Journal of Economic Dynamics & Control 40 (2014) 154–178

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Vintage human capital and learning curves

Matthias Kredler*

Departamento de Economía, Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe, Spain

A R T I C L E I N F O

Article history:
Received 23 October 2012
Received in revised form
21 October 2013
Accepted 29 December 2013
Available online 10 January 2014

JEL classification:
J01
E24

Keywords:
Vintage human capital
Tenure-wage profiles
Learning curve
Infinite-dimensional state space
Lagrange-multiplier theorem

A B S T R A C T

I study a vintage-human-capital model in which long-lived workers accumulate human capital following an exogenous learning curve. Different skill levels inside a vintage are complementary in production; this makes the ex ante homogeneous workers enter different vintages. The continuous-time framework allows me to study the timing decision for the technology phase-out differentially and to derive sharp characterization for wages and the distribution of workers in the dying technology. I show how to posit and solve a planner's problem and construct equilibrium in this way. Consistent with empirical evidence, I show that the experience premium is always positive but diminishes as a technology ages. The connection between workers' learning curves and the technology's progress curve is characterized.

1. Introduction

1.1. Motivation

There is a huge empirical literature on learning-by-doing that has established that economic agents increase their productivity at a predictable pace as they find out how to best use a new technology (see the surveys by Thompson, 2010, from the micro-economic angle and by Yelle, 1979, from the business-and-management angle). This is true both for individual workers as for firms and entire industries. I will follow Thompson's (2010) excellent survey and refer to a "learning curve" as increases in productivity that an individual worker exhibits, whereas the "progress curve" refers to the empirical relationship between a firm's current productivity and cumulative past output of a good.

The first empirical studies on learning curves date back as far as the late 19th century and come from the psychology literature. Progress curves are probably most famous because of Moore's Law, which describes advances in the production of computer chips remarkably well. However, the first empirical studies of progress curves date back to the time around World War II, when estimating and predicting productivity advances in building new types of aircraft and ships were crucial for a country's war effort.[1]

An aspect that is usually neglected in the above-mentioned empirical studies is the switching decision: when should an economic agent stop learning about an old technology and switch to a newer technology? Workers and firms are routinely

* Tel.: +34 91 624 9312.
E-mail address: matthias.kredler@uc3m.es

[1] These observations about the literature are drawn from the surveys of Thompson (2010) and Yelle (1979).

0165-1889/$ - see front matter
http://dx.doi.org/10.1016/j.jedc.2014.01.003

faced with the decision if to further exploit their accumulated knowledge in an old technology, or to upgrade to a new, more productive technology that requires learning. Examples in reality are a computer programmer switching to a new programming language, or a car company opening a new plant that implements a revolutionized manufacturing method. The technology-specific skills that agents accumulate over time are called *vintage human capital* in the literature (a term coined by Chari and Hopenhayn, 1991).

This paper aims to study the connection between learning curves and vintage human capital in a theoretical model. I take learning curves for individual workers as given and ask the following questions: when, or if at all, should workers switch to new technologies? How should firms manage the division of tasks between workers of different experience levels? What is the connection between workers′ learning curves and the firm′s/technology′s progress curve? I do this in a setting in which workers of different skill levels are complementary in production. This gives rise to a non-trivial connection between workers′ individual learning curve and the technology′s progress curve, which is a new contribution to the literature (to the best of my knowledge).

I find that workers go through cycles of learning and scrapping of skill. Technologies are shut down once returns to learning are sufficiently low, and all workers from the obsolete vintage re-locate to newer vintages. Unlike most other labor-market theories, the model predicts that workers experience initial wage cuts upon re-location. Since different skill levels are complementary and existing vintages need a supply of low-skill workers, workers enter technologies of different age cohorts and not only the frontier vintage.

The model has rich predictions on earnings profiles of workers over their careers and on the wage structure in different technologies. I show that within each technology, higher skill is always rewarded by higher wages, but that this skill premium decreases over time and eventually vanishes entirely as the technology becomes obsolete. This is driven by skill becoming less scarce as the technology ages. The model predictions are consistent with evidence provided by Michelacci and Quadrini (2009) and Kredler (in press), who find that entry wages are lower but the premium on experience is higher in fast-growing industries. They are also consistent with the premium on experience being higher in young establishments (see the evidence in Kredler, in press).

The model also predicts that a technology′s progress curve is steeper than individual workers′ learning curves. The reason is that a vintage can continuously draw on learning by newly entering workers who are on the steep part of their learning curves. Thus a firm or industry can avoid decreasing returns in learning for longer than an individual since it continuously changes its skill mix. Finally, I show that a technology′s productivity growth must fall below the economy-wide rate of technical change before it becomes obsolete and is phased out.

From the technical point of view, I exploit a continuous-time setting to study the timing decision of a technology phase-out differentially. This leads to sharp characterizations for wages and the distribution of workers in the dying vintage. I show how to state and solve a planner′s problem and construct a competitive equilibrium from there. A key problem that has to be overcome is that the distribution of workers over the state space (a choice variable for the planner) has to be linked to a collection of feasible policies by workers. To ensure this, I derive a set of constraints on the distribution that are necessary and sufficient for feasibility. I then use these constraints in a maximization problem and characterize wages and workers′ behavior, drawing heavily on the Lagrange-multiplier theorem for infinite-dimensional spaces.

The remainder of the paper is organized as follows: Section 1.2 discusses the relationship of the paper to the literature. Section 2 presents the economic environment, Section 3 defines and characterizes the competitive equilibrium. Section 4 derives further results from the vantage point of the planner′s problem and shows that the solution to the planner′s problem constitutes a competitive equilibrium. Section 5 illustrates the results in a numerical example, and Section 6 concludes.

1.2. Literature review

I first discuss the class of models most closely related to my framework: vintage-human-capital settings with multi-worker firms, which were pioneered by Chari and Hopenhayn (1991). These authors study a discrete-time setting in which workers live for two periods. There are two inputs to production: skilled and unskilled workers. The key difference to my framework is the short life span of workers, which means that there is no meaningful learning curve and that the predictions on workers′ earning profiles are less rich. This means that Chari and Hopenhayn (1991) cannot address concavity in experience-earnings profiles, which is routinely found in Mincerian earnings regressions in the data. Their model shares some predictions with the one in this paper, such as the declining skill premium in older vintages and the finite life time of technologies. However, there are also key differences: workers never experience wage losses in their setting, whereas they do occur in mine when workers re-locate to a new technology.

Even closer related to the current paper is Kredler (in press), who extends the basic setting of Chari and Hopenhayn (1991) by introducing endogenous human-capital accumulation à la Ben-Porath (1967). Many of the predictions arising from his framework are similar to mine, but they often rely importantly on the human-capital-accumulation channel (see Section 6 for a more detailed comparison between the learning-curve and the endogenous-human-capital-accumulation approach). The current paper is also able to improve upon the theoretical results in Kredler (in press) by deriving sharper characterizations of the wage structure and the link of the progress curve to the learning curve.

Second, there are models where skills are technology-specific, but which differ in other respects from mine. Both Violante (2002) and Parente (1994) consider workers who face a trade-off between a learning curve on an old technology