_t'(c(\theta)) \frac{1 + \varepsilon^u}{\varepsilon^c}$$

Taking the first order condition of the Hamiltonian with respect to z_t yields:

$$\begin{aligned}
[z_t] : &\left(-\frac{\phi_{z,t}}{u'(c_t)} + w_{z,t} l_t + \int \frac{1}{R} \frac{\phi_{z,t+1}}{u'(c_{t+1})} f^{t+1}(\theta_{t+1}|\theta_t) + \frac{1}{R} E_t \left(\frac{\mu(\theta_{t+1})}{f^{t+1}(\theta_{t+1}|\theta_t)} \frac{w_{\theta,t+1}}{w_{t+1}} l_{t+1} \phi_{lz,t+1}\right)\right) \\
&-\frac{\mu(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \left(l_t \phi_{l,t}(l_t, z_t - z_{t-1}) \frac{1}{w_t} w_{\theta z,t} - l_t \phi_{l,t} w_{\theta,t} w_{z,t} \frac{1}{w_t^2} + \frac{w_{\theta,t}}{w_t} l_t \phi_{lz,t}\right) = 0
\end{aligned}$$

Using the definition of the training time bonus, τ_{Zt} to replace for $w_{z,t} l_t$ yields:

$$-\frac{\phi_{z,t}}{u'(c_t)} + w_{z,t} l_t + \int \frac{1}{R} \frac{\phi_{z,t+1}}{u'(c_{t+1})} f^{t+1}(\theta_{t+1}|\theta_t) = -\frac{1}{(1 - \tau_{Lt})} \left(\tau_{Zt} - \tau_{Lt} \left(\frac{\phi_z}{u'(c_t)}\right)^d + P_{Zt}\right)$$

with P_{Zt} and $\left(\frac{\phi_z}{u'(c_t)}\right)^d$ as defined in the text. The FOC then becomes:

$$\begin{aligned}
[z_t] \quad : \quad t_{zt} = \frac{\left(\tau_{zt} - \tau_{Lt} \left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d + P_{Zt}\right)}{(1 - \tau_{Lt}) \left[\left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d - \tau_{zt}\right]} &= \frac{\mu(\theta_t)}{f^t(\theta_t|\theta_{t-1})} u'(c_t) \frac{\varepsilon_{w\theta,t}}{\theta_t} (1 - \rho_{\theta z,t}) \\
&- \frac{1}{\left[\left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d - \tau_{zt}\right]} \left(\frac{\mu(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \frac{\varepsilon_{w\theta,t}}{\theta_t} l_t \phi_{lz,t} - \frac{1}{R} E_t \left(\frac{\mu(\theta_{t+1})}{f^{t+1}(\theta_{t+1}|\theta_t)} \frac{\varepsilon_{w\theta,t+1}}{w_{t+1}} l_{t+1} \phi_{lz,t+1}\right)\right)
\end{aligned}$$

We can again replace the multipliers by the labor wedges, and note that

$$\frac{w_t}{w_{z,t} \phi_{l,t}(l_t)} \phi_{lz,t} = \frac{\varepsilon_{\phi z,t}}{\varepsilon_{wz,t}} \rho_{lz,t}^{\phi}$$

to obtain formula (26). In this case the elasticities are defined as $\varepsilon_{\phi z} = d\log(\phi)/d\log(z)$ and $\varepsilon_{wz} = d\log(w)/d\log(z)$.

The law of motion for the co-state $\mu(\theta)$ comes from the first-order condition with respect to the state variable $\omega(\theta)$:

$$[\omega(\theta)] : \left(-\frac{1}{u_t'(c(\theta))} + (\lambda_-) + (\gamma_-)\frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-}\frac{1}{f^t(\theta|\theta_-)}\right) f^t(\theta|\theta_-) = \dot{\mu}(\theta) \tag{35}$$

Integrating this and using the boundary condition $\mu\left(\bar{\theta}\right) = 0$, yields:

$$\mu(\theta) = \int_\theta^{\bar{\theta}} \left(\frac{1}{u_t'(c(\theta))} - (\lambda_-) - (\gamma_-)\frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-}\frac{1}{f^t(\theta|\theta_-)}\right) f^t(\theta|\theta_-) \tag{36}$$

Integrating and using both boundary conditions yields:

$$\lambda_- = \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u'(c(\theta))} f^t(\theta|\theta_-)\, d\theta \tag{37}$$

Using the envelope conditions $\frac{\partial K(v(\theta),\Delta(\theta),\theta,s(\theta),z(\theta),t+1)}{\partial v(\theta)} = \lambda(\theta)$ and $\frac{\partial K(v(\theta),\Delta(\theta),\theta,s(\theta),z(\theta),t+1)}{\partial \Delta(\theta)} = -\gamma(\theta)$, the first-order conditions with respect to $v(\theta)$ and $\Delta(\theta)$ respectively lead to:

$$[v(\theta)] : \frac{1}{u'(c)} = \frac{\lambda(\theta)}{R\beta} \tag{38}$$

and

$$[\Delta(\theta)] : -\frac{\gamma(\theta)}{R\beta} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \tag{39}$$

Using (37) and (39) in the expression for $\mu(\theta)$ from (36) yields: $\mu\left(\theta^t\right) = \kappa\left(\theta^t\right) + \eta\left(\theta^t\right)$ where

$$\kappa\left(\theta^t\right) = \int_{\theta_t}^{\bar{\theta}} \frac{1}{u'(c(\theta))}\left(1 - u'(c(\theta))\int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u'(c(m))} f(m|\theta_-)\, dm\right) f^t(\theta|\theta_-)$$

$$\eta\left(\theta^t\right) = -(\gamma_-)\int_{\theta_t}^{\bar{\theta}} \frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-} d\theta = \frac{\mu\left(\theta^{t-1}\right)}{f(\theta_{t-1}|\theta_{t-2})} R\beta \int_{\theta_t}^{\bar{\theta}} \frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-} d\theta$$

where the last equality uses the lag of (39).

Lemma 1 *Under assumption (2),* $\mu\left(\theta^t\right) \geq 0\ \forall t,\ \forall\theta^t$.

Proof of Lemma 1:

The proof is close to the one in Golosov, Tsyvinski and Troshkin (2011), for a separable utility function and with human capital. From the envelope condition and the FOC for $v(\theta)$ in (38):

$$\frac{\partial K}{\partial v} = \lambda(\theta) = \frac{R\beta}{u_t'(c(\theta))}$$

Since by assumption $v(\theta)$ is increasing in θ and $K()$ is increasing and convex in v, it must be that $\frac{\partial K}{\partial v}$ is increasing in θ, so that $\frac{1}{u_t'(c(\theta))}$ as well is increasing in θ.

Start in period $t = 1$. In this case, since θ_0 has a degenerate distribution, $\frac{\partial f^1}{\partial \theta_0}(\theta_1|\theta_0) = 0$ and

$$\mu(\theta_1) = \int_{\theta_1}^{\bar{\theta}} \left(\frac{1}{u_1'\left(c\left(\tilde{\theta}_1\right)\right)} - \lambda_- \right) f^1\left(\tilde{\theta}_1\right) d\tilde{\theta}_1$$

Choose the θ' such that $\frac{1}{u'(c(\theta'))} = \lambda_-$. Since $\frac{1}{u'(c(\theta))}$ is increasing in θ, for $\theta \geq \theta'$, $\mu(\theta) \geq 0$ (integrating over non-negative numbers only). Using the boundary condition $\mu(\underline{\theta}) = 0$, $\mu(\theta_1)$ can also be rewritten as:

$$\mu(\theta_1) = \int_{\underline{\theta}}^{\theta_1} \left(-\frac{1}{u_1'\left(c\left(\tilde{\theta}_1\right)\right)} + \lambda_- \right) f^1\left(\tilde{\theta}_1\right) d\tilde{\theta}_1$$

Since for $\theta_1 \leq \theta_1'$, $\frac{1}{u'(c(\theta_1))} \leq \lambda_-$, we again have $\mu(\theta_1) \geq 0$. Thus, for all θ_1, $\mu(\theta_1) \geq 0$.

By the first-order condition for Δ in (39):

$$-\frac{\gamma(\theta_1)}{R\beta} = \frac{\mu(\theta_1)}{f(\theta_1)}$$

so that $\gamma(\theta_1) \leq 0$, for all θ_1. Note that $\mu(\theta_2)$ is equal to:

$$\mu(\theta_2) = \int_{\theta_2}^{\bar{\theta}} \left(\frac{1}{u_2'\left(c\left(\tilde{\theta}_2\right)\right)} - (\lambda_-) - \frac{\partial f\left(\tilde{\theta}_2|\theta_1\right)}{\partial \theta_1} \frac{(\gamma_-)}{f\left(\tilde{\theta}_2|\theta_1\right)} \right) f\left(\tilde{\theta}_2|\theta_1\right)$$

Since by assumption (2) iii), $\frac{g\left(\tilde{\theta}_2|\theta_1\right)}{f\left(\tilde{\theta}_2|\theta_1\right)}$ is increasing in $\tilde{\theta}_2$, and we already showed that $\frac{1}{u'\left(c\left(\tilde{\theta}_2\right)\right)}$ is increasing in $\tilde{\theta}_2$, there is a θ_2' such that

$$\frac{1}{u_2'(c(\theta_2'))} - \frac{\partial f(\theta_2'|\theta_1)}{\partial \theta_1} \frac{(\gamma_-)}{f(\theta_2'|\theta_1)} = \lambda_-$$

and such that for $\theta_2 \geq \theta_2'$, $\mu\left(\theta^2\right) \geq 0$ (since integrating over non-negative numbers only). Rewriting $\mu\left(\theta^2\right)$ as an integral from $\underline{\theta}$ to θ_2 and using the boundary condition $\underline{\theta} = 0$, we can again show that $\mu(\theta_2) \geq 0$ also for $\theta_2 \leq \theta_2'$. Proceeding in the same way for all periods up to T shows the result.

If θ is iid, $\gamma_- = 0$ and $\eta\left(\theta^t\right) = 0$ for all t. In addition, if $u_t'(c_t) = 1\ \forall t$, then $\kappa\left(\theta^t\right) = 0$ as well.

From the boundary conditions $\mu(\underline{\theta}) = \mu\left(\bar{\theta}\right) = 0$.

The Inverse Euler Equation for capital:

Proposition 6 *At the optimum, the inverse Euler Equation holds*:

$$\frac{R\beta}{u_t'(c(\theta^t))} = \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u_{t+1}'(c(\theta^{t+1}))} f^{t+1}(\theta_{t+1}|\theta_t)\, d\theta_{t+1} \tag{40}$$

Proof of Proposition (6):

Taking integral of $\dot{\mu}(\theta)$ in equation (35) between the two boundaries, $\bar{\theta}$ and $\underline{\theta}$, and using the boundary conditions $\mu\left(\bar{\theta}\right) = \mu(\underline{\theta}) = 0$, as well as the expression for λ_- from (38), lagged by one period, yields the inverse Euler equation in (40).

Proof of Proposition (4):

In the learning-by-doing case, the training stock is a result of past labor supply:

$$\sum_t l_t = z_t$$

and

$$l_t = z_t - z_{t-1}.$$

The labor wedge is appropriately redefined here as:

$$(1-\tau_{Lt})\, w_t \left[1 + \frac{w_{zt}}{w_t} l_t\right] = \beta E\left[\left(\frac{u'(c_{t+1})}{u'(c_t)} w_{t+1} - \frac{\phi'(l_{t+1})}{u'(c_t)}\right)\right] + \frac{\phi'(l_t)}{u'(c_t)}$$

Taking the first-order condition of the planning problem with respect to z_t:

$$\begin{aligned}
[z(\theta)] \quad &: \quad -w_t\left(1 + \frac{l}{i}\varepsilon_{wi}\right) + \frac{\phi'(l)}{u'(c)} - \frac{1}{R}E\left(\frac{\phi'(l_{t+1})}{u'(c_{t+1})} - w_{t+1}\right) \\
&= \quad -\frac{\mu(\theta)}{f(\theta_t|\theta_{t-1})} l\phi'(l)\left[\frac{\partial^2 w_t}{\partial\theta_t \partial z_t}\frac{1}{w_t} - \frac{\partial w_t}{\partial \theta_t}\frac{\partial w_t}{\partial z_t}\frac{1}{w_t^2}\right] - \frac{\mu(\theta_t)}{f(\theta_t|\theta_{t-1})}\frac{w_{\theta t}}{w_t}\phi'(l)\frac{1}{\varepsilon_t} \\
&+ \quad \frac{1}{R}E\left(\frac{\mu(\theta_{t+1})}{f(\theta_{t+1}|\theta_t)}\frac{w_{\theta t+1}}{w_{t+1}}\phi'(l_{t+1})\frac{1}{\varepsilon_{t+1}}\right)
\end{aligned}$$

$$\begin{aligned}
[z(\theta)] \quad &: \quad \frac{\tau_{Lt}}{(1-\tau_{Lt})}\left[\frac{\phi'(l_t)}{u'(c_t)} + \frac{1}{R}E\left(w_{t+1} - \frac{\phi'(l_{t+1})}{u'(c_{t+1})}\right)\right] \\
&= \quad \frac{\mu(\theta)}{f(\theta_t|\theta_{t-1})}\phi'(l)\frac{\varepsilon_{w\theta,t}}{\theta_t}\left[l\frac{w_{z,t}}{w_t}\left[\rho_{\theta z,t} - 1\right] + \frac{1}{\varepsilon_t}\right] - \frac{1}{R}E\left(\frac{\mu(\theta_{t+1})}{f(\theta_{t+1}|\theta_t)}\phi'(l_{t+1})\frac{\varepsilon_{w\theta,t+1}}{\theta_{t+1}}\frac{1}{\varepsilon_{t+1}}\right) \\
&- \quad \frac{1}{(1-\tau_{Lt})}E\left[\left(\beta\frac{u'(c_{t+1})}{u'(c_t)} - \frac{1}{R}\right)\left(w_{t+1} - \frac{\phi'(l_{t+1})}{u'(c_{t+1})}\right)\right]
\end{aligned}$$

Proof of Proposition (5):

Suppose that

$$\phi_t\left(\frac{y_t}{w_t}, \frac{z_t - z_{t-1}}{\theta_t}\right)$$

The envelope condition is:

$$\dot{\omega}(\theta) = \frac{w_{\theta t}}{w_t} l_t \phi_{lt}\left(l_t, \frac{z_t - z_{t-1}}{\theta_t}\right) + \frac{z_t - z_{t-1}}{\theta_t^2}\phi_{it}\left(l_t, \frac{z_t - z_{t-1}}{\theta_t}\right)$$

Where ϕ_{it} denotes the derivative with respect to its second argument and we have the relation:

$$\phi_{zt} = \frac{1}{\theta}\phi_{it}$$

where ϕ_{zt} is the usual derivative with respect to z_t from the text.

Again, let the expenditure function: $\tilde{c}(l, \omega - \beta v, z, z_{-}, \theta)$ defines consumption indirectly as a function of labor l, current period utility ($\tilde{u} = \omega - \beta v$), training, and the current realization of the type. The Hamiltonian is:

$$
\begin{aligned}
& \left(\tilde{c}\left(l\left(\theta\right), \omega\left(\theta\right)-\beta v\left(\theta\right), z\left(\theta\right)-z_{-}, \theta\right)-w_{t}\left(\theta, z\left(\theta\right)\right) l\left(\theta\right)\right) f^{t}\left(\theta \mid \theta_{-}\right) \\
& +\frac{1}{R} K\left(v\left(\theta\right), \Delta\left(\theta\right), \theta, z\left(\theta\right), t+1\right) f^{t}\left(\theta \mid \theta_{-}\right) \\
& +\lambda_{-}\left[v-\omega\left(\theta\right) f^{t}\left(\theta \mid \theta_{-}\right)\right]+\gamma_{-}\left[\Delta-\omega\left(\theta\right) \frac{\partial f^{t}\left(\theta \mid \theta_{-}\right)}{\partial \theta_{-}}\right] \\
+ \quad & \mu\left(\theta\right)\left[\frac{w_{\theta t}}{w_{t}} l(\theta) \phi_{l t}\left(l(\theta), \frac{z(\theta)-z_{-}(\theta_{-})}{\theta}\right)+\frac{z(\theta)-z_{-}(\theta_{-})}{\theta^{2}} \phi_{z t}\left(l(\theta), \frac{z(\theta)-z_{-}(\theta_{-})}{\theta}\right)+\beta \Delta\left(\theta\right)\right]
\end{aligned}
$$

Taking the FOC with respect to labor and replacing for the definition of the labor wedge:

$$
[l_t]: \frac{\tau_L(\theta)}{1-\tau_L(\theta)}=\frac{\mu\left(\theta\right)}{f\left(\theta \mid \theta_{-}\right)} u'\left(c_t\right)\left[\frac{w_{\theta t}}{w_t}\left(\frac{1+\varepsilon^u}{\varepsilon^c}\right)+\frac{i_t}{\theta_t} \frac{\phi_{lit}}{\phi_{lt}}\right]
$$

Using the fact that

$$
\frac{i_t}{\theta_t} \frac{\phi_{lit}}{\phi_{lt}}=\rho_{lit}^{\phi} \frac{\varepsilon_{\phi i}}{\theta}
$$

yields the formula in the text:

$$
[l_t]: \frac{\tau_L(\theta)}{1-\tau_L(\theta)}=\frac{\mu\left(\theta\right)}{f\left(\theta \mid \theta_{-}\right)} u'\left(c_t\right)\left[\frac{w_{\theta t}}{w_t}\left(\frac{1+\varepsilon^u}{\varepsilon^c}\right)+\rho_{lit}^{\phi} \frac{\varepsilon_{\phi i}}{\theta}\right]
$$