

Training and Effort Dynamics in Apprenticeship*

Drew Fudenberg

MIT

drew.fudenberg@gmail.com

Luis Rayo

University of Utah and CEPR

Luis.Rayo@utah.edu

April 8, 2018

Abstract

A principal specifies time paths of knowledge transfer, effort provision, and task allocation for a cash-constrained apprentice, who is free to walk away at any time. In the optimal contract the apprentice pays for training by working for low or no wages and by working inefficiently hard. The apprentice can work on both knowledge-complementary and knowledge-independent tasks. We study how the nature of the production technology influences the length of the optimal contract and its mix of effort types, and discuss the effect of regulatory limits on how hard the apprentice can work and how long the apprenticeship can last.

*First Draft: 05/2017. We thank Daron Acemoglu, Robby Akerlof, Gonzalo Cisternas, Guido Friebel, Luis Garicano, Bob Gibbons, Niko Matouschek, Derek Neal, Klaus Schmidt, Ivan Werning, and workshop participants at Kellogg, MIT, NBER, and Toronto for valuable suggestions. Sina Moghadas for excellent research assistance, and National Science Foundation grant SES 1643517 for financial support.

1 Introduction

Both in medieval times and today, employees at the beginning of their careers (e.g. apprentice bakers, prep cooks, law firm associates, medical residents, post-docs) go through a stage where they acquire knowledge and training from their employers but do enough work that the employer gains a surplus. This raises the questions of whether the employers will specify longer training periods than strictly needed for the desired knowledge transfer, and more effort than would be socially optimal.

In this paper, we study the design of optimal (profit-maximizing) careers by a principal with commitment power, who can specify time paths of knowledge transfer, effort provision, and task allocation subject to the no-servitude condition that the agent is free to leave at any time. We assume that the agent is cash constrained, and so cannot simply purchase knowledge from the principal. Instead, the agent will undergo a form of apprenticeship, where they work hard for relatively low cash payments to compensate the principal for training them. Following Becker (1964), we are interested in environments in which the knowledge that the agent wishes to acquire takes the form of general human capital: much of what bakers, doctors, and lawyers learn in their early years is fully applicable in other firms. An important feature of our model is that the principal's ability to extract payment for transferring general human capital is constrained by the apprentice's ability to leave the firm once trained without paying the principal back.¹

In our model, the agent's effort can be split between two tasks: A "skilled task" whose productivity rises with the agent's knowledge, such as writing legal briefs, and an "unskilled task" whose productivity is independent of the agent's knowledge level – this could either be menial work such as making coffee or photocopies, or fairly sophisticated work that does not however use the knowledge that the agent is working to receive. We find that the optimal contract for the principal is inefficient both due to slow training (it would be socially efficient to transfer all the knowledge at once) and because the agent will work inefficiently hard to compensate the principal for this training.

The degree of effort distortion at any given time is determined by the speed at which the principal wishes to transfer knowledge, with greater effort corresponding to faster transfer. Because the time spent on early transfers delays the later ones, the principal becomes in less of a rush as the contract unfolds and so the transfer slows and effort

¹A related literature studies borrowing by cash constrained agents who are free to walk away with the firm's capital – e.g. Thomas and Worrall (1994), Albuquerque and Hopenhayn (2004).

becomes less distorted. The overall length of the apprenticeship is in turn governed by the degree of effort distortion and allocation of effort across tasks. If only unskilled effort is distorted above the efficient level the apprenticeship lasts $\frac{1}{r}$ years, where r is the annual interest rate, regardless of the degree of effort distortion. If skilled effort is distorted the apprenticeship lasts less than $\frac{1}{r}$ years, with a greater distortion in skilled effort leading to a shorter apprenticeship. Because the optimal contract specifies inefficiently high effort together with inefficiently lengthy training, government regulation may in principal be desirable.

To isolate the principal’s strategic (rent-extraction) motive for distorting effort and transferring knowledge inefficiently slowly, in our baseline model we abstract away from bounds on the rate of learning. As we shall see, the strategic motive on its own leads to potentially very low training rates. Thus in many cases of interest it may be this strategic motive, rather than a constraint on the agent’s learning rate, that accounts for lengthy apprenticeships. Additionally, this strategic motive appears to be consistent with the motivating examples provided in the following section, where many novices spend a considerable share of their time on menial tasks, as predicted by our model, rather than on learning new skills.

We do however extend the model to allow for a fixed upper bound on the speed of knowledge transfer. We also extend the model to include, variously: training costs, certification requirements, availability of cash to the agent, and regulations. These extensions expand the applicability of the model and suggest that our main findings are robust. Further, we believe that our main insights will extend to models where the agent needs to exert skilled effort to be effectively trained, but optimal training with such “learning-by-doing” poses additional control-theoretic complications that we do not tackle here.

Ours is the first model to consider a dynamic effort-for-knowledge exchange, and to derive the resulting time-path of excessive effort. More generally, we contribute to the study of moral hazard with an endogenously evolving participation constraint.

There is an extensive literature focusing on worker training with general human capital, but this literature abstracts away from effort choice and a fortiori from the time path of effort. In addition, Katz and Ziderman (1990), Acemoglu (1997), Acemoglu and Pischke (1998), Malcomson et al. (2003) all assume that knowledge transfer is a one-time instantaneous event. These papers focus on how market frictions may allow training to occur in equilibrium even despite the difficulties in appropriating the gains from providing

general knowledge. Garicano and Rayo (2017) also suppresses effort choice but, as we do, allows for gradual knowledge transfer and shows that even absent market frictions, the principal can profit by artificially stretching the knowledge transfer over time.²

There is also an extensive literature on the effect of uncertainty about the worker’s ability, which can lead to either more or less effort than in the first best, as in Landers et al. (1996), Holmström (1999), Dewatripont et al. (1999), Barlevy and Neal (2017), Bonatti and Hörner (2017), and Cisternas (2017). All of this work abstracts from knowledge transfers.

The remainder of the paper is organized as follows. Section 2 presents motivating examples, Section 3 sets up the model, Sections 4 and 5 derive the solution, Section 6 presents comparative statics, and Section 7 contains extensions. All proofs are in the Appendix.

2 Motivating examples

Work-for-training arrangements are common in a wide range of industries. Frequently, in these arrangements the apprentice experiences long hours and heavy workloads, and initially spends a large share of her time on menial work, in many cases unrelated to the skills she wishes to acquire from the expert. Moreover, the apprentice is often paid low (sometimes even zero) wages while being trained.

To start with, consider the restaurant industry. Star chefs possess coveted knowledge that aspiring chefs wish to acquire (e.g. Gergaud et al., 2017). Once a chef is trained, she can opt to work on her own and keep all her earnings; as a result, master chefs can only obtain rents from their apprentices before they are fully trained. In upscale restaurants, apprentices endure years of grueling work while advancing through well-defined career stages. For example, over seven-plus years, young cooks at Le Gavroche restaurant in London, under the tutelage of world-renowned chef Michel Roux Jr., gradually progress from doing, as Roux himself warns, “the jobs no one wants” (under the Apprentice position) to “workhorse” prep work (the Commis position) and eventually, if successful, to supervising activities until becoming a Head Chef, either at Le Gavroche or elsewhere.³

²Hörner and Skrzypacz (2016) study gradual information revelation by a privately informed agent about her competence; their model has neither effort nor human capital.

³As noted by Verena Lugert, a former employee of the famed Gordon Ramsey: “[Aspiring chefs] pay into a blood-sweat-and-tears account and hope for a return in form of titles: Demi Chef, Chef de Partie,

Memorable examples of menial work can be found in the documentary *Jiro Dreams of Sushi*, where Jiro, arguably the world’s top sushi chef, takes roughly ten years to train his apprentices. A large share of their time is devoted to such monotonous tasks as cleaning and preparing seafood – including massaging octopus meat for 40 to 50 minutes per batch – and toasting seaweed by hand. Apprentices spend months on some of these tasks; arguably far beyond the point where they have been mastered. Thus, rather than merely training his novices, the master appears to be strategically extracting rents from them in exchange for his knowledge.

Apprenticeships date back to at least the European trade guilds starting in the 12th century, where they served as the main source of training for artisans and merchants (Jovinelly and Netelkos, 2007). At the same time, they gave rise to opportunities for exploitation: “Master craftsmen and tradesmen took in young learners and gave them menial tasks that make filing and photocopying look plush” (Spradlin, 2009). Adam Smith considered industrial-revolution apprenticeships, which usually lasted seven years, to be excessively long and poorly paid. He viewed this arrangement as a response to the agent’s liquidity constraints: “During the continuance of the apprenticeship, the whole labour of the apprentice belongs to his master. In the mean time he must, in many cases, be maintained by his parents or relations, and in almost all cases must be cloathed by them”... “They who cannot give money [to the master], give time, or become bound for more than the usual number of years; a consideration which... is always disadvantageous to the apprentice” (Smith , 1872, p. 93).

During the industrial revolution, long hours were commonplace and even a cause for public concern. Lane (1996) notes that a 14-hour workday was typical, with frequent cases of even longer hours: “Shoemakers also theoretically worked a 14-hour day, but [apprentice] George Herbert’s memories recorded that he often worked ‘for three weeks together from three or four in the morning till ten at night’ ”. Just as with current-day apprenticeships, 17th and 18th century apprenticeships commonly began with a period of menial work: Ayres (2014) notes, “In acquiring a craft skill a youth was put through an almost military discipline. After one or two years engaged in menial tasks: fetching and carrying, sweeping the workshop floor or lighting the stove, an apprentice woodcarver

[etc.] Everything is worth it: Dedication. Burn scars. Ego-Devastation. And checking in to a parallel universe [where] ‘Only lazy [people] need sleep!’ [and] working eight hours is called ‘briefly coming in for half a day.’ ” See www.micheleroux.co.uk/working.html and Lugert (2017).

might be granted the privilege of learning to sharpen tools” (p. 350).⁴

In modern times, a number of high-end professions, including medicine, scientific research, and professional services (e.g. law, accounting, banking, architecture), exhibit some of the same features. In the medical profession, residencies constitute a form of mandatory apprenticeship. As noted by Park (2017), residencies are “structured to serve the dual, often dueling, aims of training the profession’s next generation and minding the hospital’s labor needs”, with hospitals constantly struggling to “stay on the right side of the boundary between training and taking advantage of residents.”⁵ In the U.S., residents typically endure a grueling 80-hour work week; in contrast, less than a quarter of fully trained doctors work for more than 60 hours a week (e.g. Landrigan et al., 2004, American Medical Association, 2015). A significant portion of a resident’s shift is usually spent on menial tasks, known in the medical profession as “scut” work, such as inserting IV lines, wheeling patients around, and performing lengthy administrative work, which are all valuable to the hospital but provide limited learning opportunities for the apprentices (Jauhar, 2015).⁶

In science careers, postdoctoral positions are widely used and, increasingly, a cause of public concern (e.g. Stephan and Ma, 2005, Stephan, 2013). Postdocs, especially in the life sciences, spend years working long hours (with an average 53-hour week) at low wages (around \$16/hour in 2012) in the hope of gaining skills and access to a tenure-track job.⁷ Thanks to an abundance of young aspiring scientists, many postdoc employers are able to train their postdocs very slowly, all the while using them as a form of cheap labor: “many postdoctoral scholars – especially those not funded by training grants or fellowships – are but poorly paid research assistants who receive little mentoring and have few opportunities to develop an independent research agenda.” (Stephan, 2013, p. 245.)

In professional service firms, where young professionals seek to acquire knowledge from the firms’ partners, long hours and heavy workloads are common as well (e.g. Coleman

⁴For contemporary examples of harsh apprenticeships, see UK Dept. for Business, Innovation and Skills (2013).

⁵Indeed, “[l]ong hours and hard work have been features of medical training since the modern residency program had its beginnings at the Johns Hopkins Hospital in Baltimore in the late 19th century” (Jauhar, 2015, *New York Times*).

⁶Schwartz et al. (1992) find that in-hospital hours of surgical residents averaged 98 per week, with hours slightly declining over time from around 100 hours for interns (first-year residents), 97 for junior residents, and 95 for chief residents. About 20 hours a week were spent on menial tasks.

⁷As noted by Stephan (2013): “If, instead, faculty members were to staff their labs with staff scientists, they would have to pay 50–100 per-cent more than they pay to a postdoc” (p. 245).

and Pencavel, 1993, Landers et al., 1996, Barlevy and Neal, 2017). Not unlike medieval apprentices and cooking trainees, workers in the early stages of their careers are frequently assigned mind-numbing grunt work (e.g. Maister, 1993). A *Financial Times* article notes: “There is no simple fix for an entrenched culture of overwork at professional services firms. The fact that an entry-level analyst at a Wall Street bank is required to sacrifice his or her personal life to the job – sitting at a desk until dawn, eating order-in food and correcting invisible errors in spreadsheets – has been built into the system” (Gapper, 2014).⁸

Lastly, a contemporary case of harsh apprenticeship, with conditions reminiscent of industrial-revolution apprenticeships, is that of manicurists in the New York City area. As documented by Maslin Nir (2015), “workers routinely work up to 12 hours a day, six or even seven days a week [...] enduring all manner of humiliation”. Rampant wage violations (wages below the legal minimum, tip and wage skimming, no overtime pay) have kept wages very low, with \$30-\$40 per day being typical. Furthermore, at the beginning of a typical career, aspiring manicurists “must hand over cash – usually \$100 to \$200, but sometimes much more – as a training fee. Weeks or months of work in a kind of unpaid apprenticeship follows.” As they acquire skills and pay their dues through menial work, they eventually advance through various career stages: “ ‘Little job’ is the category of the beginners. They launder hot hand towels and sweep toenail clippings. They do the work others do not want to do [...] ‘Medium Job’ workers do regular manicures [...] ‘Big Job’ employees are veterans, experts at sculpting false nails out of acrylic dust.”

3 Model

A principal (she) and an agent (they) interact over an infinite horizon. Both players have quasilinear utility in money and discount future payoffs at rate r . Time t runs continuously. At time t , the agent combines a stock of knowledge $X_t \in [0, \bar{X}]$ and two sorts of effort a_t, b_t to produce output y_t . Effort a_t is “skilled,” meaning that its productivity is increasing in the agent’s knowledge X_t . Effort b_t is “unskilled,” meaning that its productivity is independent of X_t . Let $l_t := a_t + b_t$ denote total effort. We assume

⁸While our model helps account for some aspects of these professions, it fails to account for several others – such as firms usually offering multiple career paths, many young professionals not making partner, and many leaving their firms before completing their training. Barlevy and Neal (2017) explain these patterns as the result of the firm learning about the agent. Professional service firms also tend to pay higher wages to their novices than other careers.

that both a_t and b_t are non-negative and l_t is bounded above by a constant, which we normalize to 1.

Thus total output y_t is given by

$$y_t := f(X_t, a_t) + g(b_t).$$

Exerting effort l_t imposes cost $c_t := c(l_t) \geq 0$ on the agent. We assume that f is strictly increasing in knowledge, f and g are non-decreasing in effort, c is strictly increasing, and all three functions are twice differentiable. Let $v(X) := \max_{a,b \geq 0} [f(X, a) + g(b) - c(a + b)]$ denote first-best surplus given X . Since f is strictly increasing in X , so is v . We assume throughout that $\lim_{X \rightarrow \infty} v(X) = \infty$.

The agent starts with some exogenous stock of knowledge $\underline{X} \in [0, \bar{X}]$. The agent's stock of knowledge can never decrease, and the only way it can increase is by transfers from the principal. In our baseline model, the principal is able to costlessly and instantaneously increase the agent's knowledge to any level up to \bar{X} . (Section 7 shows that the results are qualitatively unchanged if the principal faces a cost of training the agent or if there is a fixed upper bound on the rate of knowledge transfer, provided this bound is not too tight.) The agent has no other way to obtain knowledge. As a result, the principal can select any weakly increasing function X_t with range in $[\underline{X}, \bar{X}]$. Let $X_\infty := \lim_{t \rightarrow \infty} X_t$. When X_∞ is reached in finite time, we say that the agent graduates at time $T = \inf \{t : X_t \geq X_\infty\}$. Otherwise, we set $T = \infty$ and say that the agent never graduates.

The agent has access to the same output technology when working for the expert and when working on their own. The agent also has access to an alternate employment that pays \underline{v} , with $0 < \underline{v} < v(\bar{X})$, so that total surplus is maximized by fully training the agent. As a result, if the agent walks away with knowledge X , they obtain instantaneous surplus $\max\{\underline{v}, v(X)\}$, and since their knowledge will be constant from then on, their outside option at date t is $\frac{1}{r} \max\{\underline{v}, v(X)\}$. In addition, the agent has access to a savings account that pays interest r , but has no capital up-front and has no ability to borrow so they cannot purchase knowledge from the expert at date 0.

At time 0, the principal offers the agent an employment contract, denoted $S = (T, \{X_t, a_t, b_t, w_t\}_{t=0}^T)$, consisting of a graduation date T and a path that specifies for each $t \in [0, T]$ a knowledge stock X_t , effort levels (a_t, b_t) , and a money transfer w_t from principal to agent, which we call a wage. Between dates 0 and T , all output net of wages

belongs to the principal. After date T , the agent works on their own and keeps all output.⁹ While the principal can commit to this contract, the agent can walk away at any time; if the agent does so, the principal does not hire them back.

Given contract S , the principal's and agent's continuation values from date $t \leq T$ onward are

$$\begin{aligned}\Pi_t(S) &= \int_t^T e^{-r(\tau-t)} [y_\tau - w_\tau] d\tau \quad \text{and} \\ U_t(S) &= \int_t^T e^{-r(\tau-t)} [w_\tau - c_\tau] d\tau + e^{-r(T-t)} \frac{1}{r} v(X_T),\end{aligned}\tag{1}$$

where $\frac{1}{r}v(X_T)$ is the “prize” received by the agent upon graduation.

The principal can select any contract she desires subject to two constraints. First, a *participation constraint* for the agent requiring that, at each date $t \leq T$, the agent's continuation value is at least as high as their outside option:

$$U_t(S) \geq \frac{1}{r} \max \{ \underline{v}, v(X_t) \}.\tag{2}$$

Second, a *liquidity constraint* for the agent requiring that, up to any given date $t \leq T$, the agent's cumulative earnings are non-negative in present value:

$$\int_0^t e^{-r\tau} w_\tau d\tau \geq 0.\tag{3}$$

This constraint captures the assumption that the agent starts the relationship with no money and cannot borrow.

⁹Contracts where the agent continues to work for the principal after T are weakly dominated because once the knowledge transfer has ended, the agent demands wages at least equal to output. Contracts where the agent works on their own during some periods prior to T are strictly dominated because this delays the principal's profit flow.

The principal's problem is

$$\max_{T, \{X_t, a_t, b_t, w_t\}_{t=0}^T} \int_0^T e^{-rt} [y_t - w_t] dt \quad (4)$$

subject to (2) and (3)

and subject to $a_t, b_t \geq 0$, $a_t + b_t \leq 1$, and $X_t \in [\underline{X}, \overline{X}]$ non-decreasing in t . We say that a contract is optimal if it is a solution to this maximization problem.

Note that we can model situations where the principal must pay the agent a strictly positive subsistence wage with the same formalism: if the required wage is \underline{w} , we can define $\hat{f} := f - \underline{w}$, $\hat{y} := y - \underline{w}$, and $\hat{v} := v - \underline{w}$ and model the situation using (1), (2), and (3). To allow for this interpretation of the model, we will allow \hat{y} and \hat{v} to be negative. Thus the wages w_t in our model should be thought of as wages in excess of the minimum, though we will simply call them “wages” in what follows. Note that with this interpretation of v , we should not expect the constraint $v(\overline{X}) > \underline{v}$ to be satisfied when \underline{w} is very large.

4 Preliminaries

In Section 4.1 we derive general properties that any undominated contract must satisfy, without yet verifying that an optimal contract exists. In Section 4.2 we use these properties to formulate the principal's problem as one of optimal control and then verify that a maximum in (4) is indeed attained. Section 5 uses additional assumptions to show that the optimal contract is unique and then to characterize its length and effort path.

4.1 General properties of undominated contracts

Lemma 1 tells us that the principal can obtain a strictly positive profit by contracting with the agent and that she need only consider contracts that transfer all knowledge in finite time and pay 0 wages. Let $W_\infty := \int_0^\infty e^{-rt} w_t dt$ be the present value of the agent's wages.¹⁰

Lemma 1

¹⁰Recall that wages here can be interpreted as wages above subsistence or a legally required minimum. Lemma 1 shows that Lemmas 1 and 2 in Garicano and Rayo (2017) generalize to the case where the agent exerts effort and has an ex-ante outside option with arbitrary value.

1. *The principal can obtain a strictly positive profit by contracting with the agent.*
2. *Any contract with $W_\infty > 0$ is strictly dominated by a contract with 0 net payment, and any contract where the agent does not acquire all of the principal's knowledge in finite time is strictly dominated by a contract where they do. Moreover, it is without loss to require that $w_t = 0$ at all times t .*

To gain intuition for the first part of this lemma, note that the principal can train the agent to an intermediate level $X' \in (\underline{X}, \bar{X})$ such that $v(X') > \underline{v}$, and induce them to work at zero wages for some (possibly short) period of time in exchange for the remaining knowledge at the end of the contract.

To gain intuition for the second part of this lemma, recall that the efficient outcome would be for the principal to immediately transfer all of her knowledge to the agent so that the agent could work on their own. Because the agent is credit-constrained this does not occur, and instead the agent works for the principal while being trained. The key to the lemma is that any contract with $W_\infty > 0$ can be improved by a contract with an earlier graduation date, the same effort path up to graduation, 0 wages, and the same initial value for the agent. Because the new contract replaces wages with a more valuable final reward, the agent is less tempted to walk away while being trained, and so the new contract meets all of the participation constraints. Then because the two sides have the same discount rate, and it is efficient to transfer knowledge earlier, the new 0-wage contract has higher joint surplus, and since the agent is indifferent between the two contracts the new one makes the principal strictly better off. Infinite-length apprenticeships are dominated because they would require a net payment to the agent, and the principal can improve a finite-length apprenticeship that doesn't transfer all knowledge by trading the rest of it for some more work. Moreover, it is without loss to set wages to be zero at all times, because there is no gain to the principal making a loan to the agent.

The next lemma states two additional simplifying properties that we can impose when looking for the optimal contract: If the participation constraints do not hold with equality the principal can do better by increasing X_t until they do, and the principal prefers to allocate total effort $l_t = a_t + b_t$ between the two tasks to maximize output.

Let $\bar{y}(X, l) := \max_{a \in [0, l]} f(X, a) + g(l - a)$ denote the maximum possible output given knowledge and effort levels (X, l) .

Lemma 2 *Any contract is weakly dominated by a contract that at all times sets the agent's*

participation constraint to hold with equality and allocates total effort $l_t = a_t + b_t$ across tasks so as to maximize output, so that $y_t(X_t, a_t, b_t) = \bar{y}(X_t, l_t)$.¹¹

For intuition, note that since output is strictly increasing in X_t , and knowledge transfer is costless, the principal wishes to raise X_t to the point where the agent's participation constraint holds with equality. In addition, since the allocation of total effort across tasks does not affect the agent's utility, it is optimal for the expert to choose an allocation that maximizes output. Notice also that any contract with $\underline{v} > v(X_t)$ for some t can be improved by increasing X_t to $v^{-1}(\underline{v})$ over the time interval where $\underline{v} > v(X_t)$; as a result, we may assume without loss that $v(\underline{X}) \geq \underline{v}$ and write the agent's ex-ante outside option more compactly as $\frac{1}{r}v(\underline{X})$.

The principal's problem therefore simplifies to selecting a finite graduation date T and time paths of knowledge and total effort before this date. Her problem is

$$\max_{T, \{X_t, l_t\}_{t=0}^T} \int_0^T e^{-rt} \bar{y}(X_t, l_t) dt \quad (5)$$

subject to

$$U_t = \underbrace{e^{-r(T-t)} \frac{1}{r} v(\bar{X})}_{\text{discounted "graduation prize"}} - \underbrace{\int_t^T e^{-r(\tau-t)} c(l_\tau) d\tau}_{\text{"effort tax"}} = \frac{1}{r} v(X_t)$$

and subject to $0 \leq l_t \leq 1$, $X_t \in [\underline{X}, \bar{X}]$, and X_t non-decreasing in t . We denote the solution to this problem by $(T^*, \{X_t^*, l_t^*\}_{t=0}^{T^*})$.

Proposition 1 *Every optimal contract $(T^*, \{X_t^*, l_t^*\}_{t=0}^{T^*})$ is sequentially optimal: If at any given time $t < T^*$ the agent's stock of knowledge is X_t^* , then the truncated contract $(T^*, \{X_s^*, l_s^*\}_{s=t}^{T^*})$ maximizes profits from time t onward.*

This result shows that even though the principal has commitment power, she does not make commitments that she would later prefer to undo. For this reason it is sufficient for the principal to be able to commit to spot contracts.¹² This is because the agent's

¹¹The contract is strictly dominated if it fails to satisfy any of these two properties over a positive-measure fraction of time.

¹²In a discrete time model it would be sufficient for the principal to commit at the beginning of each period to transfer knowledge at the end of it, conditional on the agent exerting the agreed effort.

dynamic participation constraint is always binding, so any feasible change the principal would want to make at some time t could not make the agent worse off from then on. Thus any potential improvement to the contract could be implemented ex-ante. Proposition 1 implies that the current stock of knowledge determines the path of future effort and the speed at which knowledge is transferred. We use this observation later to provide intuition for the optimal effort and knowledge paths. Note that at this point we have not yet proven that an optimal contract exists; we will do so in the next section.

4.2 Optimal control

It is convenient to state the principal's problem as an optimal control problem in which the principal's *control variables* are the agent's effort levels $\{l_t\}_{t=0}^T$ and the choice of terminal time T , the *state variable* is the agent's continuation value, measured as a flow payoff,

$$u_t := rU_t = v(X_t),$$

with $\dot{u}_t = r[u_t + c(l_t)]$, and the agent's knowledge stock X_t is given by $\phi(u_t) := v^{-1}(u_t)$, so that the agent's participation constraint is met with equality. We assume that l_t has bounded variation.

Recall that the principal is able to instantaneously increase X to any level no greater than \bar{X} . Total surplus would be maximized if the principal instantaneously transferred all her knowledge to the agent, that is if $X_0 = \bar{X}$, but in that case the principal would get no benefit from the knowledge transfer. As we will see, in some cases the principal does want to make an instantaneous knowledge gift to the agent, but this will only occur at the initial time.¹³ The principal does this by choosing a $u_0 \in [v(\underline{X}), v(\bar{X})]$.

Note that the *state equation* $\dot{u}_t = r[u_t + c(l_t)]$ says that the value of the knowledge gained at time t equals the total opportunity cost of working for the principal, which includes both the labor cost c and the loss from postponing the outside option.

¹³An example would be a summer crash course at the beginning of a consulting career.

The optimal control problem is:

$$\begin{aligned} & \max_{u_0, T, \{l_t\}_{t=0}^T} \int_0^T e^{-rt} \bar{y}(\phi(u_t), l_t) dt & (6) \\ & \text{subject to} \\ & u_0 \in [v(\underline{X}), v(\bar{X})], \quad u_T = v(\bar{X}), \quad \dot{u}_t = r[u_t + c(l_t)], \\ & 0 \leq l_t \leq 1. \end{aligned}$$

Let λ_t denote the co-state variable, form the Hamiltonian $\mathcal{H} = e^{-rt} \bar{y}(\phi(u_t), l_t) - \lambda_t \dot{u}_t$, and adjoin the effort constraints with multipliers η_t, γ_t to form the Lagrangian $\mathcal{L} = \mathcal{H} + \eta_t [1 - l_t] + \gamma_t l_t$.

Lemma 3 *A solution to problem (6) exists. Moreover, the following system is necessary for optimality:*

$$\dot{u}_t = ru_t + rc(l_t); \quad \dot{\lambda}_t = \partial \mathcal{L} / \partial u_t \quad (7)$$

$$\partial \mathcal{L} / \partial l_t = 0 \quad (8)$$

$$\eta_t, \gamma_t \geq 0; \quad \eta_t [1 - l_t] = \gamma_t l_t = 0 \quad (9)$$

$$\mathcal{H}_T = 0; \quad \lambda_0 \geq 0; \quad \lambda_0 [u_0 - v(\underline{X})] = 0. \quad (10)$$

Expression (7) contains the state and co-state evolution equations; (8) contains the first-order condition for l_t ; (9) contains the complementary slackness conditions for the Lagrange multipliers η_t, γ_t ; and (10) contains the transversal condition for the terminal time T , where λ_0 is the multiplier for the constraint $u_0 \geq v(\underline{X})$.¹⁴ In later sections we impose additional structure on the production functions f, g that guarantee that the solution to the necessary conditions is unique and hence sufficient.

We solve the optimal control problem by starting from an arbitrary terminal time T , where the agent has all knowledge so that the state is $u_T = v(\bar{X})$, and then running time in reverse. At each time, the state determines the optimal effort level, and the state and effort level combined determine the time derivative of the state. As time moves backward, the state $v(X_t)$ continues to fall until either: (1) the state reaches the agent's ex-ante outside flow payoff $v(\underline{X})$, in which case $v(X_0) = v(\underline{X})$ and there is no initial knowledge

¹⁴We omit the constraint $u_0 \leq v(\bar{X})$ here because it cannot bind.

gift, or (2) the principal would rather give the agent an initial knowledge gift and start employing them at the current time instead of having a longer apprenticeship with an initially less trained apprentice.

5 Solution

We begin by imposing further assumptions on the output and cost functions. As we will see, these assumptions guarantee that the optimal contract is unique and that agent exerts positive effort throughout the apprenticeship.

Assumption 1

1. $f_{Xa} > 0$ and $f_{aa} \leq 0$, $g' > 0$ and $g'' \leq 0$, either $f_{aa} < 0$ or $g'' < 0$, and $\lim_{X \rightarrow \infty} f(X, 1) = \infty$.
2. $c', c'' > 0$, $c'(0) = 0$, and $c'(1) > \frac{\partial}{\partial l} \bar{y}(\bar{X}, 1)$.

Let $l^{FB}(X) := \arg \max_l \{\bar{y}(X, l) - c(l)\}$ be the first-best level of total effort, which is unique from Assumption 1, and let $a^{FB}(X) := \arg \max_a \{f(X, a) + g(l^{FB}(X) - a)\}$ be the corresponding first-best level of skilled effort, which is also unique. Finally, let $\bar{a}(X, l) := \arg \max_{a \in [0, l]} \{f(X, a) + g(l - a)\}$ denote the output-maximizing skilled effort given knowledge and total effort levels (X, l) .¹⁵

Now define the *knowledge premium* given knowledge X and skilled effort a as

$$\rho(X, a) := \frac{f_X(X, a)}{f_X(X, a^{FB}(X))}.$$

Because $f_X(X, a^{FB}(X)) = v'(X)$ from the envelope theorem, this premium measures the marginal impact of knowledge on the agent's productivity inside the relationship relative to its impact on the agent's outside option. As we shall see, the knowledge premium plays a central role in determining both the optimal effort levels and the optimal contract length. Notice in particular that for any given a , an increase in X that raises the agent's outside option $v(X)$ by one unit, raises current output by $\rho(X, a)$ units.

¹⁵Note that: (1) $\lim_{X \rightarrow \infty} v(X) = \infty$ as per our maintained assumption; (2) the efficient level of total effort $l^{FB}(X)$ is strictly positive and less than 1; and (3) $\frac{\partial}{\partial l} \bar{y}(X, l) = \max\{f_a(X, \bar{a}), g'(l - \bar{a})\}$, so in particular $\frac{\partial \bar{y}}{\partial l}$ exists.

Before we characterize the optimal contract, it is useful to describe some properties of the set of payoff vectors that can be implemented by some choice of contract. The agent's overall payoff $U = \frac{1}{r}v(X_0)$, where X_0 is the initial knowledge level inclusive of any gift, must be at least their ex-ante outside option $\frac{1}{r}v(\underline{X})$, so $U \in [\frac{1}{r}v(\underline{X}), \frac{1}{r}v(\overline{X})]$. Let $\Pi^*(U)$ denote the maximum feasible profit for the principal when the agent's initial payoff is $U \in [\frac{1}{r}v(\underline{X}), \frac{1}{r}v(\overline{X})]$ and the agent's participation constraint holds with equality at all times (so no further knowledge gifts). Call the set $\{U, \Pi^*(U) \text{ such that } U \in [\frac{1}{r}v(\underline{X}), \frac{1}{r}v(\overline{X})]\}$ the feasible payoff frontier.

Lemma 4 *Along the feasible-payoff frontier the principal's (shadow) marginal cost of raising U is*

$$-\frac{d\Pi^*(U)}{dU} = 1 - r \int_0^{T^*} \rho(X_t^*, a_t^*) dt, \quad (11)$$

where T^* , X_t^* , and a_t^* are chosen to maximize profits given U .

To understand (11), fix U and pick a feasible contract that delivers profit $\Pi^*(U)$ and denote its length by T . Next, slightly lengthen this apprenticeship by keeping the agent until $T + dT$, while specifying the same effort as before up to T and first-best effort over the interval dT , so that the agent gives up an additional $y^{FB}(\overline{X}) dT$ units of output. Set $dT = e^{rT}/y^{FB}(\overline{X})$ so that, measured in period 0 units, the agent's loss is $e^{-rT}y^{FB}(\overline{X}) dT = 1$ (that is $dU = -1$). This modification gains the principal 1 in present value (captured by the first term on the right-hand side of (11)), but since the agent is now more tempted to walk away, the principal must reduce the agent's knowledge, and hence the agent's productivity, between times 0 and T . This effect is captured by the second term on the right-hand side of (11). Specifically, because at time T the agent incurs a loss of 1 unit measured in present value, at each time t before the end of the apprenticeship, to keep the agent from walking away, the principal must lower the agent's outside option by 1 unit, which lowers the agent's productivity by $r\rho_t$ units, both measured in present value. Hence, the overall change in profit from the standpoint of time 0 is $1 - r \int_0^T \rho_t dt$.

Figure 1 depicts the feasible payoff frontier for the case where the total knowledge available for sale ($\overline{X} - \underline{X}$) is large. When $X_0 = \overline{X}$ (that is, all knowledge is gifted) the principal earns zero profits and $T = 0$. As U falls, since there is more knowledge left to sell the contract grows longer and the marginal cost $-\frac{d\Pi^*(U)}{dU}$ falls, eventually becoming negative. Intuitively, a lower U means that the principal can retain the agent longer; but it also means that the agent has a lower average productivity throughout the apprenticeship.

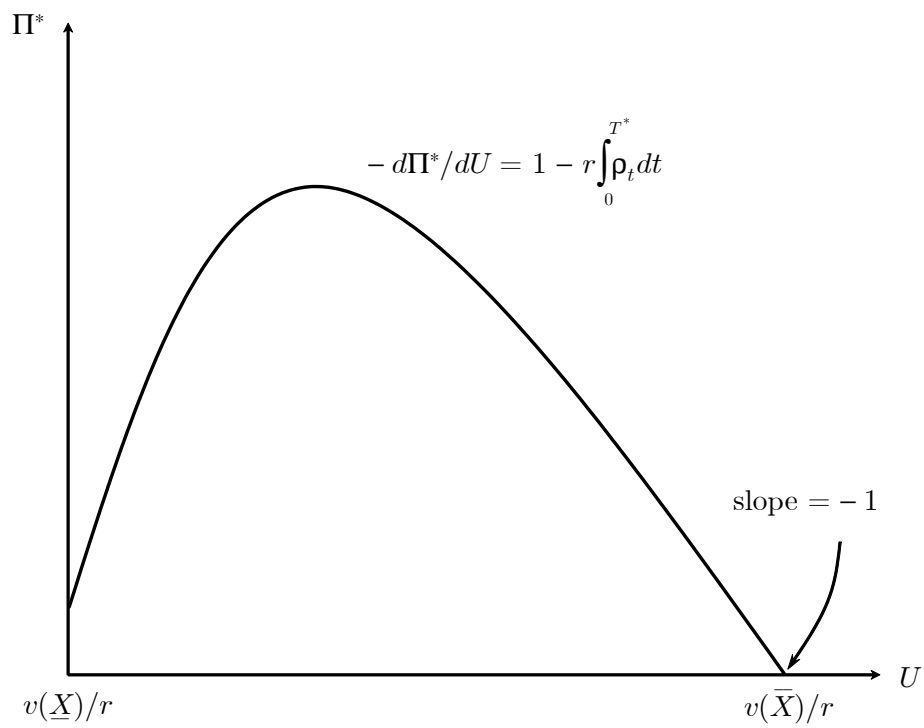


Figure 1: Feasible payoff frontier

In the decreasing portion of the frontier the first effect dominates (and so a lower U helps the principal), whereas in the increasing portion the second effect dominates.

Since knowledge can always be gifted, the Pareto frontier is the portion of the feasible-payoff frontier to the right of its peak. Notice also that the feasible payoff frontier when the agent has initial knowledge $\underline{X}' > \underline{X}$ is simply a truncated version of this one where the agent's value must be at least $\frac{1}{r}v(\underline{X}')$.

Recall from Proposition 1 that every optimal contract is sequentially optimal. Thus, as time goes by and the agent gains more knowledge, the continuation payoffs U_t and $\Pi_t = \Pi^*(U_t)$ move downward along the Pareto frontier until the agent is fully trained. Moreover, since (11) holds for all U , the (shadow) marginal cost of training the agent at each time t along the way is

$$-\left. \frac{d\Pi^*(U)}{dU} \right|_{U=U_t} = 1 - r \int_t^{T^*} \rho_\tau d\tau. \quad (12)$$

This marginal cost grows over time and is equal to 1 at the terminal time. That is, at the instant where the agent first acquires all knowledge, utility becomes transferable across players because the liquidity constraint stops binding.

It is also useful to express $-\frac{d\Pi^*(U)}{dU}$ directly in terms of the agent's promised value U . To do so, let $t(U)$ denote the time at which the agent's continuation utility is U , and let $l(U) := l_{t(U)}$ denote the effort level at this time. Upon changing the variable of integration in (12), we obtain¹⁶

$$-\frac{d\Pi^*(U)}{dU} = 1 - r \int_U^{\frac{1}{r}v(\bar{X})} \frac{\rho_{t(Z)}}{rZ + c(l(Z))} dZ.$$

By differentiating this expression with respect to U , we obtain the second derivative of the payoff frontier:

$$\frac{d^2\Pi^*(U)}{dU^2} = -\frac{\rho_{t(U)}}{U + \frac{1}{r}c(l(U))} < 0. \quad (13)$$

We now characterize the (unique) optimal contract in two theorems. Theorem 1 characterizes the effort and knowledge paths:

¹⁶To see why, notice that the state equation in (7) implies that $t'(U) = 1/\dot{U} = 1/[rU + c(l(U))]$. Moreover, $-\frac{d\Pi^*(U)}{dU} = 1 - r \int_0^T \rho_t dt = 1 - r \int_U^{\frac{1}{r}v(\bar{X})/r} \rho_{t(Z)} t'(Z) dZ$.

Theorem 1

1. The optimal effort path is efficient at the terminal time T^* , and exceeds the efficient level at all earlier times, that is $l_{T^*}^* = l^{FB}(\bar{X})$ and $l_t^* > l^{FB}(X_t^*)$ for $t < T^*$.
2. Moreover the optimal knowledge and total effort paths uniquely satisfy, for all $t \leq T^*$,

$$\frac{1}{r} \frac{d}{dt} v(X_t^*) = v(X_t^*) + c(l_t^*) \quad (14)$$

and

$$\underbrace{\frac{c'(l_t^*)}{\frac{\partial \bar{y}}{\partial l}(X_t^*, l_t^*)}}_{\text{Effort distortion}} = \min \left\{ \underbrace{\frac{1}{1 - r \int_t^{T^*} \rho_\tau d\tau}}_{\text{Target distortion}}, \underbrace{\frac{c'(1)}{\frac{\partial \bar{y}}{\partial l}(X_t^*, 1)}}_{\text{Maximum feasible distortion}} \right\}, \quad (15)$$

where $v(X_{T^*}^*) = v(\bar{X})$ and $\rho_\tau = \rho(X_\tau^*, a_\tau^*)$ is the knowledge premium at τ .

Equation (14) is the state equation. It tells us that the present value of the knowledge acquired at date t must equal the agent's total opportunity cost of working for the expert. Equation (15) characterizes the ratio $c' / (\partial \bar{y} / \partial l)$, which equals 1 in the first best. Absent an upper bound on the agent's total effort, the principal would set $c' / (\partial \bar{y} / \partial l) = \left(1 - r \int_t^{T^*} \rho_\tau d\tau\right)^{-1}$; when this is not feasible she sets it as high as possible. Notice also that the target distortion strictly decreases over time.

To gain intuition for this result, consider an optimal contract with terminal time T^* . Suppose the agent already has knowledge X_t and consider how much effort l_t to ask of the agent over the next unit of time. The marginal benefit of higher effort is the gain in current output $\partial \bar{y} / \partial l$, and since the principal pays for this effort with training, the marginal cost of higher effort is $-\frac{d\Pi^*(U_t)}{dU_t} \cdot c'$. Hence, at the terminal time where $-\frac{d\Pi^*(U_t)}{dU_t} = 1$, the principal internalizes the full effort cost, and therefore sets effort equal to the first best. At all times before then, training the agent is cheap (as it raises the agent's productivity throughout the remainder of the apprenticeship). As a result, the principal only partly internalizes the effort cost, and therefore optimally distorts effort above the first best. Finally, from (12) we can see that provided the effort upper bound is slack, $c' / (\partial \bar{y} / \partial l) = -1 / \frac{d\Pi^*(U_t)}{dU_t} = \left(1 - r \int_t^{T^*} \rho_\tau d\tau\right)^{-1}$.

Theorem 2 characterizes the (unique) optimal initial knowledge level and contract length:

Theorem 2 *The optimal initial knowledge level X_0^* and contract length T^* are unique. Moreover, for all \underline{X} , there is a threshold $0 < \Delta < \infty$ such that*

(a) *if $\bar{X} - \underline{X} > \Delta$, then there is a positive knowledge gift ($X_0^* > \underline{X}$) and $r \int_0^{T^*} \rho_t dt = 1$*

(b) *if $\bar{X} - \underline{X} \leq \Delta$, then there is no knowledge gift ($X_0^* = \underline{X}$) and $r \int_0^{T^*} \rho_t dt \leq 1$,*

where $\rho_t = \rho(X_t^*, a_t^*)$.

To understand this result, recall that when selecting the initial gift, the principal faces a trade-off between starting the apprenticeship with a more productive agent and being able to hold on to the agent for longer. Consider first the case where the total knowledge available for sale ($\bar{X} - \underline{X}$) is sufficiently large that the peak of the feasible-payoff frontier occurs at a knowledge level strictly higher than \underline{X} . In this case, the principal simply gifts enough knowledge to reach the peak at time 0, and gradually sells the remaining knowledge starting from there. Since $-\frac{d\Pi^*(U)}{dU} = 0$ at the peak, equation (12) tells us that T^* is just large enough that a 1 unit increase in revenue from keeping the agent slightly longer at the end of the apprenticeship, while exerting first-best effort, is exactly offset by the lost revenue $r \int_0^{T^*} \rho_t dt$ due the need to lower the agent's knowledge between 0 and T^* . When instead $\bar{X} - \underline{X}$ shrinks to the point where the feasible-payoff frontier always has a negative slope, then the ex-ante participation constraint $U \geq \frac{1}{r}v(\underline{X})$ binds, and here the contract begins with no gift and has a shorter length than before.

Recall from Theorem 1 that the unconstrained effort level is given by $\frac{c'(l_t)}{\frac{\partial y}{\partial l}(X_t, l_t)} = -1/\frac{d\Pi^*(U_t)}{dU_t}$. Then, when there is a knowledge gift we have $-\frac{d\Pi^*(U_0)}{dU_0} = 0$, and so the time 0 target effort distortion is infinitely large. Thus, if it were feasible, the principal would prefer to ask the agent to initially work infinitely hard for an infinitely small length of time to pay for the initial knowledge transfer, but as this is not possible the principal instead transfers the initial knowledge for free. Notice finally that the contract becomes longer as r falls. This result, obtained by Garicano and Rayo (2017) for the special case with zero effort, follows from the fact that when players become more patient knowledge becomes more valuable, so the agent is willing to work longer to acquire it.

The following corollary describes a general link between the type of effort distortions and the contract length:

Corollary 1 *When the agent's ex-ante participation constraint is slack and the contract always specifies an efficient level of skilled effort, then $T^* = \frac{1}{r}$; otherwise $T^* < \frac{1}{r}$.*

This result is immediate: efficient skilled effort at time t means that $\rho_t = 1$. The optimal length therefore satisfies $r \int_0^{T^*} dt = rT^* = 1$. Conversely, since total effort is never inefficiently low and effort is always allocated efficiently, ρ_t is never less than 1, so if it is ever greater than 1 we have $\int_0^{\frac{1}{r}} \rho_t dt > \frac{1}{r}$ so $T^* < \frac{1}{r}$. In Section 6 we provide examples that illustrate each of these cases.

Pareto-efficient contracts. Recall that the principal's optimal contract is sequentially optimal (Proposition 1). Hence in this contract, as time goes by, the players' continuation payoffs trace the Pareto frontier. Consequently, any Pareto-efficient payoffs (Π, U) are uniquely implemented by an initial knowledge gift to raise the agent's knowledge to $v^{-1}(rU)$, followed by the principal's optimal contract when the agent starts with that knowledge level. That is, every Pareto-efficient contract is simply a truncated version of the principal's optimal contract, where the agent begins the contract farther along and therefore endures a shorter apprenticeship. Notice also that along the Pareto frontier, as U grows and Π falls the effort distortion decreases, so total surplus grows monotonically, and reaches the first-best level when Π falls all the way to 0.

6 Comparative Statics

Here we study how the relative productivity of the two tasks and the curvature of the cost function impact the agreement. We begin by considering the two special cases where the agent devotes all their effort to one of the tasks. We then turn to the intermediate case where the agent works on both tasks.

A. Unskilled effort only. Suppose $f_a(\bar{X}, 0) < g'(1)$ so at each time the agent exerts effort only on the unskilled task. Suppose further that \bar{X} is large enough that the agent's ex-ante participation constraint is slack. Since at all times $a_t^* = a^{FB}(X_t) = 0$, the knowledge premia ρ_t^* are all 1, and the optimal contract length is $T^* = \frac{1}{r}$ from Corollary 1 (for example, when the annual interest rate is 5%, T^* is 20 years). Theorem 1 then implies that at all times before the terminal time, unskilled effort exceeds the first best,

and satisfies

$$\frac{c'(b_t^*)}{g'(b_t^*)} = \min \left\{ \frac{T^*}{t}, \frac{c'(1)}{g'(1)} \right\}.$$

Thus effort weakly falls over time, strictly so whenever effort is below its upper bound 1.

Figure 2.A depicts an optimal contract. The optimal effort distortion balances a loss in instantaneous surplus against a higher rate of knowledge transfer. Early on, when the principal is in most of a rush to transfer knowledge (i.e. the shadow cost of training the agent is low), effort equals its upper bound. Note that if the maximum feasible effort were reduced to b^{FB} the agent would endure an apprenticeship that is equally long (Corollary 1) but less costly per unit of time, which means the agent could be granted more knowledge throughout (as the agent is less tempted to walk away). Thus overworking the agent helps the principal, but hurts the agent and lowers social surplus.

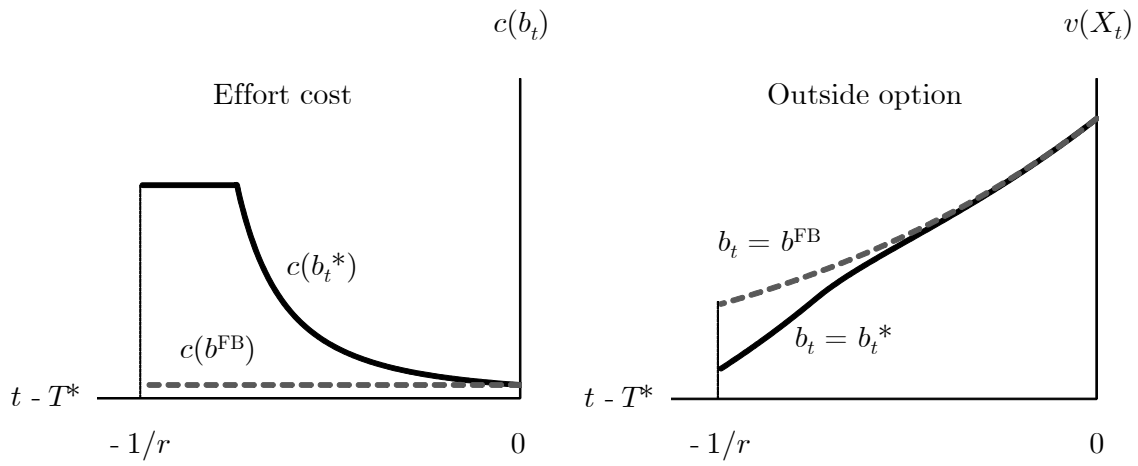
Figure 2.B illustrates how the contract changes with the curvature of the cost function c . In this figure as σ varies, α , β and the maximum feasible effort vary so that b^{FB} , $c(b^{FB})$ and the maximum feasible effort cost are held constant.¹⁷ As c becomes more linear (σ falls) effort above first best becomes less wasteful and therefore closer to a money transfer. As a result, the principal is willing to impose a higher effort cost on the agent. From (13) the second derivative of the Pareto frontier is $d^2\Pi^*/dU^2 = -1/[U + c(l(U))/r]$. As a result, as c becomes more linear the Pareto-frontier becomes more linear as well. This raises the principal's payoff, but since a higher effort cost makes the apprenticeship less attractive to the agent, the knowledge path must fall to keep the agent's outside option low enough that they do not walk away.

B. Skilled effort only. Suppose $f_a(\underline{X}, 1) > g'(0)$, so at each time the agent exerts effort only on the skilled task. Theorem 1 implies that at all times before the terminal time, skilled effort exceeds the first best, and satisfies

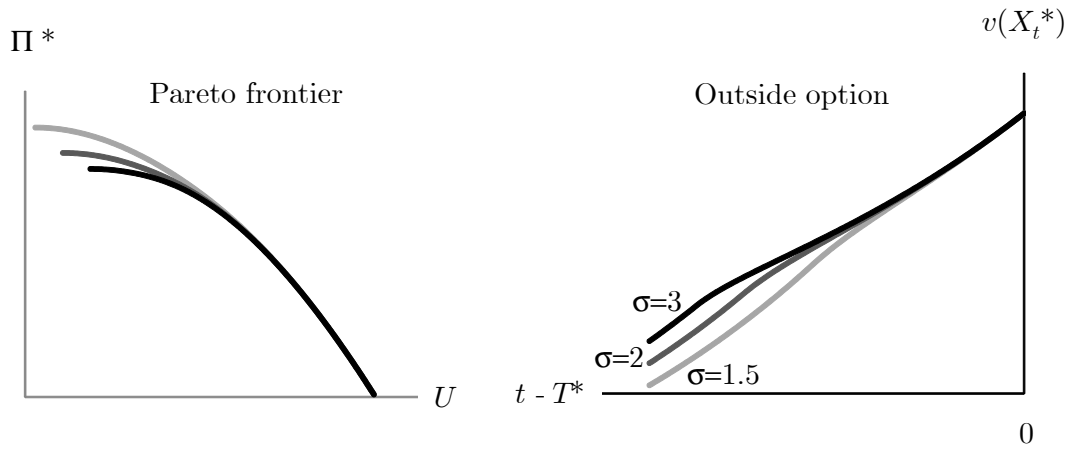
$$\frac{c'(a_t^*)}{f_a(X_t^*, a_t^*)} = \min \left\{ \frac{1}{1 - r \int_t^{T^*} \rho_\tau^* d\tau}, \frac{c'(1)}{f_a(X_t^*, 1)} \right\}.$$

Through the lifetime of the contract the effort distortion falls because $f_{Xa} > 0$, X_t^* grows, and $r \int_t^{T^*} \rho_\tau^* d\tau$ falls. However since the efficient effort is increasing over time, the optimal effort level is either weakly decreasing or non-monotone depending on the details of f and

¹⁷We keep b^{FB} and $c(b^{FB})$ constant so that the agent's steady-state payoff and the rate of knowledge transfer at the end of the apprenticeship are unchanged; we keep the maximum feasible effort cost constant so that the maximum rate of knowledge transfer does not change.



Panel A. $y = X + b$ and $c = b^2$.



Panel B. $y = X + b$ and $c = \alpha + \beta b^\sigma$.
(Lighter curves correspond to lower σ .)

Figure 2: Unskilled effort only.

*c.*¹⁸


Corollary 1 implies that the optimal apprenticeship length is always strictly less than $\frac{1}{r}$. This is because the complementarity between effort and knowledge means that increasing effort past the first-best level also increases the principal's incentive to train the agent. In contrast to the case with only unskilled effort, where the principal extracts rents from the agent by combining overwork with a lengthy apprenticeship in which training is initially very low, here overwork and slow training are conflicting sources of rents. Thus the optimal apprenticeship length is a compromise between raising the marginal productivity of apprentice's effort and keeping the apprentice longer.

Figure 3 depicts an optimal contract when there is only skilled effort. As in the unskilled case, as time passes and the principal is in less of a rush to train the agent, the effort distortion falls. In the example in the figure, because of a relatively strong complementarity between knowledge and effort, the contract is substantially shorter than in the unskilled case. Note that if the maximum feasible effort were reduced to $a^{FB}(\bar{X})$ the agent would endure an apprenticeship that is less costly per unit of time, but since the marginal value of knowledge falls, the principal would also make the apprenticeship longer. Thus overwork in the skilled task allows the principal to extract rents from the agent but has the positive countervailing effect of shortening the apprenticeship. Consequently, a ban on overwork (that is, a cap of $a_t \leq a^{FB}(\bar{X})$) can either help or hurt the agent depending on the details of the technology.

C. Both efforts. Here we specialize to the case where unskilled effort has constant returns, namely $g(b) = qb$ for some constant $q > 0$. Constant returns imply that whenever the principal specifies positive effort on both tasks, the marginal productivity of each task is $q = f_a(X_t^*, a_t^*)$, which for any given X_t^* pins down a_t^* .

As time goes by and the skilled task becomes more productive, the optimal contract in general goes through three regimes: First, over an initial (possible empty) time interval $(0, t_1)$ skilled effort is efficient (possibly zero) and unskilled effort is inefficiently high. Since $\rho_t = 1$ this regime is qualitatively similar to the special case where only unskilled effort is used. Second, over an intermediate (possible empty) time interval (t_1, t_2) both efforts are inefficiently high. Third, over a final (possible empty) time interval (t_2, T^*) skilled effort is inefficiently high and unskilled effort is first best and equal to zero. This

¹⁸We can show by example that each case can occur.



$$y = Xa \text{ and } c = a^2$$

Figure 3: skilled effort only. $\frac{1}{r} = 20$ years.

regime is qualitatively the same as the special case where only skilled effort is used.¹⁹

Note that the model's predictions about the time path of effort fit the way effort evolves in the examples of Section 2: Novices ranging from young cooks to manicurists go through well-defined career stages, initially carrying out the least desirable tasks (sweeping floors, cleaning vegetables, finding typos in spreadsheets) and then gradually progressing into more sophisticated activities.²⁰

To characterize the optimal contract length, we define the threshold $\hat{q} = f_a(\bar{X}, \hat{a})$ for $\hat{a} = \arg \max_a f(\bar{X}, a) - c(a)$. When $q \geq \hat{q}$ the unskilled task is sufficiently productive

¹⁹That the optimal contract in general goes through these regimes follows from the fact that total effort is no lower than first best, X_t^* is increasing over time, $f_{Xa} > 0$, and effort is allocated efficiently between the two tasks.

²⁰Under our general maintained assumptions, unskilled effort weakly falls over time (strictly so whenever it is positive and lower than the effort upper bound). This follows from the fact that, over time, the total effort distortion falls and the agent's knowledge grows, each of which (weakly) lowers unskilled effort.

that only regime 1 occurs. In this case, as in the case where only unskilled effort is used, the unconstrained apprenticeship length is $\frac{1}{r}$. When instead $q < \hat{q}$ the unskilled task is sufficiently unproductive that regime 3 is always non-empty. In this case, as in the case when only skilled effort is used, the unconstrained apprenticeship length is strictly less than $\frac{1}{r}$. Notice that regimes 1 and 2 are especially damaging to the agent because the overwork in the unskilled task does not have the positive countervailing effect of shortening the apprenticeship.

7 Extensions

7.1 Training costs

Here we extend the model to include a training cost. Specifically, we suppose the principal incurs cost $k \geq 0$ per unit of value $\frac{1}{r}v(X)$ acquired by the agent, and so the principal's flow payoff is now

$$\bar{y}(X, l) = \underbrace{k \frac{d}{dt} \left[\frac{v(X)}{r} \right]}_{\text{instantaneous training cost}} .$$

This simple functional form allows us to again provide a closed-form solution and obtain clean comparative statics. We assume that $k < 1$ so that it is efficient to transfer all knowledge, that is, the surplus net of the training cost, $v(X) - k[v(X) - v(\underline{X})]$, is maximized at \bar{X} . To guarantee positive profits for the expert we assume that the agent's fixed outside option \underline{v} is no greater than $v(\underline{X})$.

As we now explain, the solution to this problem is qualitatively similar to that of the base model; the main difference is that the training cost causes the principal to slow down training and overwork the agent over a longer period of time. The formal analysis of this claim is derived in Appendix A3; here we summarize the main points.

The first thing to note is that the optimal contract continues to satisfy the conclusions of Lemmas 1 and 2, so that no wages are paid, the agent is fully trained in finite time, the participation constraint holds with equality at all times except perhaps for an initial gift of knowledge, and total effort is allocated efficiently across the two tasks. Moreover, except for why the agent is fully trained, the intuition for these results is the same as in the original model. To see why the agent is fully trained suppose the agent currently has knowledge $X < \bar{X}$ and the principal offers to train the agent over an additional unit of

time while asking for first-best effort $l^{FB}(X)$. The agent's participation constraint holds with equality if $\frac{d}{dt} \left[\frac{1}{r} v(X) \right] = v(X) + c(l^{FB}(X)) = \bar{y}(X, l^{FB}(X))$. Therefore, net of the training cost, this arrangement delivers profit

$$\bar{y}(X, l^{FB}(X)) - k \frac{d}{dt} \left[\frac{1}{r} v(X) \right] = (1 - k) \bar{y}(X, l^{FB}(X)) > 0,$$

which is strictly positive because $k < 1$. This shows that the principal can pocket all of the surplus generated by the additional training by keeping the agent indifferent while being trained, and keeping effort at first best so that no surplus is wasted.

While the training cost does not alter the general form of the contract, it does change the apprenticeship length, which now satisfies

$$r \int_0^{T^*} [\rho_t - k] dt = 1.$$

This formula is almost the same condition as in the original model, except that $[\rho_t - k]$ takes the place of ρ_t . To understand why this is so, start with an apprenticeship with length T and as we did in the original model suppose the principal has the agent work a bit longer at the end of the apprenticeship, so that the principal gains 1 in present value. This change lowers the agent's continuation value throughout the apprenticeship, and so the principal needs to lower the agent's knowledge at each t so they do not walk away. As a result, the principal suffers output loss $r \int_0^T \rho_t dt$ as before, but also postpones some of the training cost and therefore her overall training costs change by²¹

$$\int_0^T e^{-rt} k \frac{d}{dt} \underbrace{\left[\frac{1}{r} v'(X_t) dX \right]}_{\text{change in value } \frac{1}{r} v(X_t)} dt = -r \int_0^T k dt.$$

The optimal length sets the output loss net of cost savings, $r \int_0^T [\rho_t - k] dt$, equal to 1.

Because of the principal's desire to backload the training cost, the apprenticeship lasts longer than in the original model. To illustrate, when the agent exerts unskilled effort only, and so $\rho_t \equiv 1$, the optimal unconstrained length is $\frac{1}{r(1-k)} > \frac{1}{r}$. This length is increasing in k because the larger the cost, the more the principal wants to postpone paying it; and is

²¹This equality follows from the fact that, to keep the agent from walking away, the principal must set $e^{rt} \frac{d}{dX_t} \left(\frac{1}{r} v(X_t) \right) = -1$.

decreasing in r , as in the original model, because as players become less patient knowledge becomes less valuable, and so the agent is not willing to work as long to acquire it.

As before, the agent is asked to exert inefficiently much effort except at the terminal time. The target effort distortion is now

$$\frac{c'(l_t)}{\frac{\partial \bar{y}}{\partial l}(\bar{X}, l_t)} = \frac{1}{1 - r \int_t^{T^*} [\rho_\tau - k] d\tau}$$

where $[\rho_\tau - k]$ takes the place of ρ_τ because a greater distortion raises the rate of knowledge transfer and so has the disadvantage of frontloading the training costs that remain to be paid. To illustrate, when the agent exerts unskilled effort only and the apprenticeship length is unconstrained, the target distortion is

$$\frac{c'(l_t)}{\frac{\partial \bar{y}}{\partial l}(\bar{X}, l_t)} = \frac{T^*}{t}.$$

This distortion depends only on the fraction of time that remains in the apprenticeship. As a result, as k grows and the apprenticeship becomes longer, the effort path is very similar to that in the original model, but is spread out over a longer period of time.

7.2 Training certificates (and indentured servitude)

So far we have assumed that the agent is unable to commit to keep working for the principal after being trained. If instead the agent had full commitment power, the optimal contract would immediately fully train the agent, specify the corresponding first-best level of effort, and require that the agent works for the principal for a time interval just long enough to extract the full value of all knowledge from the agent. The many complaints about slow training and excess effort that we discussed in Section 2 suggest that in practice the agent commonly does not have this sort of commitment ability. Nevertheless, in some situations the agent's outside opportunity is lower than $v(X)$ unless they are provided with a certificate of completion, occupational license, or letter of recommendation from the principal. Here the agent's desire to be certified in effect makes human capital at least partially firm-specific.

In the extreme case where the agent's outside option without a certificate is \underline{v} regardless of their level of training, the principal can implement the full-commitment solution with immediate training and first-best effort at all times. More generally, if the certificate adds

a fixed amount Δ to the agent's outside option, then when the agent has knowledge X their value is $\frac{1}{r}v(X)$ with a certificate and $\frac{1}{r}\max\{v(X) - \Delta, \underline{v}\}$ without it. In this case, as we show in the online appendix the optimal contract has two phases.²² Phase 1 resembles the solution for $\Delta = 0$: Here knowledge grows over time, the agent's participation constraint binds at each instant, and the agent works inefficiently hard until the last instant of the phase, which occurs when the agent is fully trained (that is when $X_t = \bar{X}$). Phase 2 corresponds to the solution for large Δ : Here the fully-trained agent exerts first-best effort and earns zero wages over a time interval just long enough to extract the full value Δ/r of the certificate from the agent.

Thus when Δ is small, the solution is very similar to the solution in our main model; the main difference is that the "graduation prize" for phase 1 is $v(\bar{X}) - \Delta$ instead of $v(\bar{X})$, and the differential version of the participation constraint is now $\dot{u}_t = r[u_t + c(l_t) - \Delta]$ instead of $\dot{u}_t = r[u_t + c(l_t)]$, where $u_t = v(X_t)$. As Δ grows, the agent is trained more quickly (phase 1 shrinks) but is kept longer after being trained (phase 2 grows). Hence, a more valuable certificate (e.g. a more demanding occupational licence) raises both total surplus and profits, but is potentially damaging to the agent.

Indentured servitude agreements may play a role similar to certificates. For instance, if the principal can threaten to impose a penalty D on an agent who walks away from the relationship, then the parties can enter a servitude agreement whereby, after being trained, the agent promises to continue working for the principal until giving up value D . This case is identical to the above case of a certificate with $\Delta = rD$. Indentured apprenticeships date back to medieval times, where apprentices were commonly bound to their masters for a number of years (e.g. Thrupp, 1989). Arguably, these arrangements are echoed in some modern apprenticeships where servitude contracts have been replaced with (more benign) certification requirements.²³

²²Online Appendix available at <http://economics.mit.edu/faculty/drewf>.

²³Indentured servitude has also been widely used, both historically and in modern times, to finance the migration of credit-constrained workers (e.g. Galenson, 1984, and Guido and Guriev, 2006). In modern times, illegal workers and their employers enter agreements where the worker spends time in a sweatshop, under coercion, until the agreed debt has been repaid. Our model suggests that such debt may in principle include not only the cost of smuggling the worker, but also the value of any knowledge acquired in the sweatshop.

7.3 Cash payments

In the baseline model the agent starts out with no cash and cannot borrow from a third party. If the agent instead begins the relationship with a cash balance M that is known to the principal, then the optimal contract is as follows (where we have assumed that $v(\underline{X}) \geq \underline{v}$ without loss):²⁴

1. If M is lower than the value of the optimal level of the knowledge gift $[v(X_0^*) - v(\underline{X})]/r$ in the baseline model, then the principal charges the agent M upfront for the right to enter the apprenticeship, and otherwise leaves the apprenticeship unaffected.
2. If M is higher than $[v(X_0^*) - v(\underline{X})]/r$, but lower than $[v(\bar{X}) - v(\underline{X})]/r$, then the principal charges the agent M and, in return, offers the agent a Pareto-efficient apprenticeship in which the agent earns payoff $U = v(\underline{X})/r + M$. Recall that this apprenticeship is the truncated version of the original apprenticeship where the agent is given a larger initial knowledge injection, and therefore starts out farther along the training path.
3. If M is higher than $[v(\bar{X}) - v(\underline{X})]/r$, then the principal sells the agent all knowledge upfront, at a price equal to its full value.

Thus we see that a higher cash level (weakly) shortens the apprenticeship, allows the agent to avoid the worst of the effort distortions (and the worst of the menial work), and raises total surplus. However, the principal's monopoly power means that, net of the cash payment, the agent does not benefit, and may even lose, from having access to this cash.

The proof for this result is as follows. Suppose first that M is lower than $[v(\bar{X}) - v(\underline{X})]/r$. Because the principal has commitment power, it is without loss to ask the agent to surrender M at the beginning of the relationship. The optimal contract then maximizes profits (above and beyond M) subject to the agent's ex-ante participation constraint $U \geq v(\underline{X}) + M$. Since raising M is equivalent to raising \underline{X} , the result follows from Theorem 2. Finally, if M is greater than $[v(\bar{X}) - v(\underline{X})]/r$, then the principal can obtain profits equal to the first-best surplus by selling all knowledge up front.

²⁴Garicano and Rayo (2017) derive a special instance of this result for the case where there is no effort choice.

As noted in Section 2, Adam Smith observed that during the industrial revolution, masters asked their novices for up-front cash payments for the right to enter an unpaid apprenticeships, and novices who could not make this payment served longer apprenticeships, as our model predicts. Similarly, in medieval apprenticeships, “Lack of payment could be made an excuse for prolonging the term of service [...] Unfortunately for the masters, the supply of labor able to pay for apprenticeship fell short of demand.” (Thrupp, 1989, 215). And as an example in modern times, aspiring New York manicurists are asked to pay to enter a type of unpaid apprenticeship, and sometimes also to acquire additional skills (“\$100 for eyebrow waxing, \$100 to learn how to apply gel and cure it with ultraviolet light ”), with discounts possible after a long enough service (Maslin Nir, 2015).

7.4 Regulating apprenticeships

Recall that any Pareto-efficient payoffs (Π, U) are uniquely implemented by an initial knowledge gift to raise the agent’s knowledge to $v^{-1}(rU)$, followed by the principal’s optimal contract when the agent starts with that knowledge level, and that along the Pareto frontier total surplus grows in U . Thus welfare would be increased if the regulator could mandate an increase in U , by for example raising the agent’s exogenous outside option. More realistically, the regulator might consider a cap on the length of the apprenticeship, a cap on effort, or a combination of the two.

A cap on apprenticeship length accelerates knowledge transfer, but also leads the principal to further distort effort, and even distort effort at the terminal date, in order to sell her knowledge more quickly.²⁵ For this reason, unless the cap is tight enough that the principal chooses to give away most of her knowledge, its impact on surplus is ambiguous. (Notice that because there is a limit on instantaneous effort, a very tight cap leads to almost the first-best outcome.) However, if the principal faces a cost when training the agent, a sufficiently tight cap may simply drive the principal away from the market.

A cap on instantaneous effort reduces the effort distortions at a given knowledge level of the agent, but may lead the principal to lengthen the apprenticeship, so here again the impact on total surplus is ambiguous. This lengthening of the apprenticeship occurs whenever the effort cap leads to a reduction in skilled effort, and therefore a reduction in the knowledge premia ρ . It also occurs whenever the agent’s initial knowledge is high

²⁵This can be seen from the first-order condition for l_t in the proof of Lemma A1, part 2.

enough that the agent’s ex-ante participation constraint binds, even when the effort cap does not affect skilled effort, as in this case there is no initial knowledge gift and so the effort cap simply slows down the rate at which knowledge is transferred.

One way to achieve a Pareto efficient allocation is to combine a cap on apprenticeship length with a time-varying cap on effort that mimics the Pareto-efficient effort path given the desired contract length, but this is unrealistic. It does however seem plausible that regulations could combine a time-invariant effort cap with a limit on the training period. This intervention will still not lead to Pareto-efficient contracts; but depending on the welfare weights used, can more often lead to higher total surplus. The reason is that a cap on contract length limits the principal’s ability to extend the apprenticeship in response to the effort cap and, at the same time, the cap on effort limits the principal’s ability to overwork the agent in response to the cap on contract length.

7.5 Bounded training speed

Our analysis allows the principal to transfer knowledge to the agent arbitrarily quickly, and when the agent’s initial knowledge is sufficiently low compared to the principal’s stock, the principal takes advantage of this flexibility to instantly train the agent to the top of the Pareto frontier (as seen in Figure 1). After that, the principal transfers knowledge at a bounded rate.²⁶ As illustrated in our numerical examples this rate can be very slow, leading to apprenticeships between 10 and 20 years long when $r = 5\%$.

If we now suppose there is a bound on the speed of knowledge transfer, then provided this bound is not too tight, the solution will be exactly as described earlier except that the principal will initially transfer knowledge to the agent as quickly as possible, while asking the agent to work as hard as possible, until the state reaches the level that maximizes the principal’s payoff on the Pareto frontier. And after that, the optimal contract will be identical to the baseline model.

Thus we see that the principal’s strategic motive for slowing down the knowledge transfer, rather than a constraint on training speed, may be what determines the transfer rate throughout most of the apprenticeship. This finding appears to be consistent with Adam Smith’s view that industrial-revolution apprenticeships (typically seven years long) were excessively lengthy and a way to take advantage of cheap labor. It also seems

²⁶Upon normalizing the units of knowledge so that $v(X) = X$, the rate of knowledge transfer is given by the state equation $\dot{X}_t = r[X_t + c(l_t)]$, and so this rate is bounded above by $r[\bar{X} + c(1)]$.

consistent with the examples in Section 2, where apprenticeships can be quite long (e.g. 10 years in Jiro’s case) and trainees spend a considerable fraction of this time on menial tasks, rather than on learning new skills.

8 Conclusion

To conclude, we briefly review our main findings. We have considered the optimal contract for a principal with commitment power to “sell” knowledge to a cash-constrained agent, or apprentice, who is free to walk away at any time. In these contracts, the agent works for the principal for low or no wages. Moreover, the principal requires the agent to work inefficiently hard. When the production function leads the principal to require excess effort in the skilled task, the period of apprenticeship decreases, while if the principal only ever requires excess effort in the unskilled task, the length of the apprenticeship is unaffected by the degree to which the agent is overworked. In some (but not all) cases, regulations that cap the agent’s maximum effort can raise surplus; effort caps combined with limits on the duration of apprenticeship can do even better.

These results follow from our assumption that the agent is unable to commit to keep working for the principal after being trained. If the agent has full commitment power, the optimal contract will immediately fully train the agent, and specify the corresponding first-best level of effort. The many complaints about slow training and excess effort that we discussed in the introduction suggest that in practice the agent commonly does not have this sort of commitment ability.

Finally, we should point out that we have abstracted away from the idea that the agent learns by doing, so that the rate of knowledge transfer depends on the amount of skilled effort, and also abstracted away from the possibility that the agent, principal, or both, are learning about the agent’s ability over time. We have also assumed that there is only a single potential agent. If the principal can only train one (or a small number) of agents at a time, then training a given agent has an opportunity cost, and in some cases this might lead to “incomplete training,” that is the principal might switch to training a new agent before the current one acquires all of the principal’s knowledge. All of these are important aspects of some apprenticeship relationships, and we plan to explore them in future work.²⁷

²⁷In ongoing work we also extend our analysis to the case where the agent’s utility of consumption is

9 Appendix A1: Proof of Lemmas 1-3, and Proposition 1

Proof of Lemma 1. The conclusion of the lemma will follow from a series of claims.

Claim 1 *The principal obtains a strictly positive profit by contracting with the agent.*

Proof. Fix $X' \in (\underline{X}, \bar{X})$, a' , and b' s.t. $y(X', a', b') > \underline{v}$ (which is feasible because $v(\bar{X}) > \underline{v}$) and then pick $T' > 0$ s.t. $e^{-rT'} v(\bar{X}) - (1 - e^{-rT'})c(a' + b') > v(X')$. Now consider the contract where the principal pays 0 wages, brings the agent's knowledge up to X' at time 0, and asks them to maintain efforts (a', b') until time T' , at which point the principal brings the agent's knowledge up to \bar{X} . This contract satisfies the agent's participation constraint (2) and the liquidity constraint (3), and gives the principal a positive payoff. ■

This proves part 1 of the lemma.

Claim 2 *Any contract where $W_\infty > 0$ is strictly dominated by some finite-duration contract where $W_\infty = 0$.*

Proof. If contract S with potentially infinite graduation date T prescribes $W_\infty > 0$ and is not strictly dominated, by the previous claim it must have $\Pi_0(S) > 0$, so $U_0(S) < \frac{1}{r}v(X_\infty)$. Now let $T' \in (0, T)$ satisfy

$$\left(e^{-rT'} - e^{-rT}\right) \frac{1}{r}v(X_\infty) + \int_{T'}^T e^{-rt} c(a_t + b_t) dt = W_\infty,$$

and consider a new contract S' where the agent earns zero wages, graduates at date T' with knowledge X_∞ , and for $t < T'$,

$$X'_t, a'_t, b'_t = X_t, a_t, b_t.$$

By construction,

$$\begin{aligned} U_0(S') &= e^{-rT'} \frac{1}{r}v(X_\infty) - \int_0^{T'} e^{-rt} c(a_t + b_t) dt \\ &= e^{-rT} \frac{1}{r}v(X_\infty) + W_\infty - \int_0^T e^{-rt} c(a_t + b_t) dt = U_0(S). \end{aligned}$$

strictly concave. Preliminary results are available upon request.

In addition, for $t < T'$,

$$\begin{aligned} U_t(S') - U_t(S) &= \left(e^{-r(T'-t)} - e^{-r(T-t)} \right) \frac{1}{r} v(X_\infty) \\ &\quad + \int_{T'}^T e^{-r(\tau-t)} c(a_\tau + b_\tau) d\tau - \int_t^T e^{-r(\tau-t)} w_\tau d\tau \\ &= e^{rt} \left[W_\infty - \int_t^T e^{-r\tau} w_\tau d\tau \right] \geq 0. \end{aligned}$$

As a result, since the original contract satisfied (2), the new contract satisfies (2) as well. And since the new contract prescribes zero wages, it satisfies (3).

Finally, we have

$$\begin{aligned} \Pi_0(S') + U_0(S') - [\Pi_0(S) + U_0(S)] \\ \geq \int_{T'}^T e^{-rt} [v(X_\infty) - v(X_t)] dt > 0, \end{aligned}$$

where the strict inequality follows from the facts that v is strictly increasing and that $X_t < X_\infty$ for all $t \in (T', T)$. Since $U_0(S') = U_0(S)$, it follows that $\Pi_0(S') > \Pi_0(S)$, and so S' strictly dominates S . ■

This proves the first clause in part 2 of the lemma.

Claim 3 *Any infinite-duration contract is strictly dominated by some finite-duration contract with $W_\infty = 0$.*

Proof. In any infinite-duration contract, constraint (2) at time 0 requires $W_\infty \geq \frac{1}{r} \max\{\underline{v}, v(\underline{X})\} > 0$, so the contract is strictly dominated by the previous claim. ■

Claim 4 *Any finite-duration contract with $W_\infty = 0$ and $X_T < \bar{X}$ is strictly dominated by some finite-duration contract with $W_\infty = 0$ and $X_T = \bar{X}$.*

Proof. If a finite-duration contract with $W_\infty = 0$ has $X_T < \bar{X}$, then there is a time interval Δ and effort levels a', b' such that $y(X', a', b') > \underline{v}$ and $e^{-r\Delta}v(\bar{X}) - (1 - e^{-r\Delta})c(a' + b') > v(X_T)$, and so the principal could obtain strictly higher profits by extending the agent's contract to $T' = T + \Delta$ paying no additional wages, setting $X_t = X_T$ and $(a_t, b_t) = (a', b')$ for $t \in [T, T')$, and setting $X_{T'} = \bar{X}$. ■

Claims 3 and 4 prove the second clause in part 2 of the lemma.

Claim 5 *Any contract is weakly dominated by some finite-duration contract with $X_T = \bar{X}$ and zero wages.*

Proof. From Claims 3 and 4, we can restrict to finite-duration contracts such that $X_T = \bar{X}$ and $W_\infty = 0$. Let S be one such contract, and consider an alternative contract S' that is identical to S except for the fact that all wages are zero.

The two contracts deliver identical profits. In addition, for all t ,

$$\begin{aligned} e^{-rt} [U_t(S) - U_t(S')] &= \int_t^T e^{-r\tau} w_\tau d\tau \\ &= W_\infty - W_t \leq 0, \end{aligned}$$

where the inequality follows from the fact that $W_t \geq 0$ (from (3)) and $W_\infty = 0$. As a result, $U_t(S') \geq U_t(S)$ and therefore S' satisfies (2) and (3) as well. ■

This proves the third clause in part 2 of the lemma and so completes its proof.

Proof of Lemma 2. We will show each clause of the lemma in turn.

Claim 6 *Any contract is weakly dominated by a contract that sets the agent's participation constraints to hold with equality.*

Proof. In a contract with zero wages, $U_t = e^{-r(T-t)} \frac{1}{r} v(\bar{X}) - \int_t^T e^{-r(\tau-t)} c(a_\tau + b_\tau) d\tau$, which is strictly increasing (because $v(\bar{X}) > 0$) and continuous. Thus if $U_t > \frac{1}{r} v(X_t)$ for some times t , the contract with the same effort path and terminal date, and $X'_t = \max\{X_t, v^{-1}(rU_t)\}$ at all times will satisfy the participation constraints and give the principal a weakly higher payoff at each date. Moreover, if the times where $U_t > \frac{1}{r} v(X_t)$ had positive measure, the new contract would give the principal a strictly higher payoff overall. ■

Claim 7 *Any contract is weakly dominated by a contract where at each t total effort $a_t + b_t$ is allocated across tasks to maximize output.*

Proof. Given any contract where at some times $y_t(X_t, a_t, b_t) \neq \bar{y}(X_t, (a_t + b_t))$, consider the alternative contract where the time paths of knowledge and total effort are the same but effort is allocated to maximize output at each time. Since the agent's knowledge stock and effort cost are the same, the participation constraints are still satisfied, and

the principal does at least as well, and strictly better if the times where $y_t(X_t, a_t, b_t) \neq \bar{y}(X_t, (a_t + b_t))$ had positive measure. ■

This completes the proof of Lemma 2.

Proof of Proposition 1. Suppose an optimal contract exists (otherwise the proposition is vacuously true). Now suppose S^* is optimal and contrary to the proposition suppose there is a date $0 < z < T^*$ and a contract S^{**} , with $X_z^{**} \geq X_z^*$, such that S^{**} delivers strictly higher profits than S^* from z onward, while satisfying the participation constraints (2) from z onward. From Lemma 2, S^* satisfies all participation constraints with equality, and therefore there is a new contract \hat{S} that is identical to S^* for all $0 \leq t < z$, and identical to S^{**} for all $z \leq t \leq T^*$, that leaves the agent weakly better off than S^* , and so satisfies all participation constraints and delivers strictly higher profits than S^* – a contradiction. ■

Proof of Lemma 3. The agent won't work longer than $T^{\max} := \frac{1}{r} \log \left(\frac{v(\bar{X})}{v(\underline{X})} \right)$ even if asked to exert 0 effort for all t . For any fixed $T \in [0, T^{\max}]$ a solution to (6) exists from Kumar (1969), and an optimal T exists because the principal's optimized profits are continuous in T .

To show the necessity of the system (7)-(10), let time run in reverse from T to 0, let u_T denote the fixed initial state, and change the signs of \dot{u} and of the co-state evolution equation (which is now $\dot{\lambda}_t = \partial \mathcal{L} / \partial u_t$). Any $T \in [0, T^{\max}]$ can be implemented by some choice of controls, and the inequality constraints on a_t and b_t are linearly independent, so the necessity of conditions (7)-(10) follows from Chachuat (2007) Theorems 3.18 and 3.33, Remark 3.19 (which notes that reachability is sufficient for the regularity constraint on the terminal conditions), and Remark 3.23 (on extending to inequality constraints on the terminal condition T). To obtain the transversality condition for λ_0 in condition (10), assign to constraint $u_0 \geq v(\underline{X})$ a multiplier ζ with associated complementary-slackness condition $\zeta \geq 0$ and $\zeta [u_0 - v(\underline{X})] = 0$. Chachuat's Theorem 3.18 implies that $\lambda_0 = \zeta$ and so $\lambda_0 \geq 0$ and $\lambda_0 [u_0 - v(\underline{X})] = 0$. ■

10 Appendix A2: Proof of Lemma 4 and Theorems 1 and 2

We begin by deriving some general properties of every solution to problem (6). Recall the system of necessary conditions in Lemma 3:

$$\left. \begin{aligned} \dot{u}_t &= ru_t + rc(l_t); \quad \dot{\lambda}_t = \partial \mathcal{L} / \partial u_t \\ \partial \mathcal{L} / \partial l_t &= 0 \\ \eta_t, \gamma_t &\geq 0; \quad \eta_t [1 - l_t] = \gamma_t l_t = 0 \end{aligned} \right\} \text{ for } 0 \leq t \leq T,$$

$$\mathcal{H}_T = 0; \quad \lambda_0 \geq 0; \quad \lambda_0 [u_0 - v(\underline{X})] = 0,$$

where $\mathcal{H} = e^{-rt} \bar{y}(\phi(u_t), l_t) - \lambda_t \dot{u}_t$ and $\mathcal{L} = \mathcal{H} + \eta_t [1 - l_t] + \gamma_t l_t$. Let $\rho_t := \rho(\phi(u_t), a_t) = \frac{f_X(X, a_t)}{f_X(\bar{X}, a^{FB}(\bar{X}))} \Big|_{X=\phi(u_t)} \geq 0$ and recall from Lemma 2 that, for all t , $a_t = \bar{a}(X_t, l_t) := \arg \max_{a \in [0, l_t]} f(X_t, a) + g(l_t - a)$.

Lemma A1

1. *The co-state evolution equation can be written as*

$$\lambda_t = e^{-rt} \left[e^{rT} \lambda_T - \int_t^T \rho_\tau d\tau \right]. \quad (16)$$

2. *In every solution $\lambda_T = e^{-rT} \frac{1}{r}$, $l_T = l^{FB}(\bar{X})$, and $\eta_T = 0$.*

Proof. *Part 1.* Using the facts that $v'(\phi(u)) = f_X(\phi(u), a^{FB}(\phi(u)))$ (from the envelope theorem) and $\phi'(u_t) = \frac{1}{v'(\phi(u))}$ (from the implicit function theorem) we obtain $\frac{d}{du} f(\phi(u), a) = f_X(\phi(u), a) \frac{1}{v'(\phi(u))} = \frac{f_X(X, a)}{f_X(\bar{X}, a^{FB}(\bar{X}))} \Big|_{X=\phi(u)}$. Thus the co-state evolution equation is $\dot{\lambda}_t = -r\lambda_t + e^{-rt} \frac{d}{du} f(\phi(u_t), a_t) = -r\lambda_t + e^{-rt} \rho_t$, which is equivalent to (16).

Part 2. The first-order condition for T is $e^{-rT} \bar{y}(\phi(u_T), l_T) - \lambda_T r [u_T + c(l_T)] = 0$ and the first-order condition for l_T implies that $\lambda_T r = [e^{-rT} \frac{\partial}{\partial l} \bar{y}(\bar{X}, l_T) - \eta_T] / c'(l_T)$. Substituting this into the first-order condition for T yields

$$\underbrace{c'(l_T) \bar{y}(\bar{X}, l_T) - \frac{\partial}{\partial l} \bar{y}(\bar{X}, l_T) [v(\bar{X}) + c(l_T)]}_{=: h(l_T)} = -e^{rT} \eta_T [v(\bar{X}) + c(l_T)]. \quad (17)$$

Notice that $h(l_T) = 0$ when $l_T = l^{FB}(\bar{X})$. In addition, since c , c' , and \bar{y} are all differentiable in l , and $\frac{\partial \bar{y}}{\partial l}$ is continuous and almost-everywhere differentiable in l , the function $h(l_T)$ is continuous, almost-everywhere differentiable, and at each point of differentiability,

$$h'(l_T) = c''(l_T) \bar{y}(\bar{X}, l_T) - \frac{\partial^2}{\partial l^2} \bar{y}(\bar{X}, l_T) [v(\bar{X}) + c(l_T)] > 0,$$

where the inequality follows from the fact that $c'' > 0$ and at each point of differentiability $\frac{\partial^2}{\partial l^2} \bar{y} \leq 0$ (since $f_{aa}, g'' \leq 0$). As result, $h(l_T)$ is strictly increasing. It follows that $\eta_T = 0$; otherwise, we would have $\eta_T > 0$ and $l_T = 1 > l^{FB}(\bar{X})$, and so the left-hand side of (17) would be positive and its right-hand side would be negative. Once we set $\eta_T = 0$ it follows that $l_T = l^{FB}(\bar{X})$ and $\lambda_T = e^{-rT} \frac{1}{r}$. ■

Lemma A2 For any given λ_t the first-order condition for l_t and complementary slackness conditions for η_t, γ_t have a unique solution, denoted $\tilde{l}(\lambda_t), \tilde{\eta}(\lambda_t), \tilde{\gamma}(\lambda_t)$. Moreover, this solution satisfies $\tilde{\gamma}(\lambda_t) = 0$ and

$$\begin{aligned} c'(\tilde{l}(\lambda_t)) &= \frac{e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), \tilde{l}(\lambda_t)) - \tilde{\eta}(\lambda_t)}{\lambda_t r}, \\ \tilde{\eta}(\lambda_t) &= \max \left\{ 0, e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), 1) - \lambda_t r c'(1) \right\}. \end{aligned} \quad (18)$$

Proof. The first-order condition for l_t is $e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), l_t) - \lambda_t r c'(l_t) - \eta_t + \gamma_t = 0$. This has a unique solution because $c'' > 0$ and \bar{y} is concave in l_t . The optimal l_t must be strictly positive because $\frac{\partial}{\partial l} \bar{y}(\phi(u_t), 0) > c'(0)$. If $c'(1) < \frac{e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), 1)}{\lambda_t r}$, the solution is that $l_t = 1$ and then $\eta_t = e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), 1) - \lambda_t r c'(1)$. If $c'(1) \geq \frac{e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), 1)}{\lambda_t r}$ then the constraint $l_t \leq 1$ is slack, so $\eta_t = 0$ and $c'(\tilde{l}(\lambda_t)) = \frac{e^{-rt} \frac{\partial}{\partial l} \bar{y}(\phi(u_t), \tilde{l}(\lambda_t))}{\lambda_t r}$. ■

We are now ready to prove Lemma 4 and Theorems 1 and 2.

Proof of Lemma 4

The time 0 co-state λ_0 measures the shadow cost of $u_0 = U/r$; hence, $-\frac{d\Pi^*(U)}{dU} = -r \frac{d\Pi^*(U)}{du_0} = r \lambda_0$. Moreover, from Lemma A1, $r \lambda_0$ equals $1 - r \int_0^{T^*} \rho_t^* dt$.

Proof of Theorem 1

Write the co-state equation (16) as

$$\lambda_t = e^{-rt} \left[\frac{1}{r} - \int_t^T \rho_\tau d\tau \right], \quad (19)$$

where we used the fact that $e^{rT} \lambda_T = \frac{1}{r}$. After substituting for λ_t into equation (18) we obtain

$$c'(l_t) = \frac{\frac{\partial}{\partial l} \bar{y}(X_t, l_t) - e^{rt} \eta_t}{1 - r \int_t^T \rho_\tau d\tau}, \quad (20)$$

$$e^{rt} \eta_t = \max \left\{ 0, \frac{\partial}{\partial l} \bar{y}(X_t, 1) - \left(1 - r \int_t^T \rho_\tau d\tau \right) c'(1) \right\}.$$

Since $\lambda_0 \geq 0$ implies that $1 - r \int_0^T \rho_t dt \geq 0$, equation (20) implies that the optimal effort path satisfies $\frac{c'(l_t^*)}{\frac{\partial}{\partial l} \bar{y}(X_t^*, l_t^*)} = \min \left\{ \left(1 - r \int_t^{T^*} \rho_\tau d\tau \right)^{-1}, \frac{c'(1)}{\frac{\partial}{\partial l} \bar{y}(X_t^*, 1)} \right\}$ with $\rho_\tau = \rho(X_\tau^*, \bar{a}(X_\tau^*, l_\tau^*))$, and the state equation implies that $\frac{1}{r} \cdot \frac{d}{dt} v(X_t^*) = v(X_t^*) + c(l_t^*)$ (as claimed in part 2 of the theorem); and therefore $l_{T^*}^* = l^{FB}(\bar{X})$ and $l_t^* > l^{FB}(X_t^*)$ for all $t < T^*$ (as claimed in part 1 of the theorem). Finally, since $c'' > 0$ and $\frac{\partial \bar{y}}{\partial l}$ is weakly decreasing in l , for any given T^* the values of X_t^* and l_t^* are unique. ■

Proof of Theorem 2

We first ignore both the complementary slackness condition for λ_0 and the constraint $u_0 \geq v(\underline{X})$, and claim that for any fixed terminal time T , the remaining necessary conditions in Lemma 3, together with the terminal condition $u_T = v(\bar{X})$, have a unique solution, which we denote u_t^T, l_t^T, ρ_t^T (and λ_t^T, η_t^T). First, write $t = T - s$ for $s \geq 0$, and use (20) to obtain

$$c'(l_{T-s}) = \begin{cases} \min \left\{ \frac{\frac{\partial}{\partial l} \bar{y}(X_t, l_{T-s})}{1 - r \int_0^s \rho_{T-\tau} d\tau}, c'(1) \right\} & \text{when } r \int_0^s \rho_{T-\tau} d\tau \leq 1, \\ c'(1) & \text{otherwise,} \end{cases} \quad (21)$$

where $\rho_{T-\tau} = \rho(\phi(u_{T-\tau}), \bar{a}(X_{T-\tau}, l_{T-\tau}))$. Second, note that since $c'' > 0$ and $\frac{\partial \bar{y}}{\partial l}$ is weakly decreasing in l , there is a unique solution $u_{T-s}^T, l_{T-s}^T, \rho_{T-s}^T$ to the system

$$u_T = v(\bar{X}), \quad \dot{u}_{T-s} = r [u_{T-s} + c(l_{T-s})], \quad \text{and (21), for } s \geq 0. \quad (22)$$

(The resulting paths λ_{T-s}^T and $\eta_{T-s}^T = \tilde{\eta}(\lambda_{T-s}^T)$ are also unique). Notice that the value of T enters this system only as a subindex for $u_{T-s}, l_{T-s}, \rho_{T-s}$. Therefore, for any given $s \geq 0$ the solution $u_{T-s}^T, l_{T-s}^T, \rho_{T-s}^T$ is independent of the chosen value of T . Moreover, since $u_0^T = u_{T-s}^T|_{s=T}$, we have $\frac{d}{dT}u_0^T = -\frac{d}{dt}u_t^T|_{t=0} = -r[u_0^T + c(l_0^T)]$. Thus, whenever u_0^T is positive, it is strictly decreasing in T . Notice also that $l_t^T \geq l^{FB}(\phi(u_t))$ (since $r \int_t^T \rho_\tau d\tau \geq 0$) and therefore $a_t^T \geq a^{FB}(\phi(u_t))$ and $\rho_t^T \geq 1$.

We now show that there is a unique T^* that satisfies the complementary slackness condition for λ_0^T , namely, $\lambda_0^T \geq 0$ and $\lambda_0^T [u_0 - v(\underline{X})] = 0$, together with the constraint $u_0^T \geq v(\underline{X})$. Since $\lambda_0^T = 1 - r \int_0^T \rho_{T-s}^T ds$ is strictly decreasing in T , there is a unique \hat{T} such that $\lambda_0^{\hat{T}} = 1 - r \int_0^{\hat{T}} \rho_t dt = 0$; and since $\lambda_0^T \geq 0$, we may restrict our search to $T \leq \hat{T}$.

There are two cases to consider: (a) $u_0^{\hat{T}} \geq v(\underline{X})$, and (b) $u_0^{\hat{T}} < v(\underline{X})$. In case (a), we must have $T^* = \hat{T}$ and therefore $r \int_0^{T^*} \rho_t dt = 1$. Otherwise $T^* < \hat{T}$, $\lambda_0^{T^*} > 0$, and $u_0^{T^*} > v(\underline{X})$, violating the complementary slackness condition. In case (b), we must have $T^* < \hat{T}$, $\lambda_0^{T^*} > 0$, and $u_0^{T^*} = v(\underline{X})$. Therefore, T^* is the unique value of T such that $u_0^T = v(\underline{X})$. Notice, finally, that for any given \underline{X} , system (22) implies that case (a) arises when \bar{X} is above a threshold $\underline{X} + \Delta$, and case (b) arises otherwise. ■

11 Appendix A3: Training cost

Here we derive the optimal contract in the extended model with training costs.

Lemma A3 *In the model with training costs, the conclusions in Lemmas 1 and 2 remain valid.*

Proof. With the exception of Claims 1 and 4, it is easy to see that the proofs of Lemmas 1 and 2 extend to this case. Claim 1 states that the principal obtains a strictly positive profit by contracting with the agent. To see why this is still true, consider a contract in which $X_0 = \underline{X}$ and $X_T = \bar{X}$, and at any time $0 \leq t \leq T$ effort is $l_t = l^{FB}(X_t)$, wages are zero, and the agent receives training dX_t/dt such that

$$v'(X_t) \frac{dX_t}{dt} = r [v(X_t) + c(l^{FB}(X_t))] = r \bar{y}(X_t, l^{FB}(X_t)).$$

This contract satisfies the agent's participation constraints with equality at all times and delivers profits

$$\int_0^T e^{-rt} \left[\bar{y}(X_t, l^{FB}(X_t)) - k \frac{1}{r} v'(X_t) \frac{dX_t}{dt} \right] dt = \int_0^T e^{-rt} (1-k) \bar{y}(X_t, l^{FB}(X_t)) dt > 0.$$

Claim 4 states that any finite-duration contract with $W_\infty = 0$ and $X_T < \bar{X}$ is strictly dominated by some finite-duration contract with $W_\infty = 0$ and $X_T = \bar{X}$. To see why this is still true, notice that if a finite-duration contract with $W_\infty = 0$ had $X_T < \bar{X}$, then the principal could obtain strictly higher profits by extending the contract to date $T' > T$, setting $X_{T'} = \bar{X}$, and for all $T < t \leq T'$ offering the same arrangement as above. ■

It follows from this lemma that with the exception of the principal's objective, the optimal control problem is the same as in the original model. The principal's objective is now

$$\int_0^T e^{-rt} \left[\bar{y}(\phi(u_t), l_t) - k \frac{1}{r} \dot{u}_t \right] dt - k \frac{1}{r} [u_0 - v(\underline{X})],$$

where the second term in the objective is the cost of the initial gift. The Hamiltonian is now $\mathcal{H} = e^{-rt} \left[\bar{y}(\phi(u_t), l_t) - k \frac{1}{r} \dot{u}_t \right] - \lambda_t \dot{u}_t$, with $\dot{u}_t = r[u_t + c(l_t)]$. Assign to the ex-ante participation constraint $u_0 \geq v(\underline{X})$ multiplier ζ .

Lemma A4 *In the model with training costs, except for the transversal condition (10), the conclusion in Lemma 3 remains valid. The transversal condition is now $\mathcal{H}_T = 0$, $\lambda_0 = -\frac{1}{r}k + \zeta$, $\zeta \geq 0$, and $\zeta [u_0 - v(\underline{X})] = 0$.*

Proof. Define $\varphi(u_0) := -k \frac{1}{r} [u_0 - v(\underline{X})]$, $\psi(u_0) := u_0 - v(\underline{X})$ and $\Phi(u_0) := \varphi(u_0) + \zeta \psi(u_0)$. The only difference relative to the proof of Lemma 3 is that Chachuat (2007) Theorem 3.18 now requires that $\lambda_0 = \Phi'(u_0) = -\frac{1}{r}k + \zeta$. ■

Lemm A5 *In the model with training costs, the conclusions in Lemmas A1 and A2 in Appendix A2 remain valid, but with co-state equation*

$$\lambda_t = e^{-rt} \left[e^{rT} \lambda_T - \int_t^T [\rho_\tau - k] d\tau \right],$$

and with $\lambda_T = e^{-rT} \frac{1}{r} [1 - k]$.

Proof. The co-state evolution equation is $\dot{\lambda}_t = -r\lambda_t + e^{-rt}[\rho_t - k]$, the first-order condition for T is $e^{-rT} \bar{y}(\phi(u_T), l_T) - (\lambda_T r + e^{-rT} k) [u_T + c(l_T)] = 0$ and the first-order condition for l_T implies that $\lambda_T r + e^{-rT} k = [e^{-rT} \frac{\partial \bar{y}}{\partial l}(\bar{X}, l_T) - \eta_T] / c'(l_T)$. Therefore, after replacing $\lambda_T r$ with $\lambda_T r + e^{-rT} k$, the proof of this lemma is identical to the proofs of Lemmas A1 and A2. ■

Proposition A1 *In the model with training costs, the conclusions in Theorems 1 and 2 remain valid, but with the target effort distortion at time t now taking the more general form*

$$\frac{c'(l_t)}{\frac{\partial \bar{y}}{\partial l}(X_t, l_t)} = \frac{1}{1 - r \int_t^{T^*} [\rho_\tau - k] d\tau},$$

and with the optimal initial knowledge level X_0^ and contract length T^* now satisfying either*

$$X_0^* > \underline{X} \quad \text{and} \quad \int_0^{T^*} [\rho_t^* - k] dt = \frac{1}{r} \quad (\text{positive knowledge gift})$$

or

$$X_0^* = \underline{X} \quad \text{and} \quad \int_0^{T^*} [\rho_t^* - k] dt \leq \frac{1}{r} \quad (\text{zero knowledge gift}).$$

Proof. We begin with two observations. First, Lemma A5 implies that the co-state evolution equation is $\lambda_t = e^{-rt} \left[\frac{1}{r} [1 - k] - \int_t^T [\rho_\tau - k] d\tau \right]$, and therefore the effort path satisfies, for all $s \geq 0$,

$$c'(l_{T-s}) = \begin{cases} \min \left\{ \frac{\frac{\partial \bar{y}}{\partial l}(X_t, l_{T-s})}{1 - r \int_0^s [\rho_{T-\tau} - k] d\tau}, c'(1) \right\} & \text{when } r \int_0^s [\rho_{T-\tau} - k] d\tau \leq 1, \\ c'(1) & \text{otherwise.} \end{cases}$$

Second, whenever the ex-ante participation constraint is slack ($\zeta = 0$), Lemma A4 implies that $\lambda_0 = -\frac{1}{r}k$, and so the co-state evolution equation implies that $\lambda_0 = \frac{1}{r} [1 - k] - \int_0^T [\rho_t - k] dt = -\frac{1}{r}k$. Consequently, the optimal unconstrained terminal date T^* satisfies $\int_0^{T^*} [\rho_t - k] dt = \frac{1}{r}$.

It follows from these two observations that after replacing ρ_t with $[\rho_t - k]$ for all t , the proof of the present proposition is identical to the proofs of Theorems 1 and 2. ■

References

1. Acemoglu, Daron (1997). “Training and Innovation in an Imperfect Labour Market.” *Review of Economic Studies* 64(3): 445-464.
2. Acemoglu, Daron, and Jorn-Steffen Pischke (1998). “Why Do Firms Train? Theory and Evidence.” *Quarterly Journal of Economics* 113(1): 79-119.
3. Albuquerque, Rui, and Hugo A. Hopenhayn (2004). “Optimal lending contracts and firm dynamics.” *Review of Economic Studies* 71(2): 285-315.
4. American Medical Association (2015). “How many hours are in the average physician workweek?” *AMA Wire*. <https://wire.ama-assn.org/life-career/how-many-hours-are-average-physician-workweek> (accessed 6.23.2017).
5. Ayres, James (2014). *Art, Artisans and Apprentices: Apprentice Painters & Sculptors in the Early Modern British Tradition*. Oxbow Books.
6. Barlevy, Gadi, and Derek Neal (2017). “Allocating Effort and Talent in Professional Labor Markets.” *Journal of Labor Economics*, forthcoming.
7. Becker, Gary S. (1964). *Human capital: A theoretical and empirical analysis with special reference to education*. University of Chicago Press.
8. Bonatti, Alessandro, and Johannes Hörner (2017). “Career concerns with exponential learning.” *Theoretical Economics* 12(1): 425-475.
9. Chachuat, Benoit (2007). *Nonlinear and Dynamic Optimization: From Theory to Practice*. Automatic Control Laboratory, EPFL, Switzerland.
10. Cisternas, Gonzalo (2017). “Career Concerns and the Nature of Skills.” *American Economic Journal: Microeconomics*, forthcoming.
11. Coleman, Mary T., and John Pencavel (1993). “Changes in work hours of male employees, 1940–1988.” *Industrial and Labor Relations Review* 46(2): 262-283.
12. Dewatripont, Mathias, Ian Jewitt, and Jean Tirole (1999). “The economics of career concerns, part I: Comparing information structures.” *The Review of Economic Studies* 66(1): 183-198.

13. Galenson, David W. (1984). "The rise and fall of indentured servitude in the Americas: an economic analysis." *Journal of Economic History* 44(1): 1-26.
14. Garicano, Luis, and Luis Rayo (2017). "Relational Knowledge Transfers." *American Economic Review* 107(9): 2695-2730.
15. Gapper, John (2014). "Bankers and Lawyers are on an unhealthy treadmill." *Financial Times* (January 15).
16. Gergaud, Olivier, Heiner Schumacher‡ Valérie Smeets, Frédéric Warzynski (2017). "A French Recipe to Become a Top Chef: Assortative Matching and Gradual Human Capital Accumulation in Creative Industries." KEDGE - Bordeaux Business School, mimeo.
17. Guido, Friebel, and Sergei Guriev. (2006). "Smuggling humans: a theory of debt-financed migration." *Journal of the European Economic Association* 4(6): 1085-1111.
18. Holmström, Bengt (1999). "Managerial incentive problems: A dynamic perspective." *Review of Economic Studies* 66(1): 169-182.
19. Hörner, Johannes, and Andrzej Skrzypacz (2016). "Selling information." *Journal of Political Economy* 124(6): 1515-1562.
20. Jauhar, Sandeep (2015). "The Real Problem With Medical Internships." *New York Times* (Jul 1).
21. Jovinelly, Joann, and Jason Netelkos (2007). *The crafts and culture of a medieval guild*. New York: The Rosen Publishing Group.
22. Katz, Eliakim, and Adrian Ziderman (1990). "Investment in General Training: The Role of Information and Labour Mobility." *Economic Journal* 100(403): 1147-1158.
23. Kumar, T. Krishna (1969). "The existence of an optimal economic policy." *Econometrica* (Oct): 600-610.
24. Landers, Renee M., James B. Rebitzer, and Lowell J. Taylor (1996). "Rat race redux: Adverse selection in the determination of work hours in law firms." *American Economic Review* (June): 329-348.

25. Landrigan, Christopher P., et al. (2004). "Effect of reducing interns' work hours on serious medical errors in intensive care units." *New England Journal of Medicine* 351(18): 1838-1848.
26. Lane, Joan (1996). *Apprenticeship in England, 1600-1914*. London: Routledge.
27. Lugert, Verena (2017). *Die Irren mit dem Messer: Mein Leben in den Küchen der Haute Cuisine*. Knauer eBook.
28. Maister, David H. (1993). *Managing the Professional Service Firm*. Free Press, New York, NY.
29. Malcomson, James M., James W. Maw, and Barry McCormick (2003). "General training by firms, apprentice contracts, and public policy." *European Economic Review* 47(2): 197-227.
30. Maslin Nir, Sarah (2015). "The Price of Nails." *New York Times* (May 7).
31. Park, Ryan (2017). "Why So Many Young Doctors Work Such Awful Hours." *The Atlantic* (Feb 21).
32. Schwartz, Robert J., et al. (1992). "Guidelines for surgical residents' working hours: intent vs reality." *Archives of Surgery* 127(7): 778-783.
33. Smith, Adam (1872). *An Inquiry Into the Nature and Causes of the Wealth of Nations: Volume 1*. A. Murray & Company.
34. Spradlin, J. Isaac (2009). "The Evolution Of Interns." *Forbes* (Apr 27).
<https://www.forbes.com/2009/04/27/intern-history-apprenticeship-leadership-careers-jobs.html> (accessed 6.24.2017).
35. Stephan, Paula (2013). "How to exploit postdocs." *BioScience* 63(4): 245-246.
36. Stephan, Paula, and Jennifer Ma (2005). "The increased frequency and duration of the postdoctorate career stage." *American Economic Review Papers & Proceedings* 95(2): 71-75.
37. Thomas, Jonathan, and Tim Worrall (1994). "Foreign direct investment and the risk of expropriation." *Review of Economic Studies* 61(1): 81-108.

38. Thrupp, Sylvia L. (1989). *The merchant class of medieval London, 1300-1500*. Vol. 72. University of Michigan Press.
39. UK Department for Business, Innovation and Skills (2013). “Apprentices’ Pay, Training and Working Hours.” *BIS*, UK Government.