

Training and Effort Dynamics in Apprenticeship*

Drew Fudenberg

MIT

drew.fudenberg@gmail.com

Luis Rayo

Northwestern University and CEPR

luis.rayo@kellogg.northwestern.edu

April 30, 2019

Abstract

A principal specifies time paths of effort provision, task allocation, and knowledge transfer 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 the optimal time path of effort distortions and their impact on the knowledge transfer, and analyze the effect of regulatory limits on the length of apprenticeships and on how much effort apprentices are allowed to provide.

*First Draft: 05/2017. We thank Daron Acemoglu, Robby Akerlof, Gonzalo Cisternas, Guido Friebel, Luis Garicano, Bob Gibbons, Niko Matouschek, Derek Neal, John Pavlus, Klaus Schmidt, Rachel Spencer, Ivan Werning, and workshop participants at Chicago, Columbia, EEA-ESEM, Harvard, Kellogg, MIT, NBER, SITE, Toronto, Toulouse, and UCLA 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 while working long hours. This raises the questions of whether the employers will specify more effort than would be socially optimal and longer training periods than strictly needed for the desired training.

In this paper, we consider the design of optimal (profit-maximizing) careers by a principal with commitment power, who can specify time paths of effort provision, task allocation, and knowledge transfer 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 where 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.

Our goal is to understand how the intensity and composition of apprentice effort vary over the course of the apprenticeship, the interaction between effort and the knowledge transfer, and the implications of our results for various possible policy interventions. To study effort composition, we suppose that 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 because the agent will work inefficiently hard to compensate the principal for their training and because this training is inefficiently slow. Inefficiently slow training makes the skilled task relatively less efficient, so the agent spends more time on unskilled or menial work. As we show, the extent of overwork decreases throughout the apprenticeship, and the agent spends a

decreasing amount of time on menial effort. These implications of the model accord well with our anecdotal observations in Section 2.

The key to understanding the time path of effort distortions is that when the principal asks the agent to exert more effort in the future the agent is more tempted to walk away at all earlier times. To offset this, she must lower the agent's outside option by training them more slowly, which reduces the agent's productivity. This productivity loss is greater for effort that is farther in the future, so over time the contracted effort path becomes less and less distorted, and indeed as we show the effort distortion shrinks to zero at the end of the apprenticeship. Moreover, because the agent is becoming more productive in the skilled task and their overall effort level is becoming less distorted, they spend less and less time on the menial task.

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 principle be desirable. As we show below, bounds on effort alone can lead to even slower training, and bounds on apprenticeship length can increase overwork. Pareto-efficient regulations require a cap on apprenticeship length along with a time-varying bound on effort, though total surplus can also sometimes be increased by a cap on apprenticeship length combined with a judiciously chosen constant upper bound on hours worked.

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 effort provision with an endogenously evolving participation constraint.

Our model builds on that of Garicano and Rayo (2017) (GR) who studied profit-maximizing training in a model without effort, and showed that for strategic reasons the principal will transfer knowledge slowly even though it would be more efficient to transfer all of it at once. Because their model does not include effort, it cannot address the empirical facts about workload and its composition that are the motivation for our analysis. Like us, GR show that the principal can induce the agent to work for little or no

pay in exchange for training, and in the main text we take this as given.¹ However, our analysis of effort distortions and their effect on the optimal contract uses entirely different techniques, based on the theory of optimal control. And in addition to being of intrinsic interest, the workload and tasks of apprentices are easier to observe than the speed of knowledge transfer, and therefore serve as potential tests of the underlying theory.

Our work and GR are both related to earlier work by Thomas and Worrall (1994) and Albuquerque and Hopenhayn (2004) on lending with limited enforcement in models without effort. In both of these papers it is also true that the agent's outside option and productivity increase gradually over time, and all consumption is postponed until output reaches its steady-state level. However, in these earlier papers repayment cannot be enforced unless the principal is able to punish the agent for walking away.² Thus our work differs both because the principal can induce the agent to work without the threat of punishment in exchange for training, and because of our analysis of the time path of overwork.

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.³

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 Akerlof (1976), 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 and does not distinguish between skilled and unskilled effort.

The remainder of the paper is organized as follows. Section 2 presents motivating examples. Section 3 sets up the model, using some simplifications of the space of feasible

¹In the Online Appendix we show that even if the principal is allowed to pay the agent wages and keep them as an apprentice forever, the optimal contract specifies 0 (or minimum) wages and trains the agent in finite time. The argument for this result does not follow from that in GR but the intuition is quite similar.

²In Thomas and Worrall the punishment is that the host country cannot produce without the help of the foreign firm; in Albuquerque and Hopenhayn the punishment is left more abstract.

³Hörner and Skrzypacz (2016) study gradual information revelation by a privately informed agent about her competence; their model has neither effort nor human capital.

contracts that we justify in Appendix A. Section 4 states and explains our main results. Section 5 analyzes extensions to cases where the need for certification gives the agent some limited commitment power, and where the agent can make initial cash payments. These extensions expand the applicability of the model and suggest that our main findings are robust. Section 6 uses our results to study the impact of various sorts of regulations. Appendix B provides proofs for all of the results that are stated in the text.

2 Motivating examples and alternative theories

2.1 Motivating examples

Work-for-training arrangements are common in a wide range of industries. Consistent with our predictions, apprentices frequently experience long hours and heavy workloads, initially spend a large share of their time on menial work unrelated to the skills they wish to acquire, and gradually progress toward more advanced tasks.

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.⁴

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.

⁴As 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, [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.michelroux.co.uk/working.html and Lugert (2017).

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 clothed 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). As we will see, this effect of cash payments is predicted by our model.

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 might be granted the privilege of learning to sharpen tools” (p. 350).⁵

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.”⁶ In the U.S.,

⁵For contemporary examples of harsh apprenticeships, see UK Dept. for Business, Innovation and Skills (2013).

⁶Indeed, “[l]ong hours and hard work have been features of medical training since the modern residency

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, which are valuable to the hospital but provide limited learning opportunities for the apprentices (Jauhar, 2015).⁷

In surgery residencies, for example, first year trainees are commonly assigned to “floor work” such as “talking to patients’ families [...] accompanying patients to tests, and obtaining patients’ signatures on consent-to-surgery forms. [...] They rarely see the inside of the OR, and when they do, it’s usually to assist with an emergency or a run-of-the-mill case” (Kirzner, 2014). Over time, residents acquire greater decision-making responsibilities and progress towards more complex procedures. At Vanderbilt, for instance, the second year experience “broadens to include opening and closing laparotomy wounds,” third year residents progress toward “renal transplantation [...] pediatric surgery, and endoscopy,” fourth years are engaged in “vascular [and] thoracic surgery,” and finally fifth years become chief residents and are “involved with major elective procedures emphasizing complex gastroenterologic, oncologic [and] vascular cases.”⁸

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.⁹ 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
program had its beginnings at the Johns Hopkins Hospital in Baltimore in the late 19th century” (Jauhar, 2015, *New York Times*).

⁷Schwartz 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.

⁸See the Vanderbilt University general surgery residency program at:
<https://ww2.mc.vanderbilt.edu/GSR/15429> (accessed 10/15/2018).

⁹As 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).

the firms’ partners, long hours and heavy workloads are common as well (e.g. Coleman 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).¹⁰

Lastly, manicurists in the New York City area go through an apprenticeship reminiscent of those of the Industrial Revolution. 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.”

2.2 Alternative theories

An alternative explanation of effort distortions is that the agent overworks not to purchase knowledge from a master, but to influence the assessment of their boss, or the labor market more generally, about their ability (see the literature cited in the introduction). Several predictions distinguish our theory from this alternative. First, our model predicts that apprentices will have a decreasing time path of menial effort (see Section 4.2), and it is not clear how that prediction would emerge from signalling or screening. Second, in our model

¹⁰While 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.

the agent becomes more knowledgeable over time, so under natural assumptions they will be assigned to increasingly advanced skilled tasks either because it is more efficient to do so or simply because there are tasks that the agent cannot do without receiving training. Finally, our model predicts that agents who can make sufficiently large cash payments can reduce the length of their apprenticeship (see Section 5.2).

Likewise, our model has different and more detailed implications than one in which workers exert effort to increase their probability of promotion in a complete-information tournament: while such tournaments may be a source of incentives it is not clear why they should lead to overwork, let alone to a decreasing time path of menial effort or to the agent working on increasingly-advanced skilled tasks over time.¹¹ This type of model would also not predict excessive effort when the agent expects to switch employer after completing their work, as in many of our examples.

Another reason that agents might work more hours than optimal for static efficiency is that doing so may allow them to learn more quickly. However, this explanation for long work hours suggests that agents ought to do little if any menial work so as to have more time for work that speeds their training; and does not predict a negative correlation between cash payments and apprenticeship duration.

3 Model

3.1 Production, Payoffs, and Knowledge

A principal (she) and an agent (they) interact over an infinite horizon in continuous time. Both players have discount rate r . 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” or “menial,” meaning that its productivity is independent of X_t . Let $q_t := a_t + b_t$ denote total effort. We assume that both a_t and b_t are non-negative and q_t is bounded above by a constant, which we normalize to 1.

¹¹Li et al. (2018) show that a firm may distort upward the number of junior workers it employs, which is conceptually similar to distorting the effort of junior workers, but their model does not deliver predictions about the time path of effort for these workers, or about its composition.

Thus total output y_t is given by

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

We assume that the agent's flow payoff function is quasi-linear in money and that exerting effort q_t imposes cost $c_t := c(q_t) \geq 0$ on the agent.

For ease of exposition, the main text will maintain all of the assumptions that are needed for any of our results. The online appendix breaks these down to show which assumptions are needed where.

Assumption 1

1. f and g are twice differentiable with $f_X > 0$, $f_a \geq 0$, $f_{Xa} > 0$, $f_{aa} \leq 0$, $g' > 0$, $g'' \leq 0$, and either $f_{aa} < 0$ or $g'' < 0$.¹²
2. c is twice differentiable with $c' \geq 0$, $c'(0) = 0$, and $c'' > 0$.

Let $v(X) := \max_{a,b \geq 0} \{f(X, a) + g(b) - c(a + b)\}$ denote first-best surplus given X . Because $f(X, a)$ is strictly increasing in X , so is $v(X)$.

Let superscript $*$ on a function denote its maximal value given the specified arguments, so that $y^*(X, q) := \max_{a \in [0, q]} \{f(X, a) + g(q - a)\}$ denotes the maximum possible output given knowledge and effort levels (X, q) . Similarly, $a^*(X, q) := \arg \max_{a \in [0, q]} \{f(X, a) + g(q - a)\}$ denotes the output-maximizing skilled effort given (X, q) . Finally, $q^*(X) := \arg \max_{q \in [0, 1]} \{y^*(X, q) - c(q)\}$ is the first-best level of total effort, which is unique from Assumption 1, and $a^*(X) := a^*(X, q^*(X))$ is the corresponding first-best level of skilled effort, which is also unique.¹³

To simplify the analysis we assume that the efficient effort level is always less than 1.

Assumption 2 $q^*(X) < 1$ for all $X \leq \bar{X}$.

¹²Note that we have assumed that $f_X(X, a) > 0$ even when $a = 0$. To justify this, imagine that the first $\underline{a} > 0$ units of skilled effort are costless to the agent (for example, because they enjoy a degree of intrinsic motivation). We can then assume, without loss, that the agent exerts at least \underline{a} units of skilled effort, and reinterpret a as skilled effort above \underline{a} .

¹³Note that $\frac{\partial y^*}{\partial q}(X, q) = \max\{f_a(X, a^*(X, q)), g'(q - a^*(X, q))\}$, so in particular $\frac{\partial y^*}{\partial q}$ exists.

The assumptions that $c'(0) = 0$ and $g' > 0$ imply that $q^*(X) > 0$ for all X . This means that our model doesn't exactly nest that of GR, though it can approximate it by setting both the productivity of effort and its cost to be very small.¹⁴

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} . (The Online Appendix shows that the results are qualitatively unchanged if the principal faces a cost of training the agent.) 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 $[\underline{X}, \bar{X}]$.

3.2 Apprenticeship Contracts

The principal has full commitment power, and offers a contract to the agent at time 0 that specifies a graduation date $T \geq 0$ and for each $t \in [0, T]$ a knowledge stock X_t and effort levels (a_t, b_t) . Between dates 0 and T , all output belongs to the principal. After date T , the agent works on their own and keeps all output. While the principal can commit to the contract, the agent can walk away at any time; if the agent does so, the principal does not hire them back. In the main text, we consider only contracts with a finite graduation date T , and where all knowledge is eventually transferred to the agent, so that $X_T = \bar{X}$. We also assume that the principal does not pay wages to the agent, and that total effort $q_t = a_t + b_t$ at each time is allocated to maximize output, so that $y_t = y^*(X_t, q_t)$. As we explain in Appendix A, such contracts are optimal in a more general model that allows $T = \infty$, wage payments, savings, an arbitrary effort allocation, and $X_T < \bar{X}$.¹⁵

In addition, we restrict attention to contracts where the time path of effort satisfies a Lipschitz condition.¹⁶ Thus the set of feasible contracts, denoted Ψ , is the set of all

¹⁴Formally, fix f, g , and c that satisfy our assumptions, and then for $\varepsilon > 0$ define new functions $f^\varepsilon, g^\varepsilon$, and c^ε by $f^\varepsilon(X, a) = f(X, \varepsilon a)$, $g^\varepsilon(b) = \varepsilon g(b)$, and $c^\varepsilon(q) = \varepsilon c(q)$. When ε is very small, effort is minimally productive and minimally costly, and as we will see below, the resulting optimal contract is close to that of GR.

¹⁵Intuitively, because the agent's payoff is linear in consumption and the parties have the same discount rate, any contract with positive wages 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. And since the allocation of total effort across tasks does not affect the agent's utility, it is optimal for the principal to choose an allocation that maximizes output.

¹⁶That is, there exists a constant K such that $|q_t - q_s| \leq K |t - s|$ for all $t, s \in [0, T]$. As we show in Lemma A3 in Appendix B, this condition guarantees existence of an optimal contract.

graduation times T and time paths of effort and knowledge $(q_t, X_t)_{t=0}^T$ with $X_t : [0, T] \rightarrow [\underline{X}, \bar{X}]$ non-increasing and right continuous, $X_T = \bar{X}$, and $q_t : [0, T] \rightarrow [0, 1]$ satisfying a Lipschitz condition. Let $\mathbf{q} := (q_t)_{t=0}^T$ and $\mathbf{X} := (X_t)_{t=0}^T$ denote the time paths of effort and knowledge, and let $\mathcal{C} := (T, \mathbf{q}, \mathbf{X})$ denote a contract.

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, and after graduation the agent's flow payoff is $v(\bar{X})$. If the agent walks away before graduation with knowledge X , they obtain instantaneous surplus $\max\{\underline{v}, v(X)\}$, which will be their flow payoff from then on.

Given a contract $\mathcal{C} \in \Psi$, the principal's and agent's continuation values from date $t \leq T$ onward are

$$\begin{aligned} \Pi_t(\mathcal{C}) &= \int_t^T e^{-r(\tau-t)} y_\tau d\tau \quad \text{and} & (1) \\ U_t(\mathcal{C}) &= \underbrace{e^{-r(T-t)} \frac{1}{r} v(\bar{X})}_{\text{discounted "graduation prize"}} - \underbrace{\int_t^T e^{-r(\tau-t)} c(q_\tau) d\tau}_{\text{"effort tax"}}, \end{aligned}$$

where $\frac{1}{r} v(\bar{X})$ is the "prize" received by the agent upon graduation.

The principal can select any contract she desires subject to the agent's *participation constraint*: At each date $t \leq T$, the agent's continuation value should be at least as high as their outside option. We show in Appendix A that once the agent accepts the contract and the principal has had the chance to raise the agent's knowledge to the desired level X_0 , this participation constraint holds with equality from then on; we simply assume that here.¹⁷ Moreover, since the principal will immediately increase the agent's knowledge to at least the point where $v(X) = \underline{v}$, we write the agent's initial outside option as $\frac{1}{r} v(\underline{X})$, instead of $\frac{1}{r} \max\{\underline{v}, v(\underline{X})\}$. Thus the participation constraint reduces to

$$U_0(\mathcal{C}) \geq \frac{1}{r} v(\underline{X}) \quad \text{and} \quad U_t(\mathcal{C}) = \frac{1}{r} v(X_t) \quad \text{for all } [0, T]. \quad (2)$$

We call the first of these conditions the *initial participation constraint*, and the second

¹⁷The principal wishes to raise X_t to the point where the agent's participation constraint holds with equality because training is costless and increases the agent's productivity.

one the *ongoing participation constraint*.

Rewriting the ongoing participation constraint in differential form, we have

$$\underbrace{dv(X_t)/dt}_{\text{Flow value of new knowledge}} = \underbrace{r[v(X_t) + c(q_t)]}_{\text{Agent's opportunity cost}}.$$

As this equation shows, when the agent is not asked to exert effort so that $c_t = 0$, then $v(X_t)$ needs to increase at rate r : the principal needs to transfer knowledge more quickly when the agent's knowledge stock and outside opportunities are higher. Moreover, the principal needs to transfer knowledge even more quickly the more she asks the agent to work.

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) and (2). To allow for this interpretation of the model, we will allow \hat{y} and \hat{v} to be negative.

The principal's problem is to select a finite graduation date T and time paths of total effort and knowledge before this date that satisfy the initial and ongoing participation constraints. Formally, her problem is

$$\max_{\mathcal{C} \in \Psi} \int_0^T e^{-rt} y^*(X_t, q_t) dt \tag{3}$$

subject to (2).

Note for future reference that any optimal contract is sequentially optimal, so while 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.¹⁸ This is because the agent's ongoing 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.

¹⁸In 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.

4 Solution

We begin by solving a relaxed version of problem (3) where the principal is allowed to select the initial knowledge level X_0 to be arbitrarily low. We then use the solution to the relaxed problem to derive the optimal contract for the original problem.

Define the *knowledge premium* given knowledge X and skilled effort a as

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

Because $f_X(X, a^*(X)) = v'(X)$ 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. 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. As we shall see, the knowledge premium plays a central role in determining both the optimal effort levels and the optimal contract length.

Relaxed problem. Here we suppose the principal is free to choose X_0 unconstrained from below. In this case, the principal's problem simplifies to that of selecting any desired contract length T and effort path \mathbf{q} while setting the knowledge path \mathbf{X} to its highest level such that the agent's ongoing participation constraint is met. That is,

$$v(X_t) = e^{-r(T-t)}v(\bar{X}) - r \int_t^T e^{-r(\tau-t)}c(q_\tau) d\tau. \quad (4)$$

Let $\rho_t := \rho(X_t, a_t)$ denote the corresponding knowledge premium at time t , where $a_t = a^*(X_t, q_t)$.

Theorem 1 characterizes the solution:

Theorem 1 *In the relaxed problem the optimal contract is unique and satisfies:*

1. $r \int_0^T \rho_t dt = 1$.

2. For all t ,

$$\frac{\partial y^*}{\partial q}(X_t, q_t) / c'(q_t) = \max \left\{ r \int_0^t \rho_\tau d\tau, \frac{\partial y^*}{\partial q}(X_t, 1) / c'(1) \right\}. \quad (5)$$

As a result, the optimal effort path is efficient at the terminal time T , and exceeds the efficient level at all earlier times.

Here is a simple heuristic derivation of this result. Consider first the choice of T . A marginal increase in T allows the principal to keep the agent's output for longer, but since this hurts the agent, it forces the principal to lower the agent's knowledge throughout the apprenticeship. That is, the principal faces a trade-off between a longer apprenticeship and a less-productive apprentice.

Suppose, in particular, that the principal lengthens the apprenticeship to $T + dT$, while specifying the same effort as before up to T , and first-best effort over the interval dT . Set dT so that this modification gains the principal, and costs the agent, 1 in present value. To meet the participation constraint, at each time $t \leq T$ the agent's outside option must fall by 1 unit in present value, which lowers their productivity by $r\rho_t$ units in present value; hence, the principal loses $r \int_0^T \rho_t dt$. It is therefore optimal to set

$$r \int_0^T \rho_t dt = 1.$$

Because $\rho_t \geq 1$, this equations tells us that the contract will last at most $1/r$ years, with a shorter apprenticeship the greater the values of ρ_t .

Consider now the choice of effort path. A marginal increase in the required effort level at time t raises output at that time, but since it hurts the agent, the principal must lower the agent's knowledge at all previous times to deter the agent from leaving. Thus the principal faces a trade-off between a higher output today and lower output at every prior date. Suppose, in particular, that the principal raises effort at time t so as to cost the agent 1 extra in present value. This raises output by $\frac{\partial y^*}{\partial q}(X_t, q_t) / c'(q_t)$ in present value. But to meet the participation constraint, at each time $\tau \leq t$ the agent's outside option must fall by 1 unit in present value, which lowers their productivity by $r\rho_\tau$ units in present value; and so the principal's overall loss is $r \int_0^t \rho_\tau d\tau$. Hence, the principal either sets

$$\frac{\partial y^*}{\partial q}(X_t, q_t) / c'(q_t) = r \int_0^t \rho_\tau d\tau, \quad (6)$$

or, when this is impossible given the upper bound on effort, she sets $q_t = 1$. The right-hand side of (6) is the principal's (shadow) cost of effort. The left-hand side of (6) is the (inverse) effort distortion. This ratio is strictly decreasing in effort, and is equal to 1

when effort is first-best. (Since effort is bounded above by 1, the lowest feasible value for the ratio is $\frac{\partial y^*}{\partial q}(X_t, 1)/c'(1)$.)

Early in the relationship the (shadow) cost of effort is so low that the principal optimally sets $q_t = 1$. Then as time goes by this cost grows and effort becomes less and less distorted relative to first best; until, at the terminal time, the principal fully internalizes the agent's true cost and effort becomes first best.

The two equations in the theorem, together with the binding participation constraint (4), define a unique contract. To see why, begin at time T where knowledge is known to be \bar{X} , effort is first best, and $\rho_T = 1$. Now move backward in time using the system (4) and (5), the relationship $r \int_0^t \rho_\tau d\tau = 1 - r \int_t^T \rho_\tau d\tau$, and $\rho_t := \rho(X_t, a^*(X_t, q_t))$ to determine the time-paths of knowledge, effort, and knowledge premia. Because (4) and (5) uniquely define X_t and q_t given the values of $\int_t^T e^{-r(\tau-t)} c(q_\tau) d\tau$ and $\int_t^T \rho_\tau d\tau$, these paths are unique.

Original problem. We now return to the original problem where X_0 is constrained to be no smaller than \underline{X} . Theorem 2 characterizes the solution:

Theorem 2 *In the original problem the optimal contract is unique. Moreover, for all \bar{X} , there is a threshold $\hat{X} < \bar{X}$ such that:*

1. *If $\underline{X} \leq \hat{X}$ the optimal contract corresponds to the relaxed solution in Theorem 1.*
2. *If $\underline{X} > \hat{X}$ the optimal contract is the truncated version of the relaxed solution obtained by eliminating from this relaxed solution all time periods where knowledge is below \underline{X} . That is, the agent begins the contract farther along, with initial knowledge $X_0 = \underline{X}$.*

The threshold \hat{X} is the initial knowledge prescribed by the unconstrained solution. Letting U_0 denote the agent's total payoff in the relaxed solution, this threshold solves $v(\hat{X})/r = U_0$.

The first part of the theorem is immediate. To gain intuition for the second part notice that because the agent's ongoing participation constraint always binds, the unconstrained solution is sequentially optimal. Hence whenever \underline{X} exceeds \hat{X} , it is optimal to begin the contract with knowledge \underline{X} and to proceed from there as prescribed by the unconstrained solution.¹⁹

¹⁹As we show in the Online Appendix, our results extend to the case where the principal incurs a

Notice that when $\underline{X} < \widehat{X}$ the principal grants the agent a knowledge gift at time zero, raising the agent's knowledge to level \widehat{X} . An example would be a summer crash course at the beginning of a consulting career.

4.1 Pareto-efficient contracts

Recall that the principal's optimal contract is sequentially optimal. Hence, as time goes by, the players' continuation payoffs trace the Pareto frontier. It follows 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. 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 total surplus grows monotonically, and reaches the first-best level when Π reaches 0.

We obtain the slope of the Pareto-frontier using a simple heuristic. Pick an arbitrary payoff $U = v(X_0)/r$ for the agent (supported by an underlying Pareto-efficient contract of length T). Now raise U by 1 by gifting the agent a money bonus at graduation worth 1 in present value.²⁰ Measured in present value, this costs the principal 1, but it also allows her to raise the agent's outside option by 1 at each time before graduation, allowing her to gain $r \int_0^T \rho_t dt$ in additional output. Hence, the net loss for the principal – which gives us the slope of the frontier – is


$$d\Pi/dU = -1 + r \int_0^T \rho_t dt.$$

This slope is -1 when the agent begins with knowledge $X_0 = \overline{X}$ (in which case $T = 0$), and becomes less and less negative as X_0 falls (and T grows). Because the principal pays for effort with knowledge, this slope is equal to the principal's (shadow) cost of effort. Hence, in any given Pareto-efficient contract, as time goes by and players move along the Pareto frontier, knowledge becomes more and more costly to the principal and hence the effort distortion falls.

cost k per unit of knowledge transferred to the agent; the only change is that in the equations for the optimal contract length and effort path, $[\rho_t - k]$ takes the place of ρ_t . Hence as k grows the apprenticeship becomes longer, but otherwise has very similar properties.

²⁰Equivalently, the principal could slightly shorten the apprenticeship.

Figure 1



Panel A. Pareto frontier.

Panel B. $\sigma = 1.5, 2, 3$.
(Lighter curves correspond to higher σ .)

Figure 1 depicts the Pareto frontier for the case where $y^*(X, q) = X + q$ and $c(q) = \alpha + \beta q^\sigma$. The slope of this frontier ranges from 0 at the principal optimum to -1 at the agent optimum (panel A). As c becomes more linear (σ falls) effort becomes closer to a money transfer and the Pareto frontier expands (panel B).²¹

4.2 Evolution of effort

Corollary 1 describes the path of unskilled effort.

Corollary 1 *The optimal path of unskilled effort is weakly decreasing, strictly so whenever unskilled effort is positive and below its upper bound of 1.*

To gain intuition, notice that as time goes by the principal's shadow cost of effort grows, and hence unskilled effort becomes less distorted relative its first-best level. At the same time, the marginal value of skilled effort grows, and hence the first-best level of unskilled effort falls. Both these effects cause b_t to fall.

In contrast, the path of skilled effort can in principle be increasing, decreasing, or even non-monotone. This is because, even though skilled effort becomes less and less distorted

²¹In panel B, as σ varies, α , β and the maximum feasible effort vary so that q^* , $c(q^*)$ and the maximum feasible effort cost are held constant.

relative its first-best level, this first-best level is itself increasing. Depending on the details of f and g , either one of these effects can dominate.²²

In practice, novices may go through an initial phase where they do very little skilled work. By imposing additional structure on the output technology we can guarantee that the optimal contract exhibits this property. To illustrate, suppose output is $X \cdot (\underline{a} + a) + g(b)$, where $\underline{a} > 0$ is a small exogenous level of skilled effort that is costless to the agent (e.g. because of intrinsic motivation), a is skilled effort above \underline{a} , and g is strictly concave. And suppose further that g satisfies

$$X_0 < g'(1) < g'(0) < \bar{X},$$

where X_0 is the optimal initial knowledge level given g .²³ In this case the novice goes through an initial phase where $X_t \leq g'(1)$ and $a_t = 0$, then an intermediate phase where $g'(1) < X_t < g'(0)$ and both a_t and b_t are potentially positive, and finally a phase where $g'(0) \leq X_t$ and $b_t = 0$.

Finally, casual empiricism suggests that as time goes by apprentices not only do less menial work, but also devote their skilled effort to more and more advanced tasks. For example, as noted in Section 2, novices ranging from young cooks to surgeons go through well-defined career stages, initially carrying out the least desirable and simplest tasks and then gradually progressing into more sophisticated activities. This pattern is readily obtained in an extension of our model with multiple skilled tasks that are ordered by difficulty, on the assumption that it is efficient to assign more knowledgeable workers to more difficult tasks.

4.3 Contract length

Corollary 2 shows how effort distortions influence the length of the optimal contract.

Corollary 2 *When the agent's initial participation constraint is slack and the contract*

²²We can show by example that each case can occur.

²³To see that such a g exists, notice first that regardless of g , profits are bounded below by $(\bar{X} - \underline{X}) \underline{a}$, because the principal can ask the agent to work over a period of time using knowledge \underline{X} and effort $q^*(\underline{X})$, promise all knowledge at the end of this period, and choose the period so that the agent earns zero rents. Now restrict to functions g such that $g(0) = 0$ and $g'(0) < \bar{X}$. Because instantaneous output is bounded above, T is bounded away from 0 (otherwise profits would not be bounded away from zero). Hence X_0 is bounded away from \bar{X} . The desired condition is then met whenever $g(0) = 0$ and both $g'(1)$ and $g'(0)$ are close to \bar{X} .

specifies a first-best level of skilled effort at all times, then $T = \frac{1}{r}$; otherwise $T < \frac{1}{r}$.

Note that if skilled effort is first best at all times, ρ_t is always equal to 1. Hence the optimal length satisfies $r \int_0^T dt = rT = 1$, which corresponds to the optimal length in GR. In our model, this case obtains for instance when $f_a(\bar{X}, 0) < g'(1)$, so that even when the agent is fully trained, it is inefficient for them to work on the skilled task.²⁴ In this “menial-effort-only” case, when menial effort becomes more productive (that is, when g is replaced by γg with $\gamma > 1$), the principal asks the agent to work harder, and the ongoing participation constraint then requires that the time path of knowledge becomes *lower* at all dates before T . For this reason, paying for knowledge with menial effort is very different from paying for it with cash. As we show in Section 5.2, when the agent has some initial cash, the principal takes it, and the time path of knowledge is weakly *higher* than without a cash payment, strictly so when the agent has significant cash.

Conversely, if skilled effort is ever higher than first best, we obtain $\int_0^{\frac{1}{r}} \rho_t dt > \frac{1}{r}$ and $T < \frac{1}{r}$. Thus the interaction of skilled effort and training leads to shorter optimal contracts than in GR.

5 Extensions

5.1 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.

²⁴This corresponds to the special case where effort and knowledge are always substitutes.

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. 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 $\dot{u}_t = r[u_t + c(q_t) - \Delta]$ instead of $\dot{u}_t = r[u_t + c(q_t)]$, where $u_t = v(X_t)$. As Δ grows, the agent can be trained more quickly, but is kept longer after being trained. Hence, a more valuable certificate (e.g. a more demanding occupational licence) raises 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.²⁵

²⁵Indentured 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.

5.2 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):²⁶

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 slope of the Pareto frontier is less than 1 in absolute value, it is optimal to ask the agent to buy M dollars worth of knowledge up front. The optimal contract then maximizes profits (above and beyond M) subject to the initial participation constraint $U \geq v(\underline{X})/r + M$. Since a larger M is equivalent to a larger $v(\underline{X})/r$, 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.

²⁶Garicano 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 apprenticeship, 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).

6 Regulating apprenticeships

Regulators have long been interested in protecting apprentices from potential abuse by their masters. In Germany and Switzerland, where apprenticeships are widespread, the government imposes heavy regulation. As Wyman (2017) notes, the apprenticeship system “is jointly governed and funded by the public and private sectors. Furthermore, cantons (state governments), the federal government, trade associations and companies collaborate closely to design curricula, standards, and training programs.” In the U.S., traditional apprenticeships are less common, but various training relationships are regulated nonetheless. For instance, to obtain a license (a form or training certificate), hair stylists “must complete a 900-hour cosmetology curriculum covering everything from chemical hair treatments and pedicures to shaving” (Williams, 1997). Another example is that of medical residencies, where, as we discuss below, regulators have been recently interested in capping the hours worked each week.

The model helps us identify some important issues in evaluating various possible forms of regulation. To frame our discussion, recall that any Pareto-efficient payoffs are uniquely implemented by a (potentially zero) initial knowledge gift followed by the principal’s optimal contract given that initial gift.

Mandated training. The government may attempt to raise surplus by mandating a higher initial knowledge gift, that is by mandating faster training early in the relationship. This may potentially justify attempts by European governments to control the curricula

of apprentices. However, in many cases the government lacks the knowledge necessary to design an optimal curriculum and the means to verify that a knowledge transfer has occurred.

Certificates and government funding. The certificates discussed in Section 5.1 can potentially promote efficient effort and rapid training. Indeed, by requiring that the agent is certified by a specific principal before they are allowed to use their knowledge in the broader labor market, or switch to another principal, the planner directly addresses the agent’s commitment problem. Such a remedy, however, faces a potential drawback outside the model: Once the agent’s has committed to a particular master, the agent has much less of a check on potentially abusive behavior.

The government can also raise surplus by relaxing the agent’s liquidity constraint (see Section 5.2). This motive may justify government funding of apprenticeships, as frequently observed in Europe. It may also justify providing student loans, though as we have seen very small loans may only benefit the principal.

Caps on effort and apprenticeship length. The regulator might also consider a cap on the length of the apprenticeship, a cap on effort, or a combination of the two.²⁷ The next result considers the idealized case where the government can impose both a cap on length and a cap on effort that varies over the life of the apprenticeship.

Proposition 1 *For every Pareto-efficient contract \mathcal{C} there is a cap on contract length and a (time-dependent) cap on effort that together induce the principal to select \mathcal{C} .*

To gain intuition, pick a Pareto-efficient contract \mathcal{C} with length T , effort path \mathbf{q} , and initial knowledge level X_0 . Notice that this contract uniquely maximizes the principal’s profits subject to the agent’s payoff being at least $v(X_0)/r$. Now impose a cap on contract length equal to T and a time varying cap on effort equal to \mathbf{q} . These two caps combined imply that the agent’s payoff can be no lower than $v(X_0)/r$ (per equation 4); hence they induce the principal to select the desired contract.

Notice that over the allowed duration of the apprenticeship, the optimal time-dependent effort cap is equal to the unconstrained principal-optimal effort path. Thus, for any given desired apprenticeship length, one could use historical data on an industry to measure effort levels, working backward from the novice’s graduation date, and use them to set the time-dependent effort cap.

²⁷GR consider caps on apprenticeship duration alone, without considering their impact on the effort path.

In some cases it may not be easy to identify or implement the optimal time-dependent effort cap in Proposition 1. For this reason it is interesting to note that when the optimal contract prescribes effort weakly above the steady-state level, welfare may be increased by regulations that combine a time-invariant effort cap with a limit on the training period, even though such regulations are not Pareto efficient. One can check this hypothesis empirically by using historical data in a given industry to compare the hours of apprentices against the hours of fully trained agents. For instance, medical residencies appear to satisfy this hypothesis, as medical residents routinely work at least 80 hours per week, whereas 90% of fully trained doctors work 70 hours or less (see American Medical Association, 2015).

Proposition 2 *Let $(T, \mathbf{q}, \mathbf{X})$ denote the principal-optimal contract and suppose $q_t \geq q^*(\bar{X})$ for all t .²⁸ Then any cap on contract length $T^{CAP} \leq T$ combined with a constant cap on effort $q^{CAP} = q^*(\bar{X})$ weakly increases surplus, strictly so whenever $T^{CAP} < T$, or $q_t > q^*(\bar{X})$ at some time t , or both.*

When the hypothesis of the proposition is satisfied, an effort cap at the steady state level forces the principal to move the effort path closer to the first-best path (which prior to graduation lies strictly below the steady state level). Moreover, because contract length is capped at the original level T (or lower), the agent faces a less costly apprenticeship and hence can be granted additional knowledge.

Notice that the improvements made by the regulations considered in the two propositions can come in part from an increase in the “knowledge gift” that the principal provides at time 0. In practice, there is a bound on training speed, and this instantaneous gift corresponds to an initial phase where knowledge is raised as quickly as possible to the desired starting level. From the viewpoint of optimal regulation, the government should always allow the principal to devote as much time as she would like to such “pure training phases” where the agent is fully devoted to learning without this time counting toward a limit on the duration of the apprenticeship.²⁹

²⁸One case of the model where this hypothesis is satisfied is when knowledge and effort are substitutes (i.e. $y^*(X, q) = X + q$) for in this case $q^*(X) = q^*(\bar{X})$ for all X . Another case is when $y^*(X, q) = X \cdot q$, c is quadratic, and the exogenous bound on effort is large relative to $q^*(\bar{X})$ (Section 4 of the Online Appendix derives an analytical solution for this case).

²⁹In an earlier version of this paper we explicitly verified that when there is a bound on training speed, under suitable assumptions the principal optimally has such a pure training phase at the beginning of the apprenticeship. And since the gift grows when we cap T (and effort), the pure training phase would also grow.

Both effort caps and caps on length are needed for these results. A cap on apprenticeship length alone, without a cap on effort, accelerates the agent’s graduation, 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.³⁰ 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. Moreover, if the principal faces a cost when training the agent, a sufficiently tight cap may simply drive the principal away from the market. Conversely, a cap on effort without a cap on apprenticeship length 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.

Moreover, as illustrated in the following remark, even very mild caps on length or effort alone can be counter-productive.

Remark 1 *Let $(T, \mathbf{q}, \mathbf{X})$ denote the principal-optimal contract absent regulation. Suppose knowledge and effort are substitutes (i.e. $y^*(X, q) = X + q$) and suppose \underline{X} is strictly above the unconstrained initial knowledge level, and so there is no initial knowledge gift. Then, every cap on contract length $T^{CAP} < T$ that is close to T reduces surplus. Similarly, every cap on effort $q^{CAP} < \sup_t q_t$ that is close to $\sup_t q_t$ reduces surplus.*

(Proof in Appendix B.)

Intuitively, the reason these “light-touch” caps fail is that when \underline{X} is above the unconstrained initial knowledge level, neither a small reduction in contract length nor a small reduction in maximum allowed effort will make the principal want to give the agent a knowledge gift. Thus these interventions lower the principal’s payoff while leaving the agent’s payoff unchanged.³¹

A recent example of regulation is the cap on hours worked by medical residents in the U.S. This cap, which limits hours to 80 per week averaged over a four week period, was first imposed in 2003 by the Accreditation Council from Graduate Medical Education

³⁰This can be seen from the first-order condition for q_T in the proof of Lemma A4, part 2.

³¹One can also find cases where caps on either length or effort raise surplus. For example in GR, because there is no effort, any cap on length $T^{CAP} < 1/r$ raises surplus. Any such cap would also raise surplus after we add effort to the model provided there is an initial knowledge gift and effort is both minimally costly and minimally productive, so that the model remains close to GR. Also, it is easy to show that any constant cap on effort q^{CAP} that is sufficiently close to 1 increases surplus when knowledge and effort are substitutes and there is an initial knowledge gift to begin with, as in this case $T = 1/r$ both before and after the cap.

(ACGME).³² Critics of this 80 hour limit feared that it would hurt patient care and hinder learning (e.g. Carroll, 2016), but in 2011 the ACGME decided to keep this limit. It noted that: “In the 7 years that the GME community has functioned under the 2003 common standards, studies have shown that residents’ clinical exposure, academic achievement, and medical knowledge have remained constant or improved slightly [...] and quality of care appears to have been unaffected or only very slightly improved by the 2003 limits.” (ACGME, 2011, p. 29.)

Our analysis highlights a risk behind such regulation. The cap, to the extent that it is adequately enforced,³³ may cause residency programs to eventually grow in length, and indeed the ACGME recently authorized a pilot program that lengthens family medicine residencies from 3 years to 4 (see Abercrombie et al., 2012).³⁴ A potential remedy, per Proposition 2, is to restrict apprenticeship length using historical data on length prior to the imposition of the weekly hour cap.

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 for the principal for low or no wages.³⁵ The optimal contract is inefficient both because the agent will work inefficiently hard to compensate the principal for their training, and because this training is inefficiently slow. The extent of overwork decreases throughout the apprenticeship, and the agent spends a decreasing amount of time on menial effort. 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

³²The ACGME was motivated by public pressure, evidence of the adverse effect of lack of sleep on performance, and a threat of legislation (ACGME, 2011). It also regulates shift length and requires minimum rest periods.

³³There is been some debate about how well the cap on residency hours has been enforced. See for example Ketcham (2008).

³⁴In defense of the pilot program, Abercrombie et al. (2012) argue that “Duty hour restrictions reduced the available teaching hours to less than 2.5 years [...] How do we produce a competent graduate with the constraints of less time?”

³⁵Appendix A shows that the optimal contract when wages are allowed is to pay the minimum allowed wage. (The minimum for apprentices and young people can be lower than the general minimum.)

which the agent is overworked. Moreover, bounds on effort alone can lead to even slower training, and bounds on apprenticeship length can increase overwork, but total surplus can be increased by a cap on apprenticeship length combined with a judiciously chosen upper bound on hours worked.

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.³⁶

References

1. Abercrombie, Stoney, Karen Hall, Sneha Chacko, Joseph Gravel, Grant Hoekzema, Lisa Maxwell, Michael Mazzone, Todd Shaffer, Michael Tuggy, and Martin Wieschhaus (2012). “Cleared for takeoff: the 4-year family medicine residency pilot.” *Annals of Family Medicine* 10(1): 84-85.
2. Acemoglu, Daron (1997). “Training and Innovation in an Imperfect Labour Market.” *Review of Economic Studies* 64(3): 445-464.
3. Acemoglu, Daron, and Jorn-Steffen Pischke (1998). “Why Do Firms Train? Theory and Evidence.” *Quarterly Journal of Economics* 113(1): 79-119.

³⁶In ongoing work we also extend our analysis to the case where the agent’s utility of consumption is strictly concave. Preliminary results are available upon request.

4. ACGME (2011). *The ACGME 2011 Duty Hour Standard*. Accreditation Council for Graduate Medical Education. Chicago, IL.
5. Akerlof, George (1976). “The economics of caste and of the rat race and other woeful tales.” *Quarterly Journal of Economics*: 599-617.
6. Albuquerque, Rui, and Hugo A. Hopenhayn (2004). “Optimal lending contracts and firm dynamics.” *Review of Economic Studies* 71(2): 285-315.
7. 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).
8. Ayres, James (2014). *Art, Artisans and Apprentices: Apprentice Painters & Sculptors in the Early Modern British Tradition*. Oxbow Books.
9. Barlevy, Gadi, and Derek Neal (2017). “Allocating Effort and Talent in Professional Labor Markets.” *Journal of Labor Economics*, forthcoming.
10. Becker, Gary S. (1964). *Human capital: A theoretical and empirical analysis with special reference to education*. University of Chicago Press.
11. Bonatti, Alessandro, and Johannes Hörner (2017). “Career concerns with exponential learning.” *Theoretical Economics* 12(1): 425-475.
12. Carroll, Aaron E. (2016). “Medical Residents, Misplaced Pride and Saner Hours.” *New York Times* (Feb 15).
13. Chachuat, Benoit (2007). *Nonlinear and Dynamic Optimization: From Theory to Practice*. Automatic Control Laboratory, EPFL, Switzerland.
14. Cisternas, Gonzalo (2017). “Career Concerns and the Nature of Skills.” *American Economic Journal: Microeconomics*, forthcoming.
15. 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.
16. 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.

17. 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.
18. Gapper, John (2014). "Bankers and Lawyers are on an unhealthy treadmill." *Financial Times* (January 15).
19. Garicano, Luis, and Luis Rayo (2017). "Relational Knowledge Transfers." *American Economic Review* 107(9): 2695-2730.
20. 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.
21. Guido, Friebel, and Sergei Guriev. (2006). "Smuggling humans: a theory of debt-financed migration." *Journal of the European Economic Association* 4(6): 1085-1111.
22. Holmström, Bengt (1999). "Managerial incentive problems: A dynamic perspective." *Review of Economic Studies* 66(1): 169-182.
23. Hörner, Johannes, and Andrzej Skrzypacz (2016). "Selling information." *Journal of Political Economy* 124(6): 1515-1562.
24. Jauhar, Sandeep (2015). "The Real Problem With Medical Internships." *New York Times* (Jul 1).
25. Jovinelly, Joann, and Jason Netelkos (2007). *The crafts and culture of a medieval guild*. New York: The Rosen Publishing Group.
26. Katz, Eliakim, and Adrian Ziderman (1990). "Investment in General Training: The Role of Information and Labour Mobility." *Economic Journal* 100(403): 1147-1158.
27. Ketcham, Braham B. (2008). "Medical Resident Maximum Hour Regulations: Overcoming Institutional Resistance for Real Reform." *Health Law and Policy Brief* 2(2): 10-21.

28. Kirzner, Amanda (2014). "Grey's Anatomy' vs. real-life residency: You already know how this turns out." The DO: <https://thedo.osteopathic.org/2014/09/greys-anatomy-vs-real-life-residency/>.
29. Kumar, T. Krishna (1969). "The existence of an optimal economic policy." *Econometrica* (Oct): 600-610.
30. 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.
31. 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.
32. Lane, Joan (1996). *Apprenticeship in England, 1600-1914*. London: Routledge.
33. Li, Jin, Michael Powell, and Rongzhu Ke (2018). "Firm Growth and Promotion Opportunities." Northwestern University, mimeo.
34. Lugert, Verena (2017). *Die Irren mit dem Messer: Mein Leben in den Küchen der Haute Cuisine*. Knaur eBook.
35. Maister, David H. (1993). *Managing the Professional Service Firm*. Free Press, New York, NY.
36. 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.
37. Maslin Nir, Sarah (2015). "The Price of Nails." *New York Times* (May 7).
38. Park, Ryan (2017). "Why So Many Young Doctors Work Such Awful Hours." *The Atlantic* (Feb 21).
39. Schwartz, Robert J., et al. (1992). "Guidelines for surgical residents' working hours: intent vs reality." *Archives of Surgery* 127(7): 778-783.
40. Smith, Adam (1872). *An Inquiry Into the Nature and Causes of the Wealth of Nations: Volume 1*. London: Alex. Murray & Co.

41. 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).
42. Stephan, Paula (2013). "How to exploit postdocs." *BioScience* 63(4): 245-246.
43. 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.
44. Thomas, Jonathan, and Tim Worrall (1994). "Foreign direct investment and the risk of expropriation." *Review of Economic Studies* 61(1): 81-108.
45. Thrupp, Sylvia L. (1989). *The merchant class of medieval London, 1300-1500*. Vol. 72. University of Michigan Press.
46. UK Department for Business, Innovation and Skills (2013). "Apprentices' Pay, Training and Working Hours." *BIS*, UK Government.
47. Williams, Lena (1997). "The Battle of the Braid Brigade." *New York Times* (Jan 26).
48. Wyman, Nicholas (2017). "Jobs Now! Learning From The Swiss Apprenticeship Model." *Forbes* (Oct 20).

8 Appendix A: General Contracts

In this appendix we show that the restricted contracts we considered in the main text are in fact optimal in the larger space of contracts where the principal can pay wages, the agent can save, apprenticeships can last infinitely long, training may be incomplete, a given amount of total effort need not be efficiently allocated between the two tasks, and the ongoing participation constraint can be slack.

We now suppose that the principal can pay nonzero wages at any time, and that the agent has access to a savings account that pays interest r , but has no capital up-front and has no ability to borrow. We do not allow the agent to go into debt, but we do allow the wage payments to be negative – this allows the agent to give money back to the principal once they have accumulated some. We model this by the *liquidity constraint* $\int_0^t e^{-r\tau} w_\tau d\tau \geq 0$ at all t .

We no longer require the graduation date T to be finite. Instead, we set $X_\infty := \lim_{t \rightarrow \infty} X_t$, and say that the agent graduates at time $T = \inf \{t : X_t \geq X_\infty\}$.

At time 0, the principal offers the agent an employment contract, denoted $\mathcal{C} = (T, \mathbf{a}, \mathbf{b}, \mathbf{w}, \mathbf{X})$ where $T \in [0, \infty]$ is the graduation date, and $\mathbf{a}, \mathbf{b}, \mathbf{w}$, and \mathbf{X} are the time paths of a, b, w , and X respectively. 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.³⁷ 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 \mathcal{C} , the principal's and agent's continuation values from date $t \leq T$ onward are now

$$\begin{aligned} \Pi_t(\mathcal{C}) &= \int_t^T e^{-r(\tau-t)} [y_\tau - w_\tau] d\tau \quad \text{and} \\ U_t(\mathcal{C}) &= \int_t^T e^{-r(\tau-t)} [w_\tau - c_\tau] d\tau + e^{-r(T-t)} \frac{1}{r} v(X_T). \end{aligned} \tag{7}$$

The principal's problem is to maximize discounted payoff $\int_0^T e^{-rt} [y_t - w_t] dt$ over con-

³⁷Contracts 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.

tracts $(T, \mathbf{a}, \mathbf{b}, \mathbf{w}, \mathbf{X})$ that satisfy the liquidity constraint and the participation constraint $U_t(C) \geq \frac{1}{r} \max \{\underline{v}, v(X_t)\}$, subject to the technological constraints $a_t, b_t \geq 0$, $a_t + b_t \leq 1$, and $X_t \in [\underline{X}, \overline{X}]$ non-decreasing and right continuous in t , and the technical condition that $a_t + b_t$ is Lipschitz continuous.

The proofs of the following two lemmas are in the Online Appendix.

Lemma A1

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 part 1 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 short period of time in exchange for receiving the remaining knowledge at the end of the contract.

The key to the second part 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.

Lemma A2 *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 $q_t =$*

$a_t + b_t$ across tasks so as to maximize output.³⁸

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.

The principal's problem therefore simplifies to the program of the main text.

9 Appendix B: Proofs

9.1 Optimal Control

We state the principal's problem as an optimal control problem with control variables T and q_t , and state variable $u_t := rU_t$. By assumption, the ongoing participation constraint holds with equality, so X_t is given by $\phi(u_t) := v^{-1}(u_t)$. The terminal value of the state is $u_T = v(\bar{X})$, while its initial value is unknown and subject only to the initial participation constraint $u_0 \geq v(\underline{X})$. Whenever $u_0 > v(\underline{X})$ there is an instantaneous knowledge gift at time 0.

From the definition of U_t we see that $\dot{u}_t = r[u_t + c(q_t)]$. Thus the optimal control problem is:

$$\max_{T, \mathbf{q}} \int_0^T e^{-rt} y^*(\phi(u_t), q_t) dt \quad (8)$$

subject to

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

$$0 \leq q_t \leq 1,$$

where q_t satisfies a Lipschitz condition.

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

³⁸The contract is strictly dominated if it fails to satisfy any of these two properties over a positive-measure fraction of time.

Lemma A3 *A solution to problem (8) exists. Moreover, the following system is necessary for optimality:*

$$\dot{u}_t = ru_t + rc(q_t); \quad \dot{\lambda}_t = \partial\mathcal{L}/\partial u_t \quad (9)$$

$$\partial\mathcal{L}/\partial q_t = 0 \quad (10)$$

$$\eta_t, \gamma_t \geq 0; \quad \eta_t [1 - q_t] = \gamma_t q_t = 0 \quad (11)$$

$$\mathcal{H}_T = 0; \quad \lambda_0 \geq 0; \quad \lambda_0 [u_0 - v(\underline{X})] = 0. \quad (12)$$

Proof. The agent won't work longer than $T^{\max} := \frac{1}{r} \log(v(\bar{X})/v(\underline{X}))$ even if asked to exert 0 effort for all t . Because we restrict to Lipschitz continuous controls, from Kumar (1969) a solution to (8) exists for any fixed $T \in [0, T^{\max}]$. Moreover, an optimal T exists because the principal's optimized profits are continuous in T .

To show the necessity of the system (9)-(12), 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, so the reachability condition of Chachuat (2007) Remark 3.19 is satisfied, and the inequality constraints on q_t are linearly independent, so the necessity of conditions (9)-(12) follows from Chachuat (2007) Theorems 3.18 and 3.33, and Remark 3.23 (on extending to inequality constraints on the terminal condition T). To obtain the transversality condition for λ_0 in condition (12), 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 $\lambda_0 [u_0 - v(\underline{X})] = 0$. ■

Expression (9) contains the state and co-state evolution equations; (10) contains the first-order condition for q_t ; (11) contains the complementary slackness conditions for the Lagrange multipliers η_t, γ_t ; and (12) contains the transversal condition for the terminal time T , where λ_0 is the multiplier for the constraint $u_0 \geq v(\underline{X})$.³⁹

9.2 Preliminary Lemmas

Recall that, for all t , $\rho_t := \rho(\phi(u_t), a_t) = \frac{f_X(X, a_t)}{f_X(X, a^*(X))} \Big|_{X=\phi(u_t)} \geq 0$ and $a_t = a^*(X_t, q_t) := \arg \max_{a \in [0, q_t]} f(X_t, a) + g(q_t - a)$.

³⁹We omit the constraint $u_0 \leq v(\bar{X})$ here because it cannot bind.

Lemma A4

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

$$\lambda_t = e^{-rt} \left[e^{rT} \lambda_T - \int_t^T \rho_\tau d\tau \right]. \quad (13)$$

2. *In every solution to the optimal control problem we have $\lambda_T = e^{-rT}/r$, $q_T = q^*(\bar{X})$, and $\eta_T = 0$.*

Proof. *Part 1.* Using the facts that $v'(\phi(u)) = f_X(\phi(u), a^*(\phi(u)))$ (from the envelope theorem) and $\phi'(u_t) = \frac{1}{v'(\phi(u))}$ (from the implicit function theorem) we obtain $\frac{d}{du} f(\phi(u), a) = f_X(\phi(u), a) \frac{1}{v'(\phi(u))} = \frac{f_X(X, a)}{f_X(X, a^*(X))} \Big|_{X=\phi(u)}$. 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 (13).

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

$$\underbrace{c'(q_T) y^*(\bar{X}, q_T) - \frac{\partial y^*}{\partial q}(\bar{X}, q_T) [v(\bar{X}) + c(q_T)]}_{=: h(q_T)} = -e^{rT} \eta_T [v(\bar{X}) + c(q_T)]. \quad (14)$$

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

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

where the inequality follows from the fact that $c'' > 0$ and at each point of differentiability $\frac{\partial^2 y^*}{\partial q^2} \leq 0$ (since $f_{aa}, g'' \leq 0$). As a result, $h(q_T)$ is strictly increasing. It follows that $\eta_T = 0$; otherwise, we would have $\eta_T > 0$ and $q_T = 1 > q^*(\bar{X})$, and so the left-hand side of (14) would be positive and its right-hand side would be negative. Once we set $\eta_T = 0$ it follows that $q_T = q^*(\bar{X})$ and $\lambda_T = e^{-rT}/r$. ■

Lemma A5 *For any given λ_t the first-order condition for q_t and complementary slackness conditions for η_t, γ_t have a unique solution, denoted $\tilde{q}(\lambda_t), \tilde{\eta}(\lambda_t), \tilde{\gamma}(\lambda_t)$. More-*

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

$$c'(\tilde{q}(\lambda_t)) = \frac{e^{-rt} \frac{\partial y^*}{\partial q}(\phi(u_t), \tilde{q}(\lambda_t)) - \tilde{\eta}(\lambda_t)}{\lambda_t r}, \quad (15)$$

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

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

9.3 Proof of Theorems 1 and 2

Our goal here is to prove that the solutions to the two optimal control problems of the theorems are given by the solutions to the necessary conditions. To do this, we argue that the necessary conditions have a unique solution. Because solutions exist from Lemma A3, this guarantees that the necessary conditions are sufficient.

Rewrite the co-state equation (13) as

$$\lambda_t = e^{-rt} \left[\frac{1}{r} - \int_t^T \rho_\tau d\tau \right], \quad (16)$$

which follows from $e^{rT} \lambda_T = \frac{1}{r}$. Upon rearranging terms, equation (15) yields

$$c'(q_t) = \frac{\frac{\partial y^*}{\partial q}(X_t, q_t) - e^{rt} \eta_t}{r e^{rt} \lambda_t}, \quad (17)$$

$$e^{rt} \eta_t = \max \left\{ 0, \frac{\partial y^*}{\partial q}(X_t, 1) - r e^{rt} \lambda_t c'(1) \right\}.$$

Next, we ignore both the complementary slackness condition for λ_0 and constraint $u_0 \geq v(\underline{X})$, and claim that for any fixed terminal time T , the remaining necessary conditions in Lemma A3, together with the terminal condition $u_T = v(\overline{X})$, have a unique solution, which we denote $u_t^T, q_t^T, \lambda_t^T$ (and ρ_t^T, η_t^T). To prove this claim, write $t = T - s$ for $s \geq 0$,

and use (17) to obtain

$$c'(q_{T-s}) = \begin{cases} \min \left\{ \frac{\partial y^*(X_{T-s}, q_{T-s})}{r e^{r(T-s)} \lambda_{T-s}}, c'(1) \right\} & \text{when } \lambda_{T-s} \geq 0, \\ c'(1) & \text{otherwise.} \end{cases} \quad (18)$$

Second, note that since $c'' > 0$ and $\frac{\partial y^*}{\partial q}$ is weakly decreasing in q , for any given s , X_{T-s} , and λ_{T-s} there is a unique q_{T-s} that satisfies (18). Therefore, because $\dot{u}_t = ru_t + rc(q_t)$ is continuously differentiable in u , there is a unique solution $u_{T-s}^T, q_{T-s}^T, \lambda_{T-s}^T$ to the system

$$\begin{aligned} u_T &= v(\bar{X}), \quad \dot{u}_{T-s} = r[u_{T-s} + c(q_{T-s})], \quad (18), \\ \lambda_T &= e^{-rT}/r, \quad \text{and } \dot{\lambda}_{T-s} = -r\lambda_{T-s} + e^{-r(T-s)}\rho_{T-s} \text{ for } s \geq 0, \end{aligned} \quad (19)$$

where $\rho_{T-s} = \rho(\phi(u_{T-s}), a^*(\phi(u_{T-s}), q_{T-s}))$. (The resulting paths $\eta_{T-s}^T = \tilde{\eta}(\lambda_{T-s}^T)$ and $\rho_{T-s}^T = \rho(\phi(u_{T-s}^T), a^*(\phi(u_{T-s}^T), q_{T-s}^T))$ are also unique.) Notice that the value of T enters this system only as a subindex for $u_{T-s}, q_{T-s}, \lambda_{T-s}$. Therefore, for any given $s \geq 0$ the solution $u_{T-s}^T, q_{T-s}^T, \lambda_{T-s}^T$ is independent of the chosen value of T . Moreover, since $u_0^T = u_{T-s}^T|_{s=T}$, we have $\frac{d}{dT}u_0^T = -\frac{d}{dt}u_t^T|_{t=0} = -r[u_0^T + c(q_0^T)]$. Thus, whenever u_0^T is positive, it is strictly decreasing in T . Notice also that $q_t^T \geq q^*(\phi(u_t^T))$ (since $r e^{rt} \lambda_t = 1 - r \int_t^T \rho_\tau d\tau \leq 1$) and therefore $\rho_t^T \geq 1$.

Proof of Theorem 1

To obtain the solution to the relaxed problem set $v(\underline{X}) = -\infty$ and $\lambda_0^T = 0$. Part 1 of the theorem follows from the equality $\lambda_0^T = 1 - r \int_0^T \rho_{T-s}^T ds = 0$. Part 2 follows from the fact that, for all $0 \leq s \leq T$, we have $r \int_0^s \rho_{T-\tau} d\tau \leq r \int_0^T \rho_{T-\tau} d\tau = 1$. Therefore eq. (18), upon setting $\lambda_t = e^{-rt} \left[\frac{1}{r} - \int_t^T \rho_\tau d\tau \right]$ (from eq. (16)) and rearranging terms, implies that for all $0 \leq s \leq T$,

$$\frac{\partial y^*}{\partial q}(X_{T-s}, q_{T-s}) / c'(q_{T-s}) = \min \left\{ 1 - r \int_0^s \rho_{T-\tau} d\tau, \frac{\partial y^*}{\partial q}(X_{T-s}, 1) / c'(1) \right\}, \quad (20)$$

which is equivalent to (5) because $r \int_0^T \rho_{T-\tau} d\tau = 1$. The claim that effort is efficient at the terminal time T , and exceeds the efficient level at all earlier times, follows from the facts that $\frac{\partial y^*}{\partial q}(X, q^*(X)) / c'(q^*(X)) = 1$ for all X , $\frac{\partial^2 y^*}{\partial q^2} \leq 0$, $c'' > 0$, and $r \int_0^s \rho_{T-\tau} d\tau$ is strictly positive for all $s > 0$.

Finally, because system (19) has a unique solution, the solution to the relaxed problem is unique. ■

Proof of Theorem 2

We claim 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 $u_0^T \geq v(\underline{X})$. Notice that 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 $T \leq \hat{T}$. Hence, 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 $u_0^T = v(\underline{X})$. Hence, T is uniquely given by the equality $u_0^T = v(\underline{X})$. Notice, finally, that for any given \bar{X} , system (19) implies that case (a) arises when \underline{X} is below a threshold \hat{X} , and case (b) arises otherwise.

For part 1 of the theorem, notice that in case (a) the solution is identical to that of the relaxed problem. For part 2 of theorem, notice that in case (b) the solution is a truncated version of the relaxed solution because, for any given $s \geq 0$, the solution $u_{T-s}^T, q_{T-s}^T, \rho_{T-s}^T$ of system (19) is independent of the chosen value of T . ■

9.4 Proofs of Corollaries 1 and 2

Proof of Corollary 1

Since total effort is allocated efficiently between the two tasks, Theorems 1 and 2 imply that

$$\frac{\partial}{\partial a} f(X_t, a_t) / c'(a_t + b_t) = (\text{resp. } \geq) (\text{resp. } \leq) 1 - r \int_t^T \rho_\tau d\tau \quad (21)$$

whenever $0 < a_t < 1$ (resp. $a_t = 1$) (resp. $a_t = 0$),

and

$$g'(b_t) / c'(a_t + b_t) = (\text{resp. } \geq) (\text{resp. } \leq) 1 - r \int_t^T \rho_\tau d\tau \quad (22)$$

whenever $0 < b_t < 1$ (resp. $b_t = 1$) (resp. $b_t = 0$).

Recall that X_t and $1 - r \int_t^T \rho_\tau d\tau$ are strictly increasing in t and continuous; as result a_t and b_t are continuous as well. Recall also that $f_{aa}, g'' \leq 0$ and $c'' > 0$.

We now claim that b_t is strictly decreasing in t whenever $b_t \in (0, 1)$, which in turn implies that b_t is weakly decreasing whenever $b_t \in \{0, 1\}$. Suppose toward a contradiction that there are two times t, s , with $t < s$, such that $0 < b_t \leq b_s < 1$. Then $g'(b_t) \geq g'(b_s)$, and because $1 - r \int_t^T \rho_\tau d\tau < 1 - r \int_s^T \rho_\tau d\tau$, condition (22) implies that $c'(a_t + b_t) > c'(a_s + b_s)$ and therefore $a_s < a_t$. It follows from (21) and (22) that

$$\frac{\partial}{\partial a} f(X_t, a_t) / c'(a_t + b_t) \geq 1 - r \int_t^T \rho_\tau d\tau = g'(b_t) / c'(a_t + b_t)$$

and

$$\frac{\partial}{\partial a} f(X_s, a_s) / c'(a_s + b_s) \leq 1 - r \int_s^T \rho_\tau d\tau = g'(b_s) / c'(a_s + b_s).$$

These two conditions imply that $\frac{\partial}{\partial a} f(X_t, a_t) \geq g'(b_t) \geq g'(b_s) \geq \frac{\partial}{\partial a} f(X_s, a_s)$, which is impossible because $X_s > X_t$ and $a_s < a_t$ imply that $\frac{\partial}{\partial a} f(X_s, a_s) > \frac{\partial}{\partial a} f(X_t, a_t)$. ■

Proof of Corollary 2

Theorems 1 and 2 imply that whenever the initial participation constraint is slack, we have $r \int_0^T \rho_t dt = 1$. If $a_t = a^*(X_t)$ for all t , then $\rho_t = 1$ for all t , and hence $r \int_0^T dt = 1$. Conversely, if $a_t > a^*(X_t)$ for some t , then because a_t and $a^*(X_t)$ are continuous in time, then $\int_0^{\frac{1}{r}} \rho_t dt > \frac{1}{r}$, and hence $T < \frac{1}{r}$. ■

9.5 Proofs of Propositions 1 and 2

Proof of Proposition 1

Let $\mathcal{C} = (T, \mathbf{q}, \mathbf{X})$ be a Pareto-efficient contract. Recall that \mathcal{C} is the (unique) principal-optimal contract such that the agent's time-zero knowledge is no smaller than X_0 , or, equivalently, such that the agent's payoff is no smaller than $v(X_0)/r$.

Now suppose the planner imposes: (1) a cap on contract length equal to T ; and (2) for any given terminal time $T' \leq T$ chosen by the principal, and any $0 \leq s \leq T'$, a cap on effort at time $T' - s$ equal to q_{T-s} . Any feasible contract $\mathcal{C}' = (T', \mathbf{q}', \mathbf{X}')$ that meets

these constraints gives the agent payoff

$$e^{-rT'} v(\bar{X}) / r - \int_0^{T'} e^{-rt} c(q'_t) dt \geq e^{-rT} v(\bar{X}) / r - \int_0^T e^{-rt} c(q_t) dt = v(X_0) / r,$$

where the inequality follows from the facts that $T' \leq T$ and $q'_{T'-s} \leq q_{T-s}$ for all $0 \leq s \leq T'$, and the equality follows from the fact that contract \mathcal{C} satisfies the agent's participation constraint with equality at all times. Hence \mathcal{C}' gives the agent a payoff no smaller than $v(X_0) / r$.

Since \mathcal{C} remains a feasible choice for the principal, it follows from the above observations that \mathcal{C} is the (unique) optimal contract for the principal given the planner's constraints. ■

Proof of Proposition 2

Let $\mathcal{C}' = (T', \mathbf{q}', \mathbf{X}')$ be a principal-optimal contract given the above caps. We claim that $q'_t \geq q^*(X'_t)$ for (almost) all t . To see why, notice that if $q'_t < q^*(X'_t)$ at some time t , the principal can raise her flow profit by raising q'_t to $q^*(X'_t)$ and using a cash payment to compensate the agent for the additional effort cost.⁴⁰

We now claim that $X'_{T'-s} \geq X_{T-s}$ for all $0 \leq s \leq T'$. To see why, notice that because the principal optimally meets the agent's ongoing participation constraint with equality at all times, for all $0 \leq s \leq T'$ we have

$$\begin{aligned} v(X'_{T'-s}) &= e^{-rs} v(\bar{X}) - r \int_0^s e^{-r(s-\tau)} c(q'_{T'-\tau}) d\tau \\ &\geq e^{-rs} v(\bar{X}) - r \int_0^s e^{-r(s-\tau)} c(q_{T-\tau}) d\tau = v(X_{T-s}), \end{aligned}$$

where the inequality follows from the fact that $q'_{T'-\tau} \leq q^{CAP} \leq q_{T-\tau}$ for all $0 \leq \tau \leq s$, and the last equality follows from the fact that the original contract $(T, \mathbf{q}, \mathbf{X})$ also meets the agent's participation constraint with equality at all times.

⁴⁰As can be seen from the proof of Lemma A1, part 2, the principal can further raise her profits by replacing this cash payment with an earlier graduation date.

It follows that \mathcal{C}' delivers surplus

$$\begin{aligned}
& \int_0^{T'} e^{-r(T'-s)} [y^*(X'_{T'-s}, q'_{T'-s}) - c(q'_{T'-s})] ds + e^{-rT'} v(\bar{X}) / r \\
& \geq \int_0^{T'} e^{-r(T'-s)} [y^*(X_{T-s}, q_{T-s}) - c(q_{T-s})] ds + e^{-rT'} v(\bar{X}) / r \\
& \geq \int_0^T e^{-r(T-s)} [y^*(X_{T-s}, q_{T-s}) - c(q_{T-s})] ds + e^{-rT} v(\bar{X}) / r,
\end{aligned} \tag{23}$$

where the first inequality follows from the fact that $q^*(X'_{T-s}) \leq q'_{T-s} \leq q^{CAP} \leq q_{T-s}$ and $X'_{T-s} \geq X_{T-s}$ for all $0 \leq s \leq T'$, and the second inequality follows from the fact that $T' \leq T^{CAP} \leq T$ and $y^*(X_{T-s}, q_{T-s}) - c(q_{T-s}) \leq v(\bar{X})$ for all $0 \leq s \leq T$.

The proposition then follows from noting that: (1) the far right-hand side of (23) is the surplus generated by $(T, \mathbf{q}, \mathbf{X})$; (2) the last inequality in (23) becomes strict whenever $T' < T$; and (3) the first inequality in (23) becomes strict whenever $T' = T$ and $q_{T-s} > q^{CAP}$ for some $0 \leq s \leq T$ (because q_{T-s} is continuous in s , and $q^*(X'_{T-s}) \leq q'_{T-s} \leq q^{CAP}$ and $X'_{T-s} \geq X_{T-s}$ for all $0 \leq s \leq T'$). ■

9.6 Proof of Remark 1

Notice that when $y^*(X, q) = X + q$ we have $\rho_t = 1$ for all t . Recall from Lemma A3 that λ_0 is the multiplier for the initial participation constraint $u_0 \geq v(\underline{X})$, and so $\lambda_0 \geq 0$ and $\lambda_0 [u_0 - v(\underline{X})] = 0$.

Now set $e^{rt} \lambda_t = \lambda_0 + t$ (from Lemma A4, part 1) and combine with Lemma A5 to obtain

$$c'(q_{T-s}) = \min \left\{ \frac{1}{r[\lambda_0 + T - s]}, c'(1) \right\}. \tag{24}$$

Before any cap is imposed, since \underline{X} is above the unconstrained initial knowledge level \widehat{X} in Theorem 2, we have $\lambda_0 > 0$ and $u_0 = v(X_0) = v(\underline{X})$. Moreover, since effort is efficient at the terminal time, we have $\lambda_0 + T = \frac{1}{r}$.

Now impose a cap $T^{CAP} = T - \varepsilon$ for some $\varepsilon > 0$. We claim that whenever ε is small, the initial participation constraint binds. If not, then after the cap is imposed, $\lambda_0 = 0$ and eq. (24) becomes

$$c'(q'_{T^{CAP}-s}) = \min \left\{ \frac{1}{r[T^{CAP} - s]}, c'(1) \right\} = \min \left\{ \frac{1}{r[T - \varepsilon - s]}, c'(1) \right\},$$

where $q'_{T^{CAP}-s}$ is the new optimal effort path. As a result the cap causes effort to jump discontinuously for all $s \leq T^{CAP}$ for which the effort upper bound was originally slack. It follows from the agent's participation constraint (4) that if T^{CAP} is close to T , the initial knowledge level X_0 drops below \underline{X} , a contradiction.

Consider now a cap on effort $q^{CAP} = \sup_t q_t - \varepsilon = q_0 - \varepsilon$. We claim that whenever ε is small, the initial participation constraint binds. If not, then after the cap is imposed, $\lambda_0 = 0$ and eq. (24) becomes

$$c'(q'_{T'-s}) = \min \left\{ \frac{1}{r[T'-s]}, c'(q^{CAP}) \right\},$$

where T' and $q'_{T'-s}$ are the new optimal length and effort path. Because effort remains efficient at the terminal time, we have $T' = 1/r$, and so the cap causes the contract's length to grow discontinuously. It follows from the agent's participation constraint (4) that if q^{CAP} is close to $\sup_t q_t$, the initial knowledge level X_0 drops below \underline{X} , a contradiction.

In sum, whenever ε is small, both types of caps leave the agent's overall payoff u_0/r unchanged. But since both types of caps reduce the principal's profits, they reduce overall surplus as well. ■