Training and Effort Dynamics in Apprenticeship[*]

Drew Fudenberg

MIT

drew.fudenberg@gmail.com

Luis Rayo

University of Utah and CEPR

Luis.Rayo@utah.edu

August 9, 2017

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.

1

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

---

[1] A 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. In some (but not all) cases, regulations that either cap the agent's maximum effort or limit the duration of the apprenticeship can raise surplus; these two policies combined can do even better.

There is an extensive literature focusing on worker training with general human capital. This literature has suppressed not only the time path of effort but the effort choice itself. 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), Holmstrom (1999), Dewatripont et al. (1999), Barlevy and Neil (2016), Bonatti and Hörner (2017), and Cisternas (2017). All of this work abstracts from worker training and knowledge transfer.

The excessive effort paths predicted by our model are consistent with the well-known empirical pattern that young professionals frequently face long hours and heavy workloads (e.g. Coleman and Pencavel, 1993, Landers et al., 1996, Barlevy and Neal, 2016) to the point where industry observers have expressed concern with their work-life balance.[3] There is also empirical support for our finding that, especially early in their careers,

---

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

[3]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).

workers are commonly instructed to perform lengthy "menial" work unrelated to the skills they wish to acquire (e.g. Maister, 1993, UK Dept. for Business, Innovation and Skills, 2013).

Medical residencies, a form of mandatory apprenticeship for young M.D.s, provide an illustration. As noted by Park (2017), they 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."[4] 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).[5]

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

---

[4]Indeed, "[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*).

[5]Schwartz 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.

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' ", and yet a legal threat by Herbert's father went nowhere because "withholding instruction was a serious threat to the apprentice" (p. 85). 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 might be granted the privilege of learning to sharpen tools" (p. 350).

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

2   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 $X_t \in [0, \overline{X}]$ 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 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 $\lim_{X \to \infty} v(X) = \infty$.

The agent starts with some exogenous stock of knowledge $\underline{X} \in [0, \overline{X})$. The agent's stock of knowledge can never decrease, and the only way it can increase is by transfers from the principal, who is able to costlessly and instantaneously increase the agent's knowledge to any level up to $\overline{X}$.[6] 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 $[\underline{X}, \overline{X}]$. Let $X_\infty := \lim_{t \to \infty} X_t$. 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 $0 < \underline{v} < v(\overline{X})$, 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.[7] 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 $t \leq T$

---

[6]This can be thought of as an idealized version of a model where the principal can costlessly transfer a flow of knowledge at a very high rate.

[7]Contracts 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.

onward are

$$\Pi_t(S) = \int_t^T e^{-r(\tau-t)} [y_\tau - w_\tau] d\tau \quad \text{and} \tag{1}$$

$$U_t(S) = \int_t^T e^{-r(\tau-t)} [w_\tau - c_\tau] d\tau + e^{-r(T-t)} \frac{1}{r} v(X_T),$$

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.

The principal's problem is

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

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 $\widehat{f} := f - \underline{w}$, $\widehat{y} := y - \underline{w}$, and $\widehat{v} := v - \underline{w}$ and model the situation using (1), (2), and (3). To allow for this interpretation of the model, we will allow $\widehat{y}$ and $\widehat{v}$ 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 $v\left(\overline{X}\right) > \underline{v}$ to be satisfied when $\underline{w}$ is very large.

3   Preliminaries

In Section 3.1 we derive general properties that any undominated contract must satisfy, without yet verifying that an optimal contract exists. In Section 3.2 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 4 uses additional assumptions to show that the optimal contract is unique and then to characterize its length and effort path.

3.1   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.[8]

Lemma 1

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 $X' \in (\underline{X}, \overline{X})$ 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.

---

[8] Recall 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.

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. Full knowledge transfer is optimal for a similar reason. Moreover, it is without loss to set wages to be zero at all times because there is no gain to the principal initially paying the agent money only to ask for it back later.

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 $\overline{y}(X, l) := \max_{a \in [0, l]} \; f(X, a) + g(l - a)$ 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 $y_t(X_t, a_t, b_t) = \overline{y}(X_t, l_t)$.*[9]

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,

---

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

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

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

subject to

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

and subject to $0 \leq l_t \leq 1$, $X_t \in \left[\underline{X}, \overline{X}\right]$, 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 undominated contract is renegotiation proof: It maximizes profits from every (t, X_t) that is reached along the optimal path.*

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.[10] This is because of the agent's dynamic participation constraint, which Lemma 2 showed is always binding. Thus, 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.

3.2 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

---

[10]In 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.

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 $\overline{X}$. Total surplus would be maximized if the principal instantaneously transferred all her knowledge to the agent, that is if $X_0 = \overline{X}$, 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 a smaller instantaneous knowledge gift to the agent, but this will only occur at the initial time.

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.

The optimal control problem is:

$$\max_{T,\{l_t\}_{t=0}^T} \int_0^T e^{-rt}\, \overline{y}(\phi(u_t), l_t)\, dt \qquad (6)$$

subject to

$$\dot{u}_t = r[u_t + c(l_t)], \quad u_T = v(\overline{X}), \quad u_0 \geq v(\underline{X}),$$

$$0 \leq l_t \leq 1.$$

Let λ_t denote the co-state variable, form the Hamiltonian $\mathcal{H} = e^{-rt}\, \overline{y}(\phi(u_t), l_t) - \lambda_t \dot{u}_t$, 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:

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

$$\partial\mathcal{L}/\partial l_t = 0 \tag{8}$$

$$\eta_t, \gamma_t \geq 0; \ \eta_t\left[1 - l_t\right] = \gamma_t l_t = 0 \tag{9}$$

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

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 conditions for the terminal time T, where λ_0 acts as a multiplier for the ex-ante participation constraint $u_0 \geq v\left(\underline{X}\right)$. 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 $u_T = v\left(\overline{X}\right)$, 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 $v\left(X_t\right)$ continues to fall until either: (1) the state reaches the agent's ex-ante outside flow payoff $v\left(\underline{X}\right)$, in which case $v\left(X_0\right) = v\left(\underline{X}\right)$ and there is no initial knowledge 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.

4   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* $\lim_{X\to\infty} f(X, 1) = \infty$.

2. $c', c'' > 0$, $c'\left(0\right) = 0$, *and* $c'\left(1\right) > \frac{\partial}{\partial l}\overline{g}(\overline{X}, 1)$.

12

Let $l^{FB}(X) := \arg\max_l \{\overline{y}(X, l) - c(l)\}$ 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. Note that $\lim_{X \to \infty} v(X) = \infty$ as per our maintained assumption. Note also that the efficient level of total effort $l^{FB}(X)$ is strictly positive and less than 1. Finally, let $\overline{a}(X, l) := \arg\max_{a \in [0, l]} \{f(X, a) + g(l - a)\}$ denote the output-maximizing skilled effort given knowledge and total effort levels (X, l), and note that $\frac{\partial}{\partial l}\overline{y}(X, l) = \max\{f_a(X, \overline{a}), g'(l - \overline{a})\}$, so in particular $\frac{\partial \overline{y}}{\partial l}$ exists.

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.

The next theorem characterizes the effort and knowledge paths in an optimal contract.

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 $l^*_{T^*} = l^{FB}(\overline{X})$ and $l^*_t > l^{FB}(X^*_t)$ for $t < T^*$.*

2. *Moreover the optimal knowledge and total effort paths uniquely satisfy, for all $t \leq T^*$,*

$$\frac{1}{r}\frac{d}{dt}v(X^*_t) = v(X^*_t) + c(l^*_t) \tag{11}$$

and

$$\underbrace{\frac{c'(l^*_t)}{\frac{\partial}{\partial l}\overline{y}(X^*_t, l^*_t)}}_{\text{Effort distortion}} = \min\left\{ \underbrace{\frac{1}{1 - r\int_t^{T^*}\rho_\tau d\tau}}_{\text{Target distortion}}, \underbrace{\frac{c'(1)}{\frac{\partial}{\partial l}\overline{y}(X^*_t, 1)}}_{\text{Maximum feasible distortion}} \right\}, \tag{12}$$

where $v(X^*_{T^*}) = v(\overline{X})$ and $\rho_\tau = \rho(X^*_\tau, \overline{a}(X^*_\tau, l^*_\tau))$ is the knowledge premium at τ.

Equation (11) is the state equation. In economic terms it states 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 (12) characterizes the ratio $c'/(\partial\overline{y}/\partial l)$, which equals 1 in the first-best case. Absent an upper bound on the agent's total effort the principal would set $c'/(\partial\overline{y}/\partial l) = \left(1 - r\int_t^{T^*} \rho_\tau d\tau\right)^{-1}$; when this is not feasible she sets it as high as possible.

By way of intuition, consider the following heuristic derivation of the optimal time path of knowledge transfers and effort distortions. Recall from Proposition 1 that the optimal choice of l and dt depend only on the stock of knowledge $\overline{X} - X$ that remains to be sold. Suppose the agent already has knowledge X and consider the problem of selling them a discrete increment of knowledge of value 1, that is $\frac{1}{r}\left[v(X + dX) - v(X)\right] = 1$. The principal must decide the time interval dt over which this knowledge is sold and the effort flow l asked of the agent during this interval.

The principal receives approximately $\overline{y}(X, l)\,dt$ from employing the agent over the interval dt. Moreover, Lemma 2 shows that in an optimal contract, after a possible initial injection of knowledge, the agent's participation constraint holds with equality. So if the principal sells additional knowledge of value one,

$$1 = [v(X) + c(l)]dt,$$

where $[v(X) + c(l)]dt$ is the agent's total opportunity cost of working for the principal over dt. This constraint, which causes the principal to internalize the agent's effort cost, tells us that a higher effort flow must be accompanied by a quicker sale (a smaller dt).

Note next that if the principal planned to sell knowledge now and then sell no further knowledge, it would be optimal to set $l = l^{FB}(X)$, which maximizes flow surplus. To see why, observe that the principal's gain from this sale is

$$\overline{y}(X, l)dt = \overline{y}(X, l)\left(\frac{1}{v(X) + c(l)}\right),$$

and that the right hand side of this equality equals 1 when $l = l^{FB}(X)$ and is smaller for other values of l because total surplus $\overline{y} - c$ is strictly concave.

When instead the principal plans to sell additional knowledge, she is in more of a rush to complete the sale, because a shorter dt means a shorter delay until she can begin collecting profits from her subsequent sales. As a result, she sets dt smaller than the above

baseline and is able to ask for effort higher than first best.

Note, finally, that the greater the agent's current knowledge X, the less knowledge the principal has left to sell, and the smaller her gain from shortening dt. As a result, the sale of the next increment of knowledge is less rushed and the effort distortion, by necessity, falls. When the principal makes her final sale she is in no rush at all; therefore, she sets effort to first best and stretches dt out as long as needed to extract all rents from the agent.

The target distortion $\left(1 - r \int_t^{T^*} \rho_\tau d\tau\right)^{-1}$ in equation (12) measures the degree to which the expert is in a rush. This target distortion is equal to 1 when $t = T^*$ (and so effort is first best) and it grows as t falls because the principal has more knowledge left to sell. In addition, for any given t, the target distortion is greater the greater are the future knowledge premia. This makes future knowledge transfers more profitable, and so the principal is more eager to carry them out. Note, finally, that whenever the target distortion is large enough, the principal sets l_t equal to its upper bound.

We are now ready to characterize the (unique) optimal initial knowledge level and contract length.

Theorem 2 *The optimal initial knowledge level X_0^* and contract length T^* are unique and satisfy either*

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

$$or$$

$$X_0^* = \underline{X} \quad and \quad \int_0^{T^*} \rho_t^* dt \leq \frac{1}{r} \quad (\text{\textbf{zero knowledge gift}}),$$

where $\rho_t^* = \rho(X_t^*, a_t^*)$. The first case arises when the principal's total knowledge $\overline{X}$ exceeds a finite cutoff $\widehat{X}$; the second case arises otherwise.

To provide intuition for the knowledge gift, we revisit our heuristic derivation. Let $\Pi\left(X, \overline{X}\right)$ denote the principal's optimized continuation payoff when the agent has knowledge X and total knowledge is $\overline{X}$. This payoff is decreasing in X and increasing in $\overline{X}$. Now suppose the agent has knowledge $X < \overline{X}$ and, as before, consider the problem of transferring knowledge worth 1 over time interval dt. The principal's total payoff, including

her continuation payoff after this transfer is complete, is approximately

$$\overline{y}(X,l)\,dt \;+\; \underbrace{(1-rdt)}_{\simeq\, e^{-rdt}}\Pi\left(X+dX,\overline{X}\right)$$

and the agent's participation constraint is $1 \geq [v(X) + c(l)]\,dt$, which we now state with an inequality to allow for the possibility that part or all of dX is gifted to the agent.

The greater $\overline{X}$ is, the greater the principal's rush to complete the transfer of dX, and therefore the greater the desired effort l. And indeed when $\overline{X}$ is sufficiently large, the principal sets l equal to its upper bound. Once this occurs the principal's choice of dt maximizes

$$\overline{y}(X,1)\,dt + (1-rdt)\,\Pi\left(X+dX,\overline{X}\right)$$

subject to $1 \geq [v(X) + c(1)]\,dt$. Since the objective is linear in dt, when $\overline{X}$ is high enough or X is low enough that $r\Pi\left(X+dX,\overline{X}\right)$ exceeds $\overline{y}(X,1)$, the principal sets $dt = 0$ and therefore instantly gifts dX to the agent. In particular, the principal will gift knowledge until the agent has knowledge X' such that $r\Pi\left(X'+dX,\overline{X}\right)$ equals $\overline{y}(X',1)$; in limit as $dX = 0$ this becomes $r\Pi\left(X',\overline{X}\right) = \overline{y}(X',1)$. Consequently, since all such gifts occur at time 0, the optimal initial knowledge X_0^* is the unique solution to

$$\Pi\left(X_0^*,\overline{X}\right) = \frac{1}{r}\overline{y}(X_0^*,1).$$

Notice that since Π is increasing in $\overline{X}$, the initial gift is increasing in $\overline{X}$ as well. We also learn that the knowledge cutoff $\widehat{X}$ in the theorem uniquely satisfies $\Pi\left(\underline{X},\widehat{X}\right) = \frac{1}{r}\overline{y}(\underline{X},1)$.

Finally, when the principal's total knowledge is below this cutoff, the agent's ex-ante participation constraint $U_0 \geq \frac{1}{r}v(\underline{X})$ binds and the agent starts working with knowledge X_0^* equal to their knowledge endowment $\underline{X}$. Since the optimal contract is renegotiation proof (Proposition 1), increasing $\underline{X}$ is like starting the contract at a later date. Thus, when the agent's knowledge endowment is higher the contract becomes shorter: since the principal has less knowledge to sell, she sells it more quickly and makes the agent work less hard for it.

To understand why the unconstrained apprenticeship length satisfies $r\int_0^{T^*} \rho_t dt = 1$, consider a contract with an arbitrary length T and first-best effort at the terminal time. Now suppose the principal has the agent work a bit longer at the end of the apprenticeship

16

(over time dT) so that the principal gains 1 in present value, namely,

$$e^{-rT}\overline{y}\left(\overline{X}, l^{FB}\left(\overline{X}\right)\right)dT = 1.$$

Since this change lowers the agent's continuation value throughout the apprenticeship, to prevent them from walking away the principal needs to lower the agent's knowledge at each t. Specifically, to keep the participation constraint satisfied with equality, for all $0 < t \le T$ the principal must set dX_t so that

$$\frac{1}{r}v'(X_t)\,dX_t = \frac{1}{r}f_X\left(X_t, a^{FB}(X_t)\right)dX_t = -e^{rt}\cdot 1 \tag{13}$$

and so $dX_t = -re^{rt}/f_X\left(X_t, a^{FB}(X_t)\right)$. The overall loss in output due to this reduction in knowledge is

$$-\int_0^T e^{-rt}f_X(X_t, a_t)\,dX_t dt = r\int_0^T \rho_t dt.$$

Since a larger T means that output must drop over a longer period of time, this loss is increasing in T. It follows that the optimal length T^* sets this loss equal to the original gain of 1.

Recall that the unconstrained effort level is given by $\frac{c'(l_t)}{\frac{\partial}{\partial l}\overline{y}(X_t, l_t)} = \left(1 - r\int_t^{T^*}\rho_\tau^* d\tau\right)^{-1}$. Then since $r\int_0^{T^*}\rho_t^* dt = 1$, the time-0 target effort distortion in Theorem 1 is infinitely large. If this 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 some initial knowledge for free. Notice also 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 for the agent, so the agent is willing to work longer to acquire it.

Corollary 1 *When the agent's ex-ante participation constraint $U_0 \ge \frac{1}{r}v(\underline{X})$ 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 $T^* < \frac{1}{r}$.

5   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 $f_a(\overline{X}, 0) < g'(1)$ so at each time the agent exerts effort only on the unskilled task. Suppose further that $\overline{X}$ 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
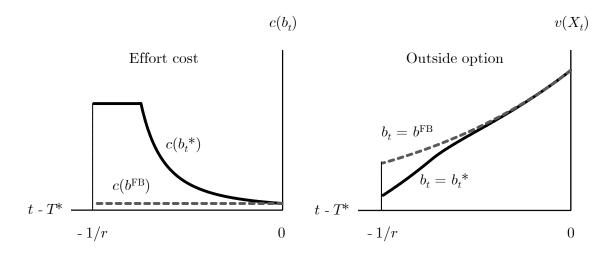
$$\frac{c'(b_t^*)}{g'(b_t^*)} = \min\left\{ \frac{T^*}{t} \ , \ \frac{c'(1)}{g'(1)} \right\}.$$

Thus effort weakly falls over time, strictly so whenever effort is below its upper bound 1.
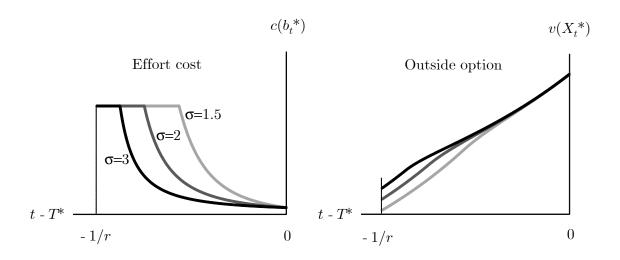
Figure 1.A depicts an optimal contract. The optimal effort distortion balances a loss in instantaneous surplus against a higher rate of knowledge transfer. Early on, when the principal is in most of a rush to transfer knowledge, effort equals its upper bound. Note that if the maximum feasible effort 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. Thus while overworking the agent allows the principal to extract more rent from the agent, this wastes effort and lowers the agent's productivity, causing social surplus to fall.

Figure 1.B illustrates how the contract changes with the curvature of the cost function c.[11] 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. This raises the principal's payoff, but makes the apprenticeship

---

[11]As σ varies, α, β and the maximum feasible effort vary so that b^{FB}, $c\left(b^{FB}\right)$ and the maximum feasible effort cost are held constant.

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 1: Unskilled effort only.

more aversive to the agent, causing both the knowledge path and social surplus to fall.

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 c.[12]
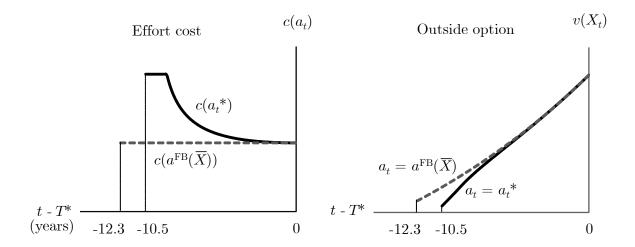
Corollary 1 implies that the optimal apprenticeship length is always strictly less than $\frac{1}{r}$. This is because the complementarity between effort and knowledge means that increasing effort past the first-best level also increases the principal's incentive to train the agent. In contrast to the 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 2.A depicts an optimal contract. 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 $a^{FB}(\overline{X})$ 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 $a_t \le a^{FB}(\overline{X})$) can either help or hurt the agent depending on the details of the technology.
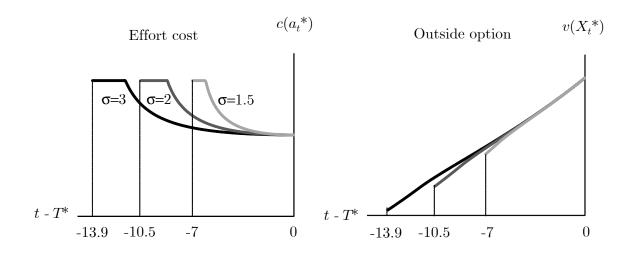
Figure 2.B illustrates how the contract changes with the curvature of the cost function c.[13] As in the unskilled case, as c becomes more linear (σ falls) so that effort is closer to

---

[12] We can show by example that each case can occur.

[13] As in figure 1.B, when σ varies, α, β and the maximum feasible effort vary so that $a^{FB}(\overline{X})$,

Panel A. $y = Xa$ and $c = a^2$.

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

Figure 2: skilled effort only. $\frac{1}{r} = 20$ years.

a money transfer, the principal is willing to impose a higher effort cost on the agent. This raises the principal's payoff and, unlike in the unskilled case, it also shortens the apprenticeship. In the example in the figure, this positive countervailing effect is strong enough that the agent's payoff grows as well.

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 regime is qualitatively the same as the special case where only skilled effort is used.[14]

Consistent with the predicted evolution of the allocation of effort, young cooks at the Le Gavroche restaurant in London, under the tutelage of world-renowned chef Michel Roux Jr., must endure seven career stages. Over seven plus years, cooks gradually progress from doing, as Roux himself warns, "the jobs no one wants" (under the Apprentice position) to "work-horse" prep work (the Commis position) and eventually, if successful, to supervising activities until becoming a Head Chef, either at Le Gavroche or elsewhere.[15]

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

---

$c\left(a^{FB}\left(\overline{X}\right)\right)$ and the maximum feasible effort cost are held constant.

[14] In regime 1, $f_a\left(X_t^*, a^{FB}\left(X_t^*\right)\right) \leq q$, whereas in regimes 2 and 3, $f_a\left(X_t^*, a^{FB}\left(X_t^*\right)\right) > q$. Moreover, in regimes 1 and 2, $f_a\left(X_t^*, a_t^*\right) \leq q$, whereas in regime 3, $f_a\left(X_t^*, a_t^*\right) > q$. That 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.

[15] See www.micheleroux.co.uk/working.html. Such patterns extend well beyond the food industry.

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.

6 Extensions

6.1 Regulating apprenticeships

Say that a contract is Pareto-efficient if it maximizes the principal's payoff for a given payoff of the agent, over all contracts that satisfy the agent's participation and liquidity constraints, and otherwise say the contract is Pareto-dominated. Because each claim in the proof of Lemmas 1 and 2 shows that a contract that does not have the stated properties is Pareto-dominated by one that does, it follows that for each feasible value of the agent's payoff, there is a Pareto-efficient contract that satisfies the claims of those lemmas. Thus any Pareto-efficient payoffs (Π, U) can be 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. Equivalently, regulations can enforce the same result with the combination of a time-varying effort cap and a limit on the duration of the apprenticeship. Neither of these alternatives seems very realistic.

Instead, we now consider a time-invariant cap on the agent's effort, beginning with the case where there is only unskilled effort. Here the principal's optimal contract length (when unconstrained by the agent's initial knowledge) is $T^* = 1/r$, independent of the size of the gift. Thus, as illustrated in Figure 1, when the agent's initial knowledge is low enough that the ex-ante participation constraint does not bind, a time-invariant cap on effort that is weakly above the first-best level leaves the contract length unchanged, and leads to a larger knowledge gift, which increases the agent's productivity and so increases social surplus. However, such contracts are typically not Pareto-efficient, because they do not correspond to a knowledge gift followed by the decreasing effort path required by a principal-optimal contract. When instead the agent's initial knowledge is high enough that the agent's ex-ante participation constraint binds, a uniform effort cap leads to a Pareto-worse outcome as the agent's equilibrium utility is not changed while the knowledge transfer is slowed.

When the agent only exerts skilled effort, the welfare implications of a uniform effort cap again depend on whether or not the agent's ex-ante participation constraint is slack.

Here, unlike with unskilled effort, the cap has an ambiguous effect when the ex-ante participation constraint is slack. As in the unskilled case excessive effort at any given time damages surplus, but here by raising the knowledge premia, overwork induces the principal to shorten the apprenticeship and therefore transfer knowledge more quickly (see Figure 2). When instead the constraint binds, then as in the unskilled case overworking the agent leaves the agent's equilibrium utility unaffected and reduces the apprenticeship length, and so a uniform effort cap lowers the principal's payoff without helping the agent.

Similarly, while a cap on the length of the apprenticeship can raise total surplus, it is not Pareto-efficient (except for the extreme case where the principal is required to give away all of her knowledge at once), as it leads the principal to distort effort even at the terminal date in order to sell her knowledge more quickly.[16]

It does however seem plausible that regulations could combine a time-invariant effort cap with a limit on the training period. This will still not lead to Pareto-efficient contracts but (depending on the welfare weights used) can more often lead to higher total surplus, because a cap on contract length limits the principal's ability to inefficiently 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.

6.2 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

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

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 $\overline{X}$. 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

---

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

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 $X < \overline{X}$ 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 $\frac{d}{dt}\left[\frac{1}{r}v(X)\right] = v(X) + c\left(l^{FB}(X)\right) = \overline{y}\left(X, l^{FB}(X)\right)$. Therefore, net of the training cost, this arrangement delivers profit

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

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

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

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, so that from (13) her overall training costs change by

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

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 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

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

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

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

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   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

26

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.

Nevertheless, it seems likely that in some cases the agent's outside opportunity is lower than $v(X)$ unless they are provided with a certificate of completion 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 certificate is $\underline{v}$ regardless of their level of training, the principal can implement the full-commitment solution with first-best effort at all times. More generally, if the agent's outside option with knowledge X is $\frac{1}{r}v(X)$ with a certificate and $\frac{1}{r}\max\{v(X) - \Delta, \underline{v}\}$ without it, one can show that the optimal contract has two phases. 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 $X_t = \overline{X}$. Phase 2 corresponds to the solution for large Δ: here the agent exerts first-best effort, receives no additional training, and works for a time interval τ that makes the agent indifferent between obtaining the certificate or leaving with value $v\left(\overline{X}\right) - \Delta$, so that the principal extracts the full value Δ/r of the certificate from the agent. When Δ is small, the solution is very similar to the solution in our main model; as Δ grows, phase 2 grows longer and the solution moves closer to the case of full commitment.

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.[17]

8   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 $X' \in (\underline{X}, \overline{X})$, a', and b' s.t. $y(X', a', b') > \underline{v}$ (which is feasible because $v(\overline{X}) > \underline{v}$) and then pick $T' > 0$ s.t. $e^{-rT'}v(\overline{X}) - (1 - e^{-rT'})c(a' + b') > v(X')$. 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 $\overline{X}$. 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,

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

In addition, for $t < T'$,

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

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

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

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 $U_0\left(S'\right) = U_0\left(S\right)$, it follows that $\Pi_0\left(S'\right) > \Pi_0\left(S\right)$, 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 $W_\infty \geq \frac{1}{r}\max\left\{\underline{v}, v\left(\underline{X}\right)\right\} > 0$, so the contract is strictly dominated by the previous claim. ∎

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

Proof. If a finite-duration contract with $W_\infty = 0$ has $X_T < \overline{X}$, then there is a time interval Δ and effort levels a', b' such that $y(X', a', b') > \underline{v}$ and $e^{-r\Delta}v(\overline{X}) - (1 - e^{-r\Delta})c(a' + b') > v(X_T)$, 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 $X_{T'} = \overline{X}$. ∎

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 $X_T = \overline{X}$ and zero wages.*

Proof. From Claims 3 and 4, we can restrict to finite-duration contracts such that $X_T = \overline{X}$ 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,

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

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, $U_t = e^{-r(T-t)}\frac{1}{r}v(\overline{X}) - \int_t^T e^{-r(\tau-t)}c(a_\tau + b_\tau) d\tau$, which is strictly increasing (because $v(\overline{X}) > 0$) 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 exist a date $z < T^*$ 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. Now consider a new contract $\widehat{S}$ that is identical to S^* for $t < z$ and identical to S^{**} for $t \geq z$. By construction, $\widehat{S}$ delivers strictly higher profits than S^* while satisfying constraints (2) from z onward. In addition,

$$U_z\left(S^{**}\right) \geq \frac{1}{r} v\left(X_z^{**}\right) \geq \frac{1}{r} v\left(X_z^*\right) = U_z\left(S^*\right),$$

where the first inequality follows from the date z participation constraint and the last equality follows from Lemma 2.

In addition, for all $0 \leq t < z$ we have

$$
\begin{aligned}
U_t\left(\widehat{S}\right) &= e^{-r(z-t)} U_z\left(S^{**}\right) - \int_t^z e^{-r(\tau-t)} c\left(l_t^*\right) d\tau \\
&\geq e^{-r(z-t)} U_z\left(S^*\right) - \int_t^z e^{-r(\tau-t)} c\left(l_t^*\right) d\tau \\
&= U_t\left(S^*\right) = \frac{1}{r} v\left(X_t^*\right),
\end{aligned}
$$

where the first two equalities follow from the fact that both S^* and $\widehat{S}$ prescribe zero wages before z. As a result, $\widehat{S}$ also satisfies all constraints (2) before z, a contradiction. ∎

31

Proof of Lemma 3. A solution to (6) exists for fixed T from Kumar (1969), and an optimal T exists because the principal's optimized profits given T are continuous in T, and $0 \leq T \leq T^{\max} := \frac{1}{r} \log \left(\frac{v(\overline{X})}{v(\underline{X})} \right)$, because even if asked to exert 0 effort for all T the agent would be unwilling to accept any $T > T^{\max}$.

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 Theorems 3.18 implies that $\lambda_0 = \zeta$ and so $\lambda_0 \geq 0$ and $\lambda_0 [u_0 - v(\underline{X})] = 0$. ∎

9   Appendix A2: Proof of 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:

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

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

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

Lemma A1

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

$$\lambda_t = e^{-rt} \left[e^{rT} \lambda_T - \int_t^T \rho_\tau d\tau \right].$$
(14)

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

Proof. *Part 1.* Using the facts that $v'\left(\phi\left(u\right)\right) = f_X\left(\phi\left(u\right), a^{FB}\left(\phi\left(u\right)\right)\right)$ (from the envelope theorem) and $\phi'\left(u_t\right) = \frac{1}{v'(\phi(u))}$ (from the implicit function theorem) we obtain $\frac{d}{du} f\left(\phi\left(u\right), a\right) = f_X\left(\phi\left(u\right), a\right) \frac{1}{v'(\phi(u))} = \frac{f_X(X,a)}{f_X(X, a^{FB}(X))}\Big|_{X=\phi(u)}$. Thus the co-state evolution equation is $\dot{\lambda}_t = -r\lambda_t + e^{-rt} \frac{d}{du} f\left(\phi\left(u_t\right), a_t\right) = -r\lambda_t + e^{-rt}\rho_t$, which is equivalent to (14).

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

$$\underbrace{c'\left(l_T\right)\overline{y}\left(\overline{X}, l_T\right) - \frac{\partial}{\partial l}\overline{y}\left(\overline{X}, l_T\right)\left[v\left(\overline{X}\right) + c\left(l_T\right)\right]}_{=:\, h(l_T)} = -e^{rT}\eta_T\left[v\left(\overline{X}\right) + c\left(l_T\right)\right].$$
(15)

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

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

where the inequality follows from the fact that $c'' > 0$ and at each point of differentiability $\frac{\partial 2}{\partial l^2}\overline{y} \leq 0$ (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 $l_T = 1 > l^{FB}\left(\overline{X}\right)$, and so the left-hand side of (15) would be positive and its right-hand side would be negative. Once we set $\eta_T = 0$ it follows that $l_T = l^{FB}\left(\overline{X}\right)$ 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 $\widetilde{l}\left(\lambda_t\right), \widetilde{\eta}\left(\lambda_t\right), \widetilde{\gamma}\left(\lambda_t\right)$. More-*

over, this solution satisfies $\widetilde{\gamma}(\lambda_t) = 0$ *and*

$$c'\left(\widetilde{l}(\lambda_t)\right) = \frac{e^{-rt}\frac{\partial}{\partial l}\overline{y}\left(\phi(u_t),\widetilde{l}(\lambda_t)\right) - \widetilde{\eta}(\lambda_t)}{\lambda_t r}, \tag{16}$$

$$\widetilde{\eta}(\lambda_t) = \max\left\{0,\ e^{-rt}\frac{\partial}{\partial l}\overline{y}\left(\phi(u_t),1\right) - \lambda_t r c'(1)\right\}.$$

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

We are now ready to prove Theorems 1 and 2.

Proof of Theorem 1

Write the co-state equation (14) as

$$\lambda_t = e^{-rt}\left[\frac{1}{r} - \int_t^T \rho_\tau d\tau\right], \tag{17}$$

where we used the fact that $e^{rT}\lambda_T = \frac{1}{r}$. After substituting for λ_t into equation (16) we obtain

$$c'(l_t) = \frac{\frac{\partial}{\partial l}\overline{y}(X_t,l_t) - e^{rt}\eta_t}{1 - r\int_t^T \rho_\tau d\tau}, \tag{18}$$

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

Since $\lambda_0 \geq 0$ implies that $1 - r\int_0^T \rho_t dt \geq 0$, equation (18) implies that the optimal effort path satisfies $\frac{c'(l_t^*)}{\frac{\partial}{\partial l}\overline{y}(X_t^*,l_t^*)} = \min\left\{\left(1 - r\int_t^{T^*}\rho_\tau d\tau\right)^{-1}, \frac{c'(1)}{\frac{\partial}{\partial l}\overline{y}(X_t^*,1)}\right\}$ with $\rho_\tau = \rho\left(X_\tau^*,\overline{a}(X_\tau^*,l_\tau^*)\right)$, 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 $l_{T^*}^* = l^{FB}\left(\overline{X}\right)$ and $l_t^* > l^{FB}(X_t^*)$ for all $t < T^*$ (as claimed in part 1 of the theorem). Finally, since $c'' > 0$ and $\frac{\partial\overline{y}}{\partial l}$ is weakly decreasing in l, for any given T^* the values of X_t^* and l_t^* are unique. ∎

34

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 $u_T = v(\overline{X})$, 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 (18) to obtain

$$
c'(l_{T-s}) = \begin{cases} \min\left\{ \dfrac{\frac{\partial}{\partial l}\overline{y}(X_t, l_{T-s})}{1 - r\int_0^s \rho_{T-\tau}d\tau}, \; c'(1) \right\} & \text{when } r\int_0^s \rho_{T-\tau}d\tau \leq 1, \\ c'(1) & \text{otherwise,} \end{cases} \tag{19}
$$

where $\rho_{T-\tau} = \rho\left(\phi(u_{T-\tau}), \overline{a}(X_{T-\tau}, l_{T-\tau})\right)$. Second, note that since $c'' > 0$ and $\frac{\partial \overline{y}}{\partial l}$ 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

$$
u_T = v(\overline{X}), \; \dot{u}_{T-s} = r\left[u_{T-s} + c(l_{T-s})\right], \text{ and (19), for } s \geq 0.
$$

(The resulting paths λ_{T-s}^T and $\eta_{T-s}^T = \widetilde{\eta}\left(\lambda_{T-s}^T\right)$ 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 $u_0^T = u_{T-s}^T\big|_{s=T}$, we have $\frac{d}{dT}u_0^T = -\frac{d}{dt}u_t^T\big|_{t=0} = -r\left[u_0^T + c(l_0^T)\right]$. 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 $\lambda_0^T\left[u_0 - v(\underline{X})\right] = 0$, 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 $\widehat{T}$ such that $\lambda_0^{\widehat{T}} = 1 - r\int_0^{\widehat{T}} \rho_t dt = 0$; and since $\lambda_0^T \geq 0$, we may restrict our search to $T \leq \widehat{T}$.

There are two cases to consider: (a) $u_0^{\widehat{T}} \geq v(\underline{X})$; and (b) $u_0^{\widehat{T}} < v(\underline{X})$. In case (a), we must have $T^* = \widehat{T}$ and therefore $r\int_0^{T^*} \rho_t dt = 1$. Otherwise $T^* < \widehat{T}$, $\lambda_0^{T^*} > 0$, and $u_0^{T^*} > v(\underline{X})$, violating the complementary slackness condition. In case (b), we must have $T^* < \widehat{T}$, $\lambda_0^{T^*} > 0$, and $u_0^{T^*} = v(\underline{X})$. Therefore, T^* is the unique value of T such that $u_0^T = v(\underline{X})$. ∎

10 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 $X_T = \overline{X}$, 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

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

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

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

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

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

where the second term in the objective is the cost of the initial gift. The Hamiltonian is now $\mathcal{H} = e^{-rt}\left[\overline{y}\left(\phi(u_t), l_t\right) - k\frac{1}{r}\dot{u}_t\right] - \lambda_t \dot{u}_t$, with $\dot{u}_t = r\left[u_t + c(l_t)\right]$. 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 $\zeta\left[u_0 - v\left(\underline{X}\right)\right] = 0$.*

Proof. Define $\varphi(u_0) := -k\frac{1}{r}\left[u_0 - v\left(\underline{X}\right)\right]$, $\psi(u_0) := u_0 - v\left(\underline{X}\right)$ 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 $\lambda_0 = \Phi'(u_0) = -\frac{1}{r}k + \zeta$. $\blacksquare$

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*

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

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

Proof. The co-state evolution equation is $\dot{\lambda}_t = -r\lambda_t + e^{-rt}\left[\rho_t - k\right]$, the first-order condition for T is $e^{-rT}\,\overline{y}\left(\phi\left(u_T\right), l_T\right) - \left(\lambda_T r + e^{-rT}k\right)\left[u_T + c\left(l_T\right)\right] = 0$ and the first-order condition for l_T implies that $\lambda_T r + e^{-rT}k = \left[e^{-rT}\frac{\partial}{\partial l}\overline{y}\left(\overline{X}, l_T\right) - \eta_T\right]/c'\left(l_T\right)$. 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. $\blacksquare$

Propoistion 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*

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

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

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

or

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

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

37

satisfies, for all $s \geq 0$,

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

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 (2016). "Allocating Effort and Talent in Professional Labor Markets." *Mimeo*, University of Chicago.

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. Department for Business, Innovation and Skills (2013). "Apprentices' Pay, Training and Working Hours." *BIS*, UK Government.

13. 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.

14. Garicano, Luis, and Luis Rayo (2017). "Relational Knowledge Transfers." *American Economic Review,* forthcoming.

15. Gapper, John (2014). "Bankers and Lawyers are on an unhealthy treadmill." *Financial Times (January 15).*

16. Holmström, Bengt (1999). "Managerial incentive problems: A dynamic perspective." *Review of Economic Studies* 66(1): 169-182.

17. Hörner, Johannes, and Andrzej Skrzypacz (2016). "Selling information." *Journal of Political Economy* 124(6): 1515-1562.

18. Jauhar, Sandeep (2015). "The Real Problem With Medical Internships." *New York Times* (Jul 1).

19. Jovinelly, Joann, and Jason Netelkos (2007). *The crafts and culture of a medieval guild.* New York: The Rosen Publishing Group.

20. Katz, Eliakim, and Adrian Ziderman (1990). "Investment in General Training: The Role of Information and Labour Mobility." *Economic Journal* 100(403): 1147-1158.

21. Kumar, T. Krishna (1969). "The existence of an optimal economic policy." *Econometrica* (Oct): 600-610.

22. 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.

23. 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.

24. Lane, Joan (2005). *Apprenticeship in England, 1600-1914.* London: Routledge.

25. Maister, David H. (1993). *Managing the Professional Service Firm.* Free Press, New York, NY.

26. 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.

27. Park, Ryan (2017). "Why So Many Young Doctors Work Such Awful Hours." *The Atlantic* (Feb 21).

28. Schwartz, Robert J., et al. (1992). "Guidelines for surgical residents' working hours: intent vs reality." *Archives of Surgery* 127(7): 778-783.

29. Smith, Adam (1872). *An Inquiry Into the Nature and Causes of the Wealth of Nations: Volume 1.* A. Murray & Company.

30. 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).

31. Thomas, Jonathan, and Tim Worrall (1994). "Foreign direct investment and the risk of expropriation." *Review of Economic Studies* 61(1): 81-108.