

Relational Knowledge Transfers*

Luis Garicano

LSE and CEPR

L.Garicano@lse.ac.uk

Luis Rayo

University of Utah and CEPR

Luis.Rayo@utah.edu

August 17, 2016

Abstract

We study how relational contracts mitigate Becker’s classic problem of providing general (non-firm-specific) human capital when training contracts are incomplete. The firm’s profit-maximizing agreement is a multi-period “apprenticeship” in which the novice is trained gradually over time and eventually receives all knowledge. The firm adopts a “ $\frac{1}{e}$ rule” whereby at the beginning of the relationship the novice is trained, for free, just enough to produce a fraction $\frac{1}{e}$ of the efficient output. After that, the novice earns all additional knowledge with labor. This rule causes inefficiently lengthy relationships that grow longer the more patient the players. We discuss policy interventions.

*We thank Ricardo Alonso, Thomas Bruss, Bob Gibbons, Richard Holden, Hugo Hopenhayn, Larry Samuelson, our discussants Jacques Crémer and Steve Tadelis, and three anonymous referees for valuable suggestions. We also thank participants at the 2013 NBER Working Group in Organizational Economics, the 2015 European Summer Symposium in Economic Theory, and the 2015 CEPR IMO Workshop, as well as seminar participants at Amsterdam, Banxico, CEMFI, Chicago, Kellogg, LBS, Milan, and USC.

1 Introduction

As noted by Becker (1994[1964]), when an expert (master) trains a novice (apprentice), there is no asset to be used as collateral. Moreover, the novice often does not have the means to pay the expert, up front, for the knowledge he or she wishes to acquire. Consequently, the first best allocation, in which knowledge is transferred as fast as technologically feasible, is not profitable for the expert. This problem is present, for example, in traditional apprenticeships as well as in professional partnerships, where novices (or “associates” in a partnership) can walk away at any time with the knowledge they have already acquired. Similarly, in international joint ventures involving technological transfers between a “northern” and a “southern” firm, the latter can potentially ignore formal agreements and establish its own operations.

In this paper, we show that a dynamic self-enforcing contract, in which the novice is trained gradually over time, mitigates Becker’s problem: it allows the novice to eventually acquire all knowledge while at the same time allowing the expert to profit. We are also interested in describing the efficiency and distributional properties of this dynamic arrangement.

We set up a simple model in which an expert and a novice, both of whom are risk-neutral, interact repeatedly over time. The expert (she) has a stock of general-purpose, perfectly-divisible knowledge. The novice (he) has no knowledge, and therefore is not able to produce output; he also has no cash, and therefore is not able to purchase knowledge from the expert. By transferring knowledge, the expert raises the novice’s productivity. The complication, however, is that at any time the novice may choose to leave the relationship with the knowledge already acquired and enjoy the output he is able to produce, on his own, with this knowledge. Since knowledge is non-contractable and general-purpose, the only repercussion is an end to the players’ interactions.

To build a productive relationship in the face of both the novice’s cash constraint and his temptation to abandon the expert, players rely on a self-enforcing agreement in which knowledge is transferred gradually over time. In such an arrangement, the novice is willing to remain in the relationship, working for the expert, as long as the expert has sufficient additional knowledge to grant him in return.

The profit-maximizing contract is an “apprenticeship” (labor-for-training) arrangement in which all knowledge is eventually transferred. After an initial knowledge gift, which jump-starts the novice’s productivity, the novice is asked to work for the expert

and is paid only by means of additional training. Each period, the (present) value of the additional knowledge received by the novice is just high enough to compensate him for the output he gives up while working for the expert. The overall length of this apprenticeship is controlled by the size of the initial knowledge gift, with a larger gift leading to an earlier graduation.

When selecting the knowledge gift, the expert faces the following trade-off: by raising the novice’s productivity, a larger gift allows the expert to more quickly extract revenues from the novice; but a larger gift also reduces the remaining knowledge that the expert is able to sell during the labor-for-training exchange. This trade-off favors a lengthy arrangement in which, despite the novice being trained slowly, the expert is able to charge for most of her knowledge.

We find that, no matter how patient the players, and regardless of the details of the output technology, the profit-maximizing knowledge gift allows the novice to produce, at the beginning of the relationship, a fraction $\frac{1}{e}$ of the efficient output level (where e is the mathematical constant). This “ $\frac{1}{e}$ rule” implies that, for realistic discount rates, the apprenticeship lasts many years. For example, when the annual interest rate is 10% (resp. 5%), training takes approximately 10 years (resp. 20 years) to complete. Regardless of the interest rate, no less than a quarter of all potential output is wasted.

As players become more patient, the apprenticeship grows longer, knowledge is transferred more slowly, and less output is produced while the novice is being trained. The reason is that, when patience increases, knowledge becomes more valuable in the margin, as the novice can use the acquired knowledge during every subsequent period of his life. Consequently, in any given period, the novice is willing to work for the expert in exchange for less additional knowledge; a fact that the expert exploits by (inefficiently) slowing down the speed of training and keeping the novice’s output for longer.

Next, we consider two policy experiments. These experiments are motivated by the expert’s preference for artificially lengthy apprenticeships, as well as by general commentary on real-world masters “exploiting” their apprentices by means of contracts with low wages and slow training (discussed in Section 2). First, we force the expert to pay the novice a minimum wage during training. The result is an efficiency gain: this policy leaves the contract length unaffected (an implication of the expert’s $\frac{1}{e}$ rule) and, at the same time, uniformly accelerates the novice’s training. Second, we force the expert to contain his interactions with the novice within a shorter horizon. The result is also an

efficiency gain: the policy alters the expert's optimal balance between knowledge gifted and knowledge sold in favor of a larger gift and a faster sale. We also illustrate how both of these policies may backfire when the expert does not enjoy rents to begin with.

Finally, we study several extensions of the model. First, we show that when the novice has concave utility, the contract remains very similar except for the fact that expert grants the novice positive and increasing wages while being trained. Such wages represent a compromise between delaying consumption (which allows the expert to more quickly extract output from the novice) and smoothing consumption (which helps the novice endure the apprenticeship). Secondly, we consider some brief extensions of practical interest: the expert facing training costs, the novice arriving with capital, and the novice causing externalities on the expert.¹ These modifications alter the contract exclusively via the size of the initial knowledge gift. Lastly, we show that every Pareto-efficient agreement has the same overall structure as the profit-maximizing one, with the novice's Pareto-weight again affecting only the size of the initial knowledge gift. Taken together, these extensions suggest that the model's core results are robust.

The human capital acquisition literature, since Becker's (1994[1964]) classic analysis, shows that firms, in principle, will not pay for the general human capital of their workers – if they did so, firms would not recoup their investment, as workers can always move to another firm. A large literature has relied on market imperfections to explain, under these circumstances, firms' incentives to train their workers. These imperfections include: imperfect competition for workers (e.g. Stevens, 1994, Acemoglu, 1997, and Acemoglu and Pischke, 1999a,b); asymmetric information about a worker's training (e.g. Katz and Ziderman, 1990, Chang and Wang, 1996, and Acemoglu and Pischke, 1998); and matching frictions (Burdett and Smith, 1996, and Loewenstein and Spletzer, 1998). In our analysis, in contrast, it is the timing of training, with gradual training combined with promises of further training down the road, that supports the knowledge transfer.²

A different literature studies the complementary problem of a firm that, to reward the investments of its workers in specific human capital, attempts to build credible promises.

¹Examples of externalities include an expert partner in a law firm benefitting when a novice associate becomes a more effective problem solver, e.g., Garicano 2000, and an expert firm losing profits when training a novice firm who then becomes a more effective competitor.

²Alternatively, in learning-by-doing models (following, e.g., Heckman, 1971, Weiss, 1972, Rosen, 1972, Killingsworth, 1982, and Shaw, 1989) skill accumulation is a by-product of work. Unlike in these models, our principal has the flexibility to determine the rate at which learning takes place independently of the amount of time the agent spends working.

Prendergast (1993) argues that, when firms can commit to pay different wages across tasks, the promise of promotions provides a solution. Relatedly, Kahn and Huberman (1988) and Waldman (1990) argue that an up-or-out rule leads to credible promises, even if the promoted worker has a similar productivity in all jobs.

Malcomson et al. (2003) study the training of workers using long-term apprenticeship contracts with an initial period of low wages during which the training firm earns rents, allowing it to recover its training costs. They study how asymmetric information, concerning both the worker’s intrinsic ability and the firm’s training costs, which are absent in our model, impact the worker’s training. In their model, all training occurs at the start of the relationship, before the period of low wages is over.³ (In this setting, workers do not leave before the low-wage period is over because their ability is not observed by competing firms.) In our model, in contrast, the timing of training is endogenous, allowing us to study how knowledge transfers are optimally spread out over time.

Our work is also related to the literature on principal-agent models with relational contracts. There, akin to our model, self-enforcing rewards motivate the agent (a few examples of this growing literature are Bull, 1987, MacLeod and Malcomson, 1989, 1998, Baker, Gibbons, and Murphy, 1994, Levin, 2003, Rayo, 2007, Halac, 2012, Li and Matouschek, 2013, Barron et al., 2015).⁴ This literature focuses on eliciting effort from the agent while treating the agent’s skill level as stationary and exogenous. In contrast, we treat the agent’s skill as persistent and endogenous while assuming away effort costs.⁵

Hörner and Skrzypacz (2010) study a separate challenge underlying knowledge transfers: asymmetric information regarding the value of the knowledge to be sold. They show that in an environment with limited enforceability, a privately-informed seller benefits

³In this setting, apprenticeships involve a commitment to future wages (which is not possible in our model). The authors show that a regulator can promote training by subsidizing firms and simultaneously forcing them to offer contracts with longer periods of low wages after training is over (which is possible in their setting because of information asymmetries). In our setting, in contrast, a regulator can increase surplus by forcing firms to limit their knowledge transfers to a shorter training horizon, a consideration absent in Malcomson et al. (2003). In our setting, since training is gradual, a second policy – a minimum wage during training – may also be beneficial.

⁴In an alternative setting, Bar-Isaac and Ganuza (2008) study the effect of training on effort in the presence of career concerns.

⁵The “dynamic enforcement” constraint that governs the provision of self-enforcing incentives takes a different form across the two settings: in the costly-effort setting, this constraint typically indicates that self-enforcing money bonuses cannot exceed the (stationary) future surplus created by the relationship; in the knowledge-transfer setting, it indicates that the output that can be extracted from the novice cannot exceed the (shrinking) value of the knowledge yet to be gained by the novice (which represents only a fraction of future surplus).

from gradual revelation as a way to provide evidence regarding the quality of her information, and therefore raise the price of the information yet to be sold.⁶ In our model, in contrast, the value of information is known to all and gradual transmission is instead a consequence of the buyer being liquidity-constrained – i.e. requiring knowledge to produce output and compensate the seller with it.

Finally, a related literature studies lender/borrower contracting under limited enforceability (e.g. Thomas and Worrall, 1994, Albuquerque and Hopenhayn, 2004, DeMarzo and Sannikov, 2006, Biais et al., 2007, and DeMarzo and Fishman, 2007). Limited enforceability means that the borrower’s access to capital is restricted; therefore, his output can grow at most gradually over time. In this lender/borrower setting, transactions involve a single good (capital), whereas in our setting players trade knowledge for capital (or, equivalently, for labor). As a result, the equilibrium contracts take a different form. In the lender/borrower setting, absent uncertainty, players write debt contracts in which debt payments are enforced via the threat of direct punishments on the borrower (i.e. legal penalties and/or a reduction in the borrower’s access to the productive technology). In our setting, in contrast, after an initial knowledge gift – rather than a loan – players engage in a series of knowledge-for-labor sales, and the reason they remain in the relationship is to benefit from future sales, rather than to avoid punishments.⁷

As Bulow and Rogoff (1989) show, in the lender/borrower setting, self-enforcing debt contracts are only possible when direct punishments are available (otherwise, the agent eventually prefers to unilaterally reinvest his output rather than using it to honor his debt). In our setting, with knowledge being noncontractable and general-purpose, such direct punishments are absent and, yet, are not needed to sustain a productive relationship. Also novel to our setting is the economic trade-off at the heart of the model: the fraction of knowledge that the expert sells, rather than gifts, to the novice.

The rest of the paper is organized as follows. Section 2 describes some stylized facts in expert-novice relationships. Sections 3 and 4 present the baseline model and derive profit-maximizing (and Pareto-efficient) contracts. Section 5 considers policy experiments and Section 6 considers extensions of the baseline model. All proofs are in the Appendix.

⁶Anton and Yao (2002) also consider the sale of information of unknown quality. In their model, to signal quality, the seller reveals part of her information up front. After that, two firms compete to purchase the remaining knowledge in a one-shot transaction.

⁷In both cases, provided he is risk-neutral, the agent postpones all consumption until after output has reached its efficient level. In the lender/borrower setting, foregoing consumption helps the agent more quickly honor his debt; in our setting, it helps the agent more quickly purchase additional knowledge.

2 Some stylized facts

Here we present some empirical observations, concerning knowledge transfers within and between firms, that serve to motivate our analysis. These observations illustrate the difficulties caused by a weak contracting environment and suggest that, often, the resulting knowledge transfers are inefficiently slow.

2.1 Apprenticeships, professional partnerships, and slow knowledge transfers

It has long been observed that apprenticeships may be inefficiently lengthy, and training inefficiently slow. According to Adam Smith, “long apprenticeships are altogether unnecessary... [If they were shorter, the] master, indeed, would be a loser. He would lose all the wages of the apprentice, which he now saves, for seven years together” (Smith, 1863:56). During the industrial revolution, in extreme cases, training would slow to a crawl: “[S]ome masters exploited these apprentices’ helpless situations, demanding virtual slave labour, providing little in the way of food and clothing, and failing to teach the novices the trade” (Goloboy, 2008:3). Regarding musical trainees, McVeigh (2006:184) notes: “Since the master received any earnings from concert appearances, apprentices were inevitably subject to exploitation [...] Other apprentices he set to menial tasks. Burney [the apprentice] recorded with irritation the drudgery he undertook for Arne [the master] in the mid 1740s: Music copying, coaching singers and so on.”

Similar observations are often made of present-day training relationships. According to a UK government inquiry: “Several apprentices reported that they were being used as cheap labour [...] Typical responses from apprentices were that [...] they were used to do menial tasks around the workplace” (Dept. for Business, Innovation and Skills, 2013).

An important example of training relationships are those in professional service firms – e.g. law, consulting, architecture. These firms provide a wide range of general skills to junior consultants, usually called associates (see Richter et al., 2008). While in training, associates “pay their dues” by “grinding” through menial tasks; in the process, rents are extracted from their work in exchange for the promise of future training and eventual promotion (see, for example, Maister, 1993).⁸

⁸Levin and Tadelis (2005) provide an alternative view of partnerships. There, partnerships serve as a commitment device to provide high-quality service in a context of imperfect observability.

Consider law firms. There, as described by numerous blog posts and articles, associates are frequently required to perform time-consuming menial tasks. According to a former litigator, “this recession may be the thing that delivers them from more 3,000-hour years of such drudgery as changing the dates on securitization documents and shuffling them from one side of the desk to the other”... “it often takes a forced exit to break the leash of inertia that collars so many smart law graduates to mind-numbing work” (Slater, 2009). The more time associates spend on menial work, the less time they devote to learning the advanced tasks – such as building a case and rainmaking – that they aspire to perform in the future. As a result, training slows down. As an Australian Justice observes, “young solicitors are being exploited and overworked by law firms that have lost sight of their traditional duty to nurture the next generation of lawyers” (Merritt, 2013).⁹

In closing, the recent documentary “Jiro Dreams of Sushi” provides a vivid example of menial tasks and slow training in the restaurant industry. Jiro, a three-Michelin-star sushi chef, possesses coveted skills. His apprentices must endure years of grueling work: cleaning fish, cooking eggs, massaging octopus meat for 40-50 minutes at a time. As noted by the food writer M. Yamamoto: “When you work for Jiro, he teaches you for free. But, you have to endure 10 years of training.” All the while, the apprentices provide valuable work (all the support work in the restaurant). These lengthy apprenticeships, however, eventually pay off. In the words of Yamamoto: “If you persevere for 10 years, you will acquire the skills to be recognized as a first-rate chef [and to have your own place].”¹⁰

2.2 International joint ventures and limited contractability

International joint ventures between a “northern” firm in a developed country (the expert) and a “southern” firm in a developing country (the novice) frequently involve a technology transfer in exchange for a cash flow. Often, owing to weak institutions in the

⁹Similar patterns of drudgery are found in a variety of industries, such as architecture, politics, and entertainment (e.g. Ingalls, 2015, and Kasperkevic, 2016). For example, in entertainment, in a recent class-action regarding the intern program at Fox Searchlight Pictures, a movie studio, the judge found that the plaintiff “did not receive any formal training or education during his internship. He did not acquire any new skills aside from those specific to Black Swan’s back office, such as how it watermarked scripts or how the photocopier or coffee maker operated”. See US District Court of the Southern District of New York, “Eric Glatt et al., plaintiffs against Fox Searchlight Pictures.” 11 Civ. 6784 (WHP).

¹⁰Unpaid traineeships are a crucial part of the business model in upscale restaurants. For example, at El Bulli, unpaid interns outnumbered paid chefs 2 to 1. Noma has 25 full-time chefs and around 30 interns. Aspiring chefs often learn little while spending long hours on menial tasks. In the words of the famous Swedish chef Niklas Erdstad, the traineeship at El Bulli was “close to slavery [...] You might as well have been picking strawberries or peeling potatoes.” (Fox, 2015).

developing country, the partners cannot rely on legally-enforced contracts. As a result, their relationship becomes analogous to one in which knowledge is transferred between two individuals, with the novice free to walk away at any time with the acquired knowledge.

A notable example is the failed partnership of Danone and Wahaha. Their relationship began in 1996 when Danone, a French drink and yogurt producer, established a joint venture with the Hangzhou Wahaha group, a Chinese producer of milk drinks for children. (See, for example, *Financial Times*, April 2007.)¹¹ For Danone, the venture was a way to profit from the growing Chinese market; for Wahaha, it was a means to learn Danone's technology. Initially, the joint venture was highly successful, contributing 5-6% of Danone's entire operating profits. However, in 2007, after Wahaha learned what it needed, it set up a parallel organization that served its clients outside of the joint venture.

Danone appeared, legally, to have the upper hand, as it owned 51% of the joint venture. However, this apparent power was not real. As noted by the press, "the joint venture depends on Mr. Zong's [Wahaha's boss] continuing cooperation. Not only is he chairman and general manager of the joint venture, but he is the driving force behind the entire Wahaha organization. Furthermore, in China, employees in private enterprises often feel a stronger loyalty to the boss than the organization itself. Winning in the courts or pushing out Mr. Zong, therefore, are not solutions to Danone's problems." (*Financial Times*, April 2007). Workers were strongly behind Zong: "We formally warn Danone and the traitors they hire, we will punish your sins. We only want Chairman Zong. Please get out of Wahaha!" (*Financial Times*, June 2007).¹² In the end, Danone lost all its court battles in China, and with them its trademarks.

The Danone-Wahaha case is far from unique; indeed, anecdotal evidence suggests that these types of disputes are quite common. For example, in a case involving two industrial machinery manufacturers, Ingersoll-Rand claimed that Liyang Zhengchang had breached their joint-venture agreement by manufacturing and selling imitation processing equipment based on Ingersoll-Rand's patents.¹³ Once again, the Chinese authorities sided with the Chinese partner.

In the previous examples, the northern partner appears to have underestimated the weakness of the legal institutions in question; as a result, it failed to appreciate the dynamic inconsistency of the exchange. A case in which the northern partner seems fully

¹¹ "How Danone's China venture turned sour." *FT*. April 11, 2007. By Geoff Dyer.

¹² "Still waters run deep in dispute at Wahaha." *FT*. June 12, 2007. By Geoff Dyer.

¹³ 2000 U.S. Dist. LEXIS 18449.

aware of such challenges is the auto-manufacturing alliance between General Motors (GM) and the Chinese manufacturer SAIC. As GM’s chairman points out: “We have a good and viable relationship and partnership. But to make it work, you have to have needs on both sides of the table” (Wall Street Journal, 2012).¹⁴ GM was careful to provide enough knowledge to make the relationship valuable for SAIC: “SAIC [...] went into the partnership with big dreams but little know-how. Today the companies operate much more like equals.” At the same time, presumably mindful of the self-enforcing nature of the relationship, GM has not yet transferred some of its key knowledge: “GM is holding tight to its more valuable technology. Beijing is eager to tap into foreign auto companies’ clean-energy technologies. But GM doesn’t want to share all its research with its Chinese partner”. Indeed, “SAIC could use GM expertise and technology to transform itself into a global auto powerhouse that challenges [GM] down the road.”

3 Baseline model

There are two players: an expert (she) and a novice (he). Players interact over infinite periods $t = 0, 1, \dots$ and discount future payoffs using a common interest rate $r > 0$. Let $\delta = \frac{1}{1+r}$ denote the players’ discount factor. The expert possesses one unit of general-purpose knowledge. This knowledge is perfectly divisible, does not depreciate, and can be transferred from the expert to the novice at any speed desired by the expert.

Let $X_t \in [0, 1]$ denote the novice’s stock of knowledge at the beginning of period t , and let $x_t = X_{t+1} - X_t$ denote the additional knowledge transferred during period t . At first, the novice has no knowledge ($X_0 = 0$). For the time being, to highlight the expert’s desire for an artificially slow knowledge transfer, we assume that the expert faces no costs when training the novice.

During period t , the novice produces output $y_t = f(X_t)$, with f continuous and increasing.¹⁵ We assume that $f(0) = 0$ and so, in period 0, knowledge can be transferred but no output is yet produced. One interpretation of the function f is that the novice’s output originates from a variety of tasks, with more valuable tasks requiring more knowledge. As the novice acquires knowledge, he efficiently spends less time on menial tasks

¹⁴“Balancing the Give and Take in GM’s Chinese Partnership.” WSJ. August 19, 2012. By Sharon Terlep.

¹⁵Notice that the additional knowledge x_t learned in period t is not put to use until the following period. The alternative assumption ($y_t = f(X_t + x_t)$) delivers the same qualitative results and very similar quantitative results (see Supplement).

and more on advanced ones.

Each period, the novice may either work for the expert or work for himself. Since knowledge is general, output is the same in both cases. In what follows we assume that, unless players separate, the novice works for the expert. (As we shall see, this assumption is without loss.) The expert earns $f(X_t)$ and compensates the novice by means of a money transfer $w_t \in \mathbb{R}$, which we call a wage, and a transfer of additional knowledge $X_{t+1} - X_t$.

At the beginning of time, players agree on a *relational contract*: a self-enforcing agreement that specifies, for each period, a knowledge stock X_t and a wage w_t , conditional on the players remaining together. We denote a relational contract by $\mathcal{C} = (X_t, w_t)_{t=0}^{\infty}$ (or equivalently by $\mathcal{C} = (y_t, w_t)_{t=0}^{\infty}$). For brevity, we call \mathcal{C} a contract.

Both players are risk-neutral and care only about the present value of the money they earn. For any given contract \mathcal{C} , let $\Pi_t(\mathcal{C})$ and $V_t(\mathcal{C})$ denote, respectively, the expert's and novice's continuation payoffs from the standpoint of the beginning of period t . These continuation payoffs are given by

$$\Pi_t(\mathcal{C}) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} [f(X_{\tau}) - w_{\tau}] \quad \text{and} \quad V_t(\mathcal{C}) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} w_{\tau}.$$

At the beginning of each period, players are free to walk away from the relationship. When either player walks away, the expert earns no money from that period onwards and the novice earns the present value of all output he can produce with the knowledge already acquired, namely, $\frac{1}{1-\delta} f(X_t)$. Therefore, in order for a contract to be honored, it must satisfy the following two incentive constraints (the first one for the expert and the second one for the novice):

$$\Pi_t(\mathcal{C}) \geq 0 \quad \text{for all } t, \tag{1}$$

$$V_t(\mathcal{C}) \geq \frac{1}{1-\delta} f(X_t) \quad \text{for all } t. \tag{2}$$

We assume that players cannot walk away from the relationship mid-period. In other words, players can commit to honor single-period transactions – in which they trade labor for wages and training – but have no commitment power beyond that.¹⁶

¹⁶The following two alternative settings deliver the same results as the setting above: (1) the expert has full commitment power and selects \mathcal{C} up front; and (2) the expert can only commit to one-period contracts and, at the beginning of each period, proposes – via a take-it-or-leave-it offer – a contract (x_t, w_t) for that period.

The novice has no access to credit and begins the relationship without any cash. (This assumption implies that the novice cannot simply buy all knowledge from the expert upfront.) As a result, the contract \mathcal{C} must also satisfy the following liquidity constraint:

$$\sum_{\tau=0}^t (1+r)^{t-\tau} w_{\tau} \geq 0 \text{ for all } t, \quad (3)$$

which states that the novice's cumulative earnings up to any given date t (including interest) must be non-negative.

We have assumed that, unless players renege on the agreement, the novice works for the expert. This assumption is without loss because a contract in which the novice works for himself in any given period t is equivalent to a contract in which the novice instead works for the expert during that period and earns a wage w_t equal to his output y_t .

Throughout most of the paper, we assume that the expert chooses the contract \mathcal{C} with the goal of maximizing her profits (namely, the expert has full bargaining power). This case is relevant, for instance, when the market in question has a large supply of potential novices and each expert is able to train only a limited number of them. The expert's profit-maximization problem is

$$\begin{aligned} \max_{\mathcal{C} = (X_t, w_t)_{t=0}^{\infty}} \quad & \sum_{t=0}^{\infty} \delta^t [f(X_t) - w_t] \\ \text{s.t.} \quad & (1), (2), (3) \end{aligned} \quad (I)$$

(and subject to the knowledge stock $X_t \in [0, 1]$ being nondecreasing over time). At the end of Section 4, we also study the broader set of Pareto-efficient contracts that maximize a weighted sum of the two players' payoffs.

Notice that since the players' combined surplus is $\sum_{t=0}^{\infty} \delta^t f(X_t)$, the first-best allocation calls for a full knowledge transfer in period 0. While this allocation is ideal from the perspective of the novice, it creates no value for the expert: after receiving all knowledge, the novice would only work for the expert if he is paid a wage at least equal to output.

Preliminaries

The initial period $t = 0$ is used by the expert to provide the novice an initial level of training (which raises his knowledge stock from zero to X_1) in order to jump-start the novice's productivity. Since the novice has no cash, the expert must provide this initial

training for free. We refer to X_1 , the magnitude of such training, as the initial *knowledge gift*.

Once the initial knowledge gift is in place, the expert can profit from the novice by trading knowledge for work, but is bound by the novice's incentive constraint. After rearranging terms, this constraint tells us that from any period t onwards, the overall profits the expert can extract from the novice cannot be greater than the value of all remaining knowledge transfers:¹⁷

$$\sum_{\tau=t}^{\infty} \delta^{\tau-t} \underbrace{\frac{1}{r} [f(X_{\tau+1}) - f(X_{\tau})]}_{\text{value of period-}\tau \text{ knowledge transfer}} \geq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \underbrace{[f(X_{\tau}) - w_{\tau}]}_{\text{period-}\tau \text{ profits}}. \quad (2')$$

The term $\frac{1}{r} [f(X_{\tau+1}) - f(X_{\tau})]$ measures the (present) value of the knowledge $X_{\tau+1} - X_{\tau}$ transferred in period τ . (This term is multiplied by $\frac{1}{r}$ because the additional knowledge permanently raises the novice's productivity.) Notice that since knowledge is general, the novice is only willing to work for future knowledge transfers, not for past ones.

This constraint exposes the expert's predicament: a larger knowledge stock in the hands of the novice means that, each period, he can afford to purchase more knowledge with his work; but it also means the expert is left with less knowledge to sell.

Finally, if the expert completes the knowledge transfer in finite time, we say that the novice graduates. Let the novice's *graduation date* T be the first period in which knowledge is no longer transferred. Notice that from period T onwards, the players' combined incentive constraints require that the novice earns a per-period wage equal to output (namely, $w_t = f(X_T)$).

4 Profit-maximizing contracts

Here we derive the expert's optimal contract. To build intuition, we begin by studying simple two-period contracts. We then turn to the general case in which the expert is free to choose a contract of any length.¹⁸

¹⁷This expression is obtained from (2) by adding the term $\sum_{\tau=t}^{\infty} \delta^{\tau-t} y_{\tau}$ to both sides, noting that $\sum_{\tau=t}^{\infty} \delta^{\tau-t} [y_{\tau} - y_t] = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \frac{1}{r} [y_{\tau+1} - y_{\tau}]$, and rearranging terms.

¹⁸Throughout, to avoid knife-edge cases in which there is more than one solution, we assume δ is "generic" in the sense that $\frac{1}{1-\delta} \neq n$ for all $n \in \mathbb{N}$.

4.1 Benchmark: two-period contracts

Here we consider two-period contracts in which, by assumption, the expert transfers all knowledge by the end of period 1. These are the shortest contracts that allow the expert to profit. From period 2 onward the novice earns a wage equal to output. As a result, the expert's only source of profits is the output produced in period 1.

The expert must decide how to split the knowledge transfer between periods 0 and 1. This decision is captured by the size of the initial knowledge gift X_1 . A larger gift raises the novice's period-1 productivity – and therefore his ability to purchase additional knowledge with his work – but it also lowers the knowledge $1 - X_1$ that is left for the expert to sell. The expert must also select a period-1 wage w_1 . (In period 0, the expert optimally pays zero wages.)

The expert's problem is to maximize her period-1 profits:

$$\begin{aligned} \max_{X_1, w_1} \quad & f(X_1) - w_1 \\ \text{s.t.} \quad & w_1 \geq 0, \\ & \underbrace{\frac{1}{r} [f(1) - f(X_1)]}_{\text{value of remaining knowledge}} \geq \underbrace{f(X_1) - w_1}_{\text{net output}}. \end{aligned}$$

The last constraint is the novice's period-1 incentive constraint. It requires that the novice be transferred knowledge $1 - X_1$ with a value no smaller than the output he gives up, net of wages, for the expert. (All other constraints in the expert's original problem (I) are automatically met.)

The solution, derived in Figure 1, is to set $w_1 = 0$ and meet the incentive constraint with equality. Namely, the expert just barely compensates the novice for his labor, and compensates him exclusively by means of new knowledge: while positive wages can be used to raise the novice's gross output, they are not cost-effective.

The profit-maximizing knowledge gift satisfies

$$\frac{f(X_1)}{f(1)} = \delta.$$

When players meet frequently, or are otherwise patient, almost 100% of knowledge is gifted. Therefore, the arrangement is close to first best. The reason is that, when it is used frequently, knowledge has great value, allowing the transfer $1 - X_1$ to be small.

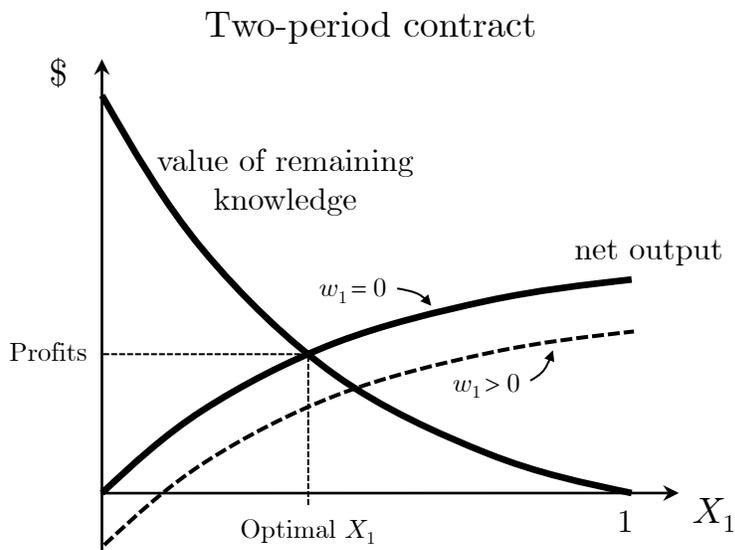


Figure 1: two-period contract. As a function of X_1 , the value of the remaining knowledge is decreasing and net output is increasing. For any given wage w_1 , the optimal X_1 is at the intersection of the two functions (where the novice’s incentive constraint binds). Raising w_1 shifts the net output function downward, reducing both the amount of knowledge sold and the expert’s equilibrium profits.

This two-period arrangement, however, is of limited value for the expert: having only one period to exchange knowledge for work, she pockets only one period of output. Next, we turn to longer (multi-period) arrangements. In these, the expert is able to lower the knowledge gift and trade more knowledge for work. As a result, she pockets multiple periods of output. (In this case, as we shall see, the arrangement is no longer close to first best. Moreover, as players become more patient, the overall knowledge transfer slows down, rather than speeding up.)

4.2 General case: multi-period contracts

Here we return to the general case in which the expert is free to train the novice over any number of periods. We begin by ruling out contracts in which training takes infinitely long:

Lemma 1 *In every profit-maximizing contract, the novice graduates in finite time and is transferred all available knowledge, namely, $X_T = 1$ for some T .*

Intuition is as follows:

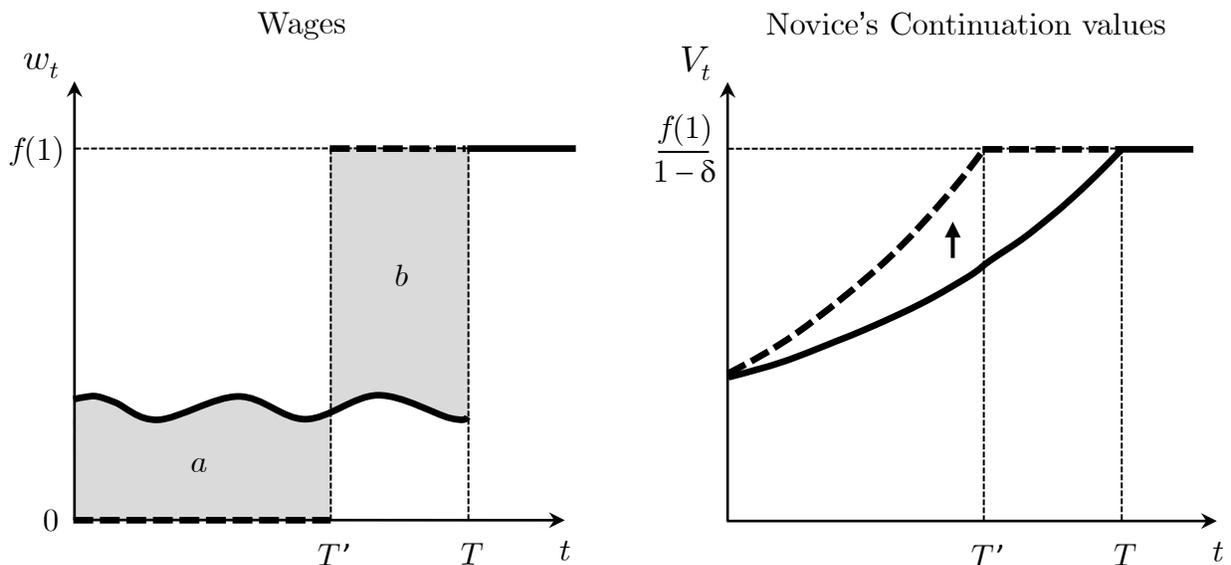


Figure 2: zero wages before graduation. Original contract prescribes graduation T and positive wages before graduation (bold functions). New contract prescribes graduation T' and zero wages before graduation (dashed functions). In both contracts, the novice receives the same total payoff – namely, after discounting, area $a =$ area b .

- *Finite graduation date.* Suppose instead that training takes infinitely long. Recall that from any date onward, the overall profits the expert can extract from the novice are no greater than the value of the knowledge yet to be transferred. Moreover, since knowledge is finite, this value must necessarily approach zero as time goes by. As a result, once the novice can afford to buy all remaining knowledge with a single period of work, the expert can choose to end the contract early and sell all this knowledge at once. By doing so, the expert benefits from earlier revenues.
- *Full knowledge transfer.* Since knowledge has positive value, the expert profits from selling it all to the novice, in exchange for his work.

Next, we show that before graduation the novice earns zero wages and is instead compensated through additional training:

Lemma 2 *In every profit-maximizing contract, the novice earns zero wages before graduation. As a result, during training, the novice is paid for his work exclusively by means of additional knowledge.*

For intuition, see Figure 2. The bold functions represent an arbitrary contract with positive wages prior to graduation T . The dashed functions represent an alternative contract with no pre-graduation wages and an earlier graduation T' . In present value, both contracts prescribe the same total wage bill. In the new contract, however, because wages are shifted toward the future, the novice has greater incentives to remain in the relationship while being trained (i.e. his continuation values V_t grow and his incentive constraints $V_t \geq \frac{1}{1-\delta} f(X_t)$ are relaxed). As a result, the expert can trust the novice with additional knowledge and, since the wage bill is the same as before, she can pocket all additional output.

Notice that once all pre-graduation wages are set to zero, the novice's incentive constraints become (using (2')):

$$\sum_{\tau=t}^{T-1} \delta^{\tau-t} \frac{1}{r} [f(X_{\tau+1}) - f(X_{\tau})] \geq \sum_{\tau=t}^{T-1} \delta^{\tau-t} f(X_{\tau}).$$

We learn that in order to retain the novice, it suffices that period-by-period the knowledge transfer $X_{\tau+1} - X_{\tau}$ has a value equal to the novice's output $f(X_{\tau})$. When instead the value of the period- τ knowledge transfer exceeds $f(X_{\tau})$, we say that a portion of this transfer is gifted.

We are now ready to describe the structure of every profit-maximizing contract:

Proposition 1 *In the baseline model, every profit-maximizing contract has the following structure. In period 0, the novice receives a knowledge gift. Next, during every period after the gift and before graduation, the novice works for zero wages and the value of the additional knowledge he learns is equal to the output he produces for the expert:*

$$\underbrace{\frac{1}{r} [f(X_{t+1}) - f(X_t)]}_{\text{value of knowledge transfer}} = f(X_t). \quad (4)$$

This training process continues until 100% of knowledge has been transferred, after which the novice graduates and earns, each period, a wage equal to the maximum output $f(1)$.

Proposition 1 tells us that a profit-maximizing contract is a type of ‘‘apprenticeship’’ in which the novice receives no wages while in training. After an initial knowledge gift, the novice receives just enough additional knowledge, every period, to compensate him

for the output he gives up. The overall length of the apprenticeship is controlled by the size of the knowledge gift. Since a larger gift increases the novice's productivity – and all subsequent knowledge transfers – it leads to an earlier graduation.

Equation (4) in the proposition tells us that, beyond period 0, the novice does not receive any knowledge gifts. Intuitively, recall that the novice does not require any gifts to endure his training. As a result, the only benefit the expert derives from such gifts is a higher productivity before graduation. But since an earlier gift raises productivity sooner, it is best to place all gifts up front, at the beginning of the relationship.¹⁹

Before we proceed, two remarks are in order:

1. When selecting the initial knowledge gift, the expert faces the following trade-off: by raising the novice's productivity, a larger gift means that the novice can afford to buy (with his work) more knowledge per period of training, leading to earlier revenues for the expert, but it also means that the expert has less knowledge left to sell overall. As we shall see, this trade-off results in lengthy contracts in which significant output is wasted.
2. Equation (4) can be written more compactly as $\frac{1}{r}\Delta y_{t+1} = y_t$, where $\Delta y_{t+1} = f(X_{t+1}) - f(X_t)$. Consequently, after the initial knowledge gift, output grows at rate r until it reaches its maximum value. To see why, notice that for each dollar of output Δy_{t+1} produced with new knowledge, the novice gains, in present value, $\frac{1}{r}$ dollars. Therefore, for each dollar of output y_t produced today, it suffices to grant the novice, as compensation, new knowledge worth r dollars per-period.

4.3 Profit-maximizing contracts in closed form: a $\frac{1}{e}$ rule

Here we derive the profit-maximizing knowledge gift and apprenticeship length. We begin by expressing the expert's problem in reduced form, as a function of X_1 . In what follows, we make use of both the effective interest rate r and the nominal (instantaneous) interest rate r_0 (satisfying $e^{r_0} = 1 + r$).

Let $\mathcal{T}(X_1)$ denote the novice's graduation date. Recall that after X_1 has been gifted, output grows at rate r until all knowledge has been transferred. Therefore, the novice's

¹⁹Formally, when the novice receives a knowledge gift during period $t \geq 1$, namely, $\frac{1}{r}[f(X_{t+1}) - f(X_t)] > f(X_t)$, his incentive constraint for that period is slack. As a result, the expert can increase profits by raising $f(X_t)$ while holding $f(X_{t+1})$ constant, which in turn is achieved by lowering x_t and increasing x_{t-1} – namely, by moving the knowledge gift one period ahead.

last period of training, $\mathcal{T}(X_1) - 1$, corresponds to the number of periods of compound growth at rate r required for output to reach $f(1)$, starting from level $f(X_1)$.²⁰

During each period after the initial gift and before graduation, since output grows at rate r , the expert's discounted revenues $\delta^{t-1} f(X_t)$, measured from the standpoint of period 1, remain constant and equal to $f(X_1)$. Consequently, the expert's reduced-form problem is

$$\max_{X_1} \underbrace{f(X_1)}_{\text{discounted revenues per period}} \cdot \underbrace{[\mathcal{T}(X_1) - 1]}_{\text{\# periods before graduation}}, \quad (\text{II})$$

subject to the graduation date being an integer. This simple objective function summarizes the expert's trade-off: a higher gift raises the novice's productivity and so the expert enjoys higher revenues during each period while the novice is being trained (first term), but it also means that the expert has less knowledge left to trade for labor – while also trading more of this knowledge every period – and therefore the novice graduates sooner (second term). Proposition 2 describes the solution:

Proposition 2 *Up to an integer constraint for the novice's graduation date, the profit-maximizing knowledge gift X_1 satisfies*

$$\frac{f(X_1)}{f(1)} = \frac{1}{e},$$

where e is Euler's number. As a result, the novice's graduation date is $\frac{1}{r_0} + 1$, which is decreasing in r_0 (namely, increasing in δ) and independent of the production technology.²¹

This result tells us that no matter how patient the players, and regardless of the details of the output technology, the expert optimally balances her conflicting goals by allowing the novice to produce, at the start of the apprenticeship, a share $\frac{1}{e}$ of the efficient output level. Indeed, upon solving for the value of $\mathcal{T}(X_1) - 1$, the expert's problem is equivalent to that of maximizing the average logarithm $\frac{1}{z} \log z$ of the output ratio $z = \frac{f(1)}{f(X_1)}$. The maximum is achieved when this ratio is e .

This solution is reminiscent of the solution to the “secretary problem” in which a recruiter who faces a queue of job applicants of unknown quality must decide what fraction

²⁰Namely, $\mathcal{T}(X_1) - 1$ satisfies $\frac{f(1)}{f(X_1)} = (1 + r)^{\mathcal{T}(X_1) - 1}$.

²¹When $\frac{1}{r_0} + 1$ is not an integer, the knowledge gift is adjusted until the graduation date equals the integer directly above/below $\frac{1}{r_0} + 1$. When players meet frequently (r_0 is close to zero) such adjustment is small. (An analogous observation applies to Corollaries 1 and 3 and to the extensions in Section 6.2.)

of applicants to sample before making any hiring decision (e.g. Bruss, 1984). As the total number of applicants tends to infinity, the optimal sample converges to a fraction $\frac{1}{e}$ of all applicants (a result sometimes called a “ $\frac{1}{e}$ law”). The two problems, however, do not appear to have any direct economic link. In addition, unlike in the secretary problem, we obtain a $\frac{1}{e}$ rule for transactions of finite duration.²²

Notice that the profit-maximizing apprenticeship is longer, and knowledge is transferred more slowly, the more patient the players. Intuitively, when patience increases knowledge becomes more valuable in the margin – as the novice can use what he learns throughout the rest of his life. As a result, the novice is willing to work, each period, for less additional knowledge. The expert takes advantage of this fact by stretching out the novice’s training and pocketing his output for longer.

Consistent with real-world practices noted in Section 2, the expert’s $\frac{1}{e}$ rule causes lengthy apprenticeships. To illustrate, suppose players meet n times per year and let $r_A = n \cdot r_0$ denote the annual nominal interest rate. Then, beyond the initial knowledge gift, the novice’s training takes $\frac{1}{r_A}$ years to complete – for example, 10 years when $r_A = 10\%$ and 20 years when $r_A = 5\%$.²³ Finally, the deadweight loss caused by the expert is at least $\frac{e-2}{e} \cdot \frac{1}{r} f(1)$ dollars. Namely, a share no smaller than $\frac{e-2}{e} \simeq 26\%$ of the highest attainable surplus $\frac{1}{r} f(1)$.²⁴

We now consider the continuous-time version of the model (which other than eliminating the integer constraint for T delivers identical results) and the broader class of Pareto-efficient contracts.

Continuous-time model. Suppose time t runs continuously. Accordingly, $f(X_t)$ and w_t represent the instantaneous output and wages at time t . As before, T is the date of graduation and $\delta = e^{-r_0}$. (We start counting time at $t = 1$ so that, as in the baseline model, $t = 1$ is the first period with positive output.)

As we show in Proposition 4 in the Supplement, the profit-maximizing contract is akin to that in the baseline model: After the novice is gifted knowledge X_1 , his output grows continuously at rate r_0 until it reaches $f(1)$. (Namely, analogous to (4), at each instant the value of the knowledge transfer $\frac{1}{r_0} \frac{d}{dt} f(X_t)$ equals $f(X_t)$.) As a result, output is $f(X_t) = \delta^{T-t} f(1)$ for all $t \leq T$. Moreover, wages are zero before T .

²²We are grateful to Thomas Bruss for providing insights into the secretary problem.

²³Beyond the gift, training consumes $\frac{1}{n} (\mathcal{T}(X_1) - 1)$ years, which from Proposition 2 equals $\frac{1}{n \cdot r_0}$.

²⁴Measured in date-0 dollars, the deadweight loss – output lost relative to first best – is $\delta \sum_{t=1}^{r_0-1} [\delta^{t-1} - \frac{1}{e}] f(1)$. As r falls, this expression grows, but it grows slower than $\frac{1}{r} f(1)$ does.

Notice, therefore, that the expert's problem boils down to maximizing a simple objective (expressed now as a function of T):

$$\underbrace{\int_1^T \delta^{t-1} f(X_t) dt}_{\text{output while novice is in training}} = \underbrace{\delta^{T-1} f(1)}_{f(X_1)} \cdot (T - 1).$$

This objective is maximized by setting $T - 1 = \frac{1}{r_0}$. Consequently, the novice graduates at date $\frac{1}{r_0} + 1$ and the knowledge gift satisfies $\frac{f(X_1)}{f(1)} = \frac{1}{e}$, exactly as in the baseline model – but with no integer constraint for T .

Pareto-efficient contracts. Here we describe the broader class of Pareto-efficient contracts. Namely, contracts that maximize a weighted sum $\lambda V_0(\mathcal{C}) + \Pi_0(\mathcal{C})$ of the players' payoffs, subject to the same constraints as the original problem (I). The parameter $\lambda \geq 0$ is the novice's Pareto weight.²⁵ Note that this exercise describes the outcome of maximizing the payoff of one player subject to guaranteeing a given payoff for the other. (For example, a guaranteed payoff for the expert may allow her to recover a fixed cost; and a guaranteed payoff for the novice may serve as compensation for a forgone opportunity.)

Corollary 1 *Every Pareto-efficient contract satisfies the properties in Proposition 1. Moreover, up to an integer constraint for the novice's graduation date, it prescribes a knowledge gift such that*

$$\frac{f(X_1)}{f(1)} = \min \left\{ \frac{1}{e^{1-\lambda A}}, 1 \right\}$$

where $A = \frac{r_0}{1-e^{-r_0}}$. (When r_0 is small, $A \approx 1$.)

In these contracts, the novice obtains rents only by means of the knowledge gift, rather than wages, and this gift is concentrated up front. This feature ensures that the novice's productivity is raised as early as possible. (The constant A , which arises from the fact that time is discrete, is slightly greater than 1 and converges to 1 when the length of each period converges to zero.)

We learn that the knowledge gift is increasing in λ until 100% of knowledge is gifted, which occurs when λA reaches 1 (for which $\lambda = 1$ suffices). In such case, the allocation is first best in the sense of maximizing the players' combined surplus.²⁶

²⁵Below, to guarantee a unique solution, we assume that $\frac{1-\lambda}{1-\delta} \neq n$ for all $n \in \mathbb{N}$.

²⁶The allocation of surplus ranges from an approximately 50-50 split (which is exact when time is continuous) to 100% of surplus going to the novice.

5 Policy experiments

Governments are interested in encouraging firms to offer apprenticeships that grant significant benefits to apprentices. For instance, in a recent meeting in Guadalajara, Mexico, the G20 ministers declared themselves committed to “promote, and when necessary, strengthen quality apprenticeship systems that ensure high level of instruction [...] and avoid taking advantage of lower salaries” (OECD, 2012).²⁷ As the OECD put it, “Quality apprenticeships require good governance to prevent misuse as a form of cheap labour.”²⁸

Motivated by such concerns, we consider two policy experiments: a minimum wage during training and a limit on the apprenticeship’s duration. The discussion that follows presumes that the expert earns sufficient rents from the relationship that she remains interested in training the novice even after the loss in profits caused by these policies. If instead the expert earns no rents, the policies may easily backfire, as illustrated below.

Minimum wage. Suppose a planner forces the expert to pay the novice, during each period of the relationship, a wage no smaller than w_{\min} . Assume that $0 < w_{\min} < f(1)$.

Corollary 2 tells us that the profit-maximizing contract retains the basic properties of the apprenticeship characterized in Proposition 1:

Corollary 2 *When the expert is required to pay a minimum wage w_{\min} , every profit-maximizing contract has the following structure. In period 0, the novice receives a knowledge gift. Next, during every period after the gift and before graduation, the novice receives wage w_{\min} and the value of the additional knowledge he learns is equal to the net output $f(X_t) - w_{\min}$ he produces for the expert:*

$$\underbrace{\frac{1}{r} [f(X_{t+1}) - f(X_t)]}_{\text{value of knowledge transfer}} = f(X_t) - w_{\min}. \quad (5)$$

This training process continues until 100% of knowledge has been transferred, after which the novice graduates and earns, each period, a wage equal to the maximum output $f(1)$.

²⁷“OECD Note on ‘Quality Apprenticeships’ for the G20 Task Force on Employment.” September, 2012. (<http://www.oecd.org/els/emp/OECD%20Apprenticeship%20Note%2026%20Sept.pdf>.)

²⁸Governments have long been interested in regulating apprenticeships. See, for example, Malcomson et al. (2003) for a discussion and Elbaum (1989) for a historical perspective. Malcomson et al. consider a type of regulation (which we do not consider) whereby firms are forced to pay low wages over a minimum time period after training is over. They show that this seemingly counter-intuitive regulation may be beneficial when information asymmetries prevent workers from leaving the firm, and the firm is capable of committing to future wages.

As in the baseline model, the expert concentrates the knowledge gift in the first period, after which she extracts all surplus from the relationship. The difference is that, since the novice is compensated with a combination of money and new knowledge, the minimum wage partially crowds out the transfer of new knowledge. Consequently, for any fixed knowledge gift, the policy delays the novice's graduation.

We now express the expert's problem in reduced form, with X_1 serving as her choice variable. Recall that in the baseline model, after the initial gift, output grows at rate r until training is complete. Under the minimum wage policy, it is net output $f(X_t) - w_{\min}$ that grows at rate r (which can be seen from (5)). As a result, throughout the training process, the expert's discounted profits remain constant and equal to $f(X_1) - w_{\min}$ (in period-1 dollars). Letting $\mathcal{T}(X_1, w_{\min})$ denote the novice's graduation date, the expert's reduced-form problem is

$$\max_{X_1} \underbrace{[f(X_1) - w_{\min}]}_{\text{discounted revenues per period}} \cdot \underbrace{[\mathcal{T}(X_1, w_{\min}) - 1]}_{\text{\# periods before graduation}}, \quad (\text{III})$$

subject to $\mathcal{T}(X_1, w_{\min})$ being an integer.²⁹ Corollary 3 describes the solution:

Corollary 3 *Suppose the expert is required to pay a minimum wage w_{\min} . Up to an integer constraint for the novice's graduation date, the profit-maximizing knowledge gift X_1 satisfies*

$$\frac{f(X_1) - w_{\min}}{f(1) - w_{\min}} = \frac{1}{e}.$$

Consequently, the novice's graduation date (given by $\frac{1}{r_0} + 1$) is independent of w_{\min} . Moreover, the novice's output before graduation is uniformly increasing in w_{\min} .

This result tells us that the expert confronts the policy by raising the initial knowledge gift while holding constant the length of the apprenticeship. As a result, the policy shifts the output path upward and, at the same time, reduces its slope. The policy, therefore, is surplus enhancing. The policy also increases the novice's payoff, as he enjoys positive earnings before graduation.

Intuitively, the minimum wage is analogous to reducing the novice's productivity. To partially offset this loss, the expert raises the knowledge gift. Moreover, recall from the baseline model that the novice's graduation date is not affected by the details of f ;

²⁹The graduation date now satisfies $\frac{f(1) - w_{\min}}{f(X_1) - w_{\min}} = (1 + r)^{\mathcal{T}(X_1, w_{\min}) - 1}$.

therefore, being analogous to a reduction in f , the minimum wage does not affect the date of graduation.

Limit on apprenticeship duration. Suppose the planner requires, instead, that the expert ends the apprenticeship before some maximum number T_{\max} of periods. This policy, when binding, reduces the number of periods during which the expert can exchange knowledge for work, forcing the expert to speed up the exchange. The result is a higher knowledge gift and a uniformly higher level of output. As a result, the policy is also surplus enhancing.

It is worth noting that the above policies may backfire when the expert does not enjoy rents to begin with. For a simple example, suppose many experts compete for the novice and each one must pay a fixed cost F when contracting with him. In this case, experts offer the novice an apprenticeship just long enough for them to recover F . If experts are now required to pay a minimum wage, they must delay the novice's graduation, and slow down his training, simply to break even (assuming they are still able to do so). As a result, there is an efficiency loss.³⁰ Even worse, if experts are forced to reduce the duration of the apprenticeship below its original equilibrium level, it is impossible for them to recover F . As a result, they do not enter the market.

6 Extensions

We first study the case in which the novice has concave utility. We then consider some simple extensions of practical interest.

6.1 Consumption smoothing

Suppose the novice has concave utility. Let c_t denote the novice's period t consumption and let the novice's payoff from period t onward be $V_t = \sum_{\tau=t}^{\infty} \delta^{\tau-t} u(c_{\tau})$. For tractability, we assume that the novice has constant intertemporal elasticity of substitution (CIES), namely, $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ for some $\sigma > 0$.

Since the novice has no ability to borrow, we set $c_t = w_t$ if the novice works for the

³⁰The knowledge gift is now given by the largest value of X_1 such that $[f(X_1) - w_{\min}] \cdot [\mathcal{T}(X_1, w_{\min}) - 1] \geq F$. The reader can verify that as w_{\min} grows, X_1 falls, the graduation date grows, and output $f(X_t)$, which satisfies $f(X_t) - w_{\min} = (1+r)^{t-1} [f(X_1) - w_{\min}]$, uniformly falls during training.

expert, and $c_t = y_t$ if he walks away from the relationship.³¹ Other than the novice's incentive constraints, which are now $\frac{1}{1-\delta}u(y_t) \leq V_t$, the expert's problem is identical to that in the baseline model.

Proposition 3 *Suppose the novice has CIES utility with parameter σ . Then, every profit-maximizing contract is such that, after the initial gift and before graduation,*

$$\underbrace{\frac{1}{r} [u(y_{t+1}) - u(y_t)]}_{\text{value of knowledge transfer}} = \underbrace{u(y_t) - u(w_t)}_{\text{cost of working for expert}}$$

Moreover, $w_t = (1 - \delta)^{\frac{1}{\sigma}} Y_t$, with $Y_t = \left(\sum_{\tau=1}^t y_{\tau}^{\sigma}\right)^{\frac{1}{\sigma}}$.

Akin to the baseline model, the value of the period- t knowledge transfer equals the cost of working for the expert, now measured in utility terms. Moreover, the contract now prescribes an increasing wage path. Intuitively, this path is a compromise between delaying wages (which accelerates the knowledge transfer) and smoothing consumption (which helps the novice endure his training). Moreover, since knowledge is more valuable the earlier it is acquired, the consumption path is skewed toward the future. As an example, in the case of log utility ($\sigma = 1$), Y_t is the cumulative output produced up to period t . As a result, the wage path is both increasing and convex.

We derive the remaining details numerically (see Figure 3 below).³² The results are as follows:

1. As the novice's intertemporal elasticity of substitution $\frac{1}{\sigma}$ falls, he consumes a higher fraction of output during training. Consequently, the knowledge transfer slows down and the apprenticeship becomes longer – and is always longer than in the baseline model (Figure 3.A).
2. As in the baseline model, when players become more patient, training is slowed down and output is uniformly reduced. The reason is that, when δ grows, knowledge becomes more valuable, leading the expert to take longer to sell it (Figure 3.B).

³¹Since the wages prescribed by the optimal contract are non-decreasing over time (as we show below), the novice has no incentive to save either.

³²We obtain the optimal contract by searching (numerically) for the profit-maximizing value of T while imposing that, for each candidate T , the profile $(y_t, w_t)_{t=1}^{T-1}$ solves the $2(T-1)$ equations in the proposition (which have a unique solution).

3. Also as in the baseline model, imposing a minimum wage increases surplus by uniformly increasing output. The reason is that the expert partially counteracts the expense caused by the minimum wage by transferring additional knowledge – especially early in the relationship, when the minimum wage binds (Figure 3.C).

6.2 Other motives for altering the apprenticeship length

The baseline model can be readily extended along several other dimensions. Here we describe some examples. In all of them, the profit-maximizing contracts are apprenticeships with the properties in Proposition 1.³³ Therefore, they differ from one another only in the initial knowledge transfer X_1 and, potentially, a money transfer.

Training costs. When training costs are introduced, the apprenticeship may further slow down. Suppose, for example, that transferring additional knowledge $X' - X$, starting from stock X , costs the expert $\beta \frac{1}{r} [f(X') - f(X)]$ for some constant $\beta \in (0, 1)$ (namely, the cost is a constant fraction of the value of the additional knowledge).³⁴ Once this cost is considered, the profit-maximizing knowledge gift satisfies

$$\frac{f(X_1)}{f(1)} = \frac{1}{e^{1+B}},$$

where $B = \frac{\beta}{1-\beta} \left[\frac{r_0}{1-e^{-r_0}} \right]$ (recall that the constant in brackets, resulting from time being discrete, is approximately 1 when r_0 is small). Consequently, a higher β results in a smaller gift and a more distant graduation.

Novice's liquidity. Suppose that, at the beginning of the relationship, the novice has capital L , with $0 < L < \frac{1}{r} f(1)$.³⁵ In this case, the expert asks the novice to surrender all his capital up front and, in return, offers him an apprenticeship with the features in Proposition 1 – namely, a Pareto-efficient contract. Recall that, in such an apprenticeship, the novice enjoys (gross) rents $\frac{1}{r} f(X_1)$ (the present value of X_1). Therefore, the novice's participation constraint is $\frac{1}{r} f(X_1) \geq L$.

³³The proof of this claim, available upon request, is a straightforward extension of the proof of Proposition 1.

³⁴The expert's reduced-form objective, viewed from the standpoint of period 1, is now $(1 - \beta)f(X_1)[T(X_1) - 1] - \beta \frac{1+r}{r} f(X_1)$ (since, from Proposition 1, the per-period training costs, measured in period 1 dollars, are $\delta^{t-1} \beta \frac{1}{r} [f(X_{t+1}) - f(X_t)] = \beta f(X_1)$ for $t > 0$, and $\beta \frac{1+r}{r} f(X_1)$ for $t = 0$). The optimal X_1 maximizes this expression.

³⁵When $L \geq \frac{1}{r} f(1)$, the expert simply sells all knowledge up front.

Notice that when X_1 takes its baseline value ($\frac{f(X_1)}{f(1)} = \frac{1}{e}$), the novice enjoys rents $\frac{1}{r} \frac{f(1)}{e}$. Therefore, the profit-maximizing X_1 is obtained as follows: if the baseline rents exceed L , the expert simply pockets the novice’s capital; otherwise, the expert raises X_1 above its baseline level until $\frac{1}{r} f(X_1) = L$. As a result, the novice’s access to capital weakly accelerates the knowledge transfer. Moreover, it weakly *decreases* the novice’s payoff.

Externalities. In practice, the expert may experience externalities as the novice gains knowledge. For example, a partner in a law firm benefits when an associate becomes a more effective helper (e.g. Garicano, 2000). Alternatively, a northern firm loses profits when a southern firm learns from it and becomes a stronger competitor (as may potentially occur in the GM-SAIC case discussed in Section 2). For a simple formalization, suppose that, in addition to collecting revenues from the novice, the expert herself produces output $\gamma f(X_t)$ each period, with $\gamma > -1$ capturing the magnitude and sign of the externality.³⁶ In this case, the profit-maximizing knowledge gift satisfies

$$\frac{f(X_1)}{f(1)} = \frac{1}{e^{1-C}},$$

where $C = \frac{\gamma}{1+\gamma} \left[\frac{r_0}{1-e^{-r_0}} \right]$. As expected, this gift is increasing in γ : a larger externality accelerates training. Moreover, when γ is sufficiently large, training ends in one period; and when γ approaches -1 , training is stretched to infinity.

7 Conclusion

We have shown that a multi-period training arrangement, in which the novice is trained gradually over time, mitigates the classic problem of transferring general human capital. Such arrangement allows the novice to eventually acquire all knowledge while at the same time allowing the expert to profit.

Since the expert uses the promise of future knowledge transfers to retain the novice, she artificially prolongs the novice’s training. Likely instances of this inefficiency can be found in both traditional apprenticeships and in high-end professional partnerships. In the latter case, juniors appear, anecdotally, to spend years “paying their dues” to the

³⁶The expert’s reduced-form objective, viewed from the standpoint of period 1, is now $[\mathcal{T}(X_1) - 1] (1 + \gamma) f(X_1) + \gamma \frac{\delta^{\mathcal{T}(X_1)-1}}{1-\delta} f(1)$, which captures the impact of the externality during training (embedded in the first term) and during every period after that (second term). The optimal X_1 maximizes this expression.

firms' partners. During those years, juniors are involved in menial work, rather than being more quickly trained to perform high-value tasks.

We find that, as players become more patient, training takes longer to complete and output falls uniformly. The reason is that, as patience increases, knowledge becomes more valuable in the margin and the expert can keep the novice around with smaller transfers of additional knowledge. Thus, features that are traditionally considered to affect the discount factor, such as having more reliable partners, lead to slower transfers and a lower productivity while the novice is being trained.

Beyond apprenticeships, the model has implications for knowledge transfers in international alliances and joint ventures. When institutions are weak, contracts between companies exhibit the same lack of commitment found in training relationships between individuals. In this case, the "expert" partner may benefit from slowing down the knowledge transfer to ensure incentive compatibility while extracting maximum rents.

The present model can be used as a building block for other models of human-capital acquisition. In future work, we expect to study training hierarchies, where an expert can train a number of other agents, who in turn can train others. Also of interest is to study how firms may strengthen their relationships in order to facilitate knowledge sharing between them. An example is the use of cross-share holdings, typical for instance of Japanese Keiretsus.

Finally, the empirical evidence we have mentioned is by necessity anecdotal. Future empirical work is needed to study the extent to which experts artificially slow training down.

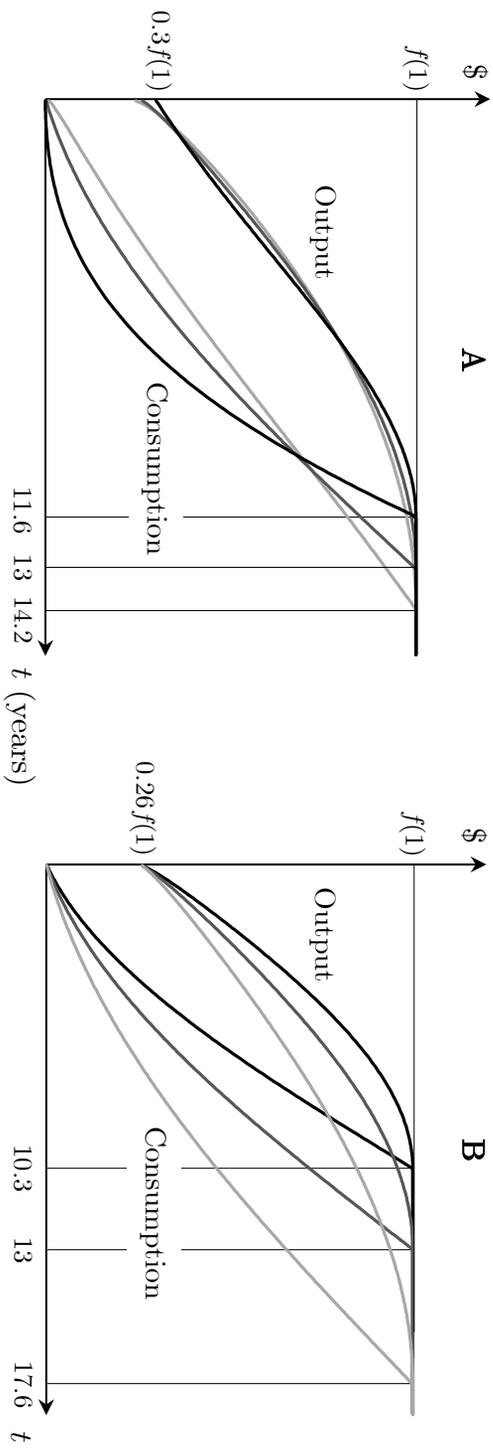


Figure 3. Apprenticeships with CES utility

A. $\sigma \in \{0.5, 1, 1.5\}$; $\delta = 0.9$
(lighter curves correspond to higher σ)

B. $\delta \in \{0.875, 0.9, 0.925\}$; $\sigma = 1$
(lighter curves correspond to higher δ)

C. $w_{\min} / f(1) \in \{0.05, 0.15, 0.25\}$; $\sigma = 1$; $\delta = 0.9$
(lighter curves correspond to lower w_{\min})

REFERENCES

1. Acemoglu, D. (1997). “Training and Innovation in an Imperfect Labour Market.” *Review of Economic Studies* 64(3): 445-464.
2. Acemoglu, D., and J.-S. Pischke (1998). “Why Do Firms Train? Theory and Evidence.” *Quarterly Journal of Economics* 113(1): 79-119.
3. Acemoglu, D., and J.-S. Pischke (1999a). “Beyond Becker: Training in Imperfect Labour Markets.” *Economic Journal* 109(453): 112-142.
4. Acemoglu, D., and J.-S. Pischke (1999b). “The structure of wages and investment in general training.” *Journal of Political Economy* 107(3): 539-572.
5. Albuquerque, R., and H. A. Hopenhayn (2004). “Optimal lending contracts and firm dynamics.” *Review of Economic Studies* 71(2): 285-315.
6. Anton, J. J., and D. A. Yao (2002). “The sale of ideas: Strategic disclosure, property rights, and contracting.” *Review of Economic Studies* 69(3): 513-531.
7. Baker, G., R. Gibbons, and K. Murphy (1994). “Subjective Performance Measures in Optimal Incentive Contracts.” *Quarterly Journal of Economics* 109(4): 1125-1156.
8. Bar-Isaac, H., and J.-J. Ganuza (2008). “Recruitment, training, and career concerns.” *Journal of Economics & Management Strategy* 17(4): 839-864.
9. Barron, D., R. Gibbons, R. Gil, and K. J. Murphy (2015). “Reel Authority: Relational Adaptation in the Movie Industry.” Northwestern University, mimeo.
10. Becker, G. S. (1994[1964]). *Human capital: A theoretical and empirical analysis with special reference to education*. University of Chicago Press.
11. Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications.” *Review of Economic Studies* 74(2): 345-90.
12. Bruss, F. Thomas (1984). “A Unified Approach to a Class of Best Choice Problems with an Unknown Number of Options.” *Annals of Probability* 12(3): 882-889.

13. Bull, C. (1987). “The Existence of Self-Enforcing Implicit Contracts.” *Quarterly Journal of Economics* 102(1): 147-160.
14. Bulow, J., and K. Rogoff (1989). “Sovereign Debt: Is to Forgive to Forget?” *American Economic Review* 79(1): 43-50.
15. Burdett, K., and E. Smith (1996). “Education and Matching Externalities.” In *Acquiring Skills*. Booth, A. L., and D. J. Snower, eds. Cambridge: Cambridge University Press: 65-80.
16. Chang, C., and Y. Wang (1996). “Human Capital Investment under Asymmetric Information: The Pigovian Conjecture Revisited.” *Journal of Labor Economics* 14(3): 505-519.
17. DeMarzo, P. M., and Y. Sannikov (2006). “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model.” *Journal of Finance* 61(6): 2681-724.
18. DeMarzo, P. M., and M. J. Fishman (2007). “Optimal Long-Term Financial Contracting.” *Review of Financial Studies* 20(6): 2079-2128.
19. Department for Business, Innovation and Skills (2013). “Apprentices’ Pay, Training and Working Hours.” *BIS*, UK Government.
20. Elbaum, Bernard (1989). “Why Apprenticeship Persisted in Britain But Not in the United States.” *Journal of Economic History* 49(2): 337-349.
21. Fox, K. (2015). “The kitchen apprentices: have knives, will travel.” *The Guardian*.
22. Garicano, L. (2000). “Hierarchies and the Organization of Knowledge in Production.” *Journal of Political Economy* 108(5): 874-904.
23. Gergaud, O., V. Smeets, and F. Warzynski (2014). “Learning by Cooking and Reputation Building: A French Recipe to Become a Top Chef.” KEDGE- Bordeaux Business School, mimeo.
24. Goloboy, J. L. (2008). *Industrial Revolution: People and Perspectives*. ABC-CLIO.
25. Halac, M. (2012). “Relational contracts and the value of relationships.” *American Economic Review* 102(2): 750-779.

26. Heckman, James (1971). "Three Essays on Household Labor Supply and the Demand For Market Goods." Chapter 1, Ph.D. Dissertation. Princeton University.
27. Hörner, J., and A. Skrzypacz (2010). "Selling Information." Cowles Foundation Discussion Paper No. 1743.
28. Ingalls, J. (2015) "Is Twitter the Architect's intern unofficial union? Exposing the reality behind unpaid internships across borders and industries." *Architect*.
29. Kahn, C., and G. Huberman (1988). "Two-Sided Uncertainty and "Up-or-Out" Contracts." *Journal of Labor Economics* 6(4): 423-444.
30. Kasperkevic, J. (2016). "Just saying no: US interns challenge employers over exploitation". *The Guardian*.
31. Katz, E., and A. Ziderman (1990). "Investment in General Training: The Role of Information and Labour Mobility." *Economic Journal* 100(403): 1147-1158.
32. Killingsworth, M. (1982) "Learning By Doing and Investment in Training: A Synthesis of Two Rival Models of the Life Cycle." *Review of Economic Studies* 49(2): 263-272.
33. Levin, J. (2003). "Relational Incentive Contracts." *American Economic Review* 93(3): 835-857.
34. Levin, J., and S. Tadelis (2005). "Profit sharing and the role of professional partnerships." *Quarterly Journal of Economics* 120(1): 131-171.
35. Li, J., and N. Matouschek (2013). "Managing conflicts in relational contracts." *American Economic Review* 103(6): 2328-2351.
36. Loewenstein, M. A., and J. R. Spletzer (1998). "Dividing the Costs and Returns to General Training." *Journal of Labor Economics* 16(1): 142-171.
37. MacLeod, W. B., and J. Malcomson (1989). "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57(2): 447-480.
38. MacLeod, W. B., and J. Malcomson (1998). "Motivation and Markets." *American Economic Review* 88(3): 388-411.

39. Maister, D. H. (1993). *Managing the Professional Service Firm*. Free Press, New York, NY.
40. Malcomson, J. M., J. W. Maw, B. McCormick (2003). "General training by firms, apprentice contracts, and public policy." *European Economic Review* 47(2): 197-227.
41. McVeigh, S. (2006). *Concert Life in London from Mozart to Haydn*. Cambridge University Press.
42. Merrit, C. (2013). "Firms exploiting young lawyers." *The Australian*.
43. Prendergast, C. (1993), "The Role of Promotion in Inducing Specific Human Capital Acquisition," *Quarterly Journal of Economics* 108(2): 523-534.
44. Rayo, L. (2007). "Relational incentives and moral hazard in teams." *Review of Economic Studies* 74(3): 937-963.
45. Richter, A., M. Dickmann, and M. Graubner (2008). "Patterns of human resource management in consulting firms." *Personnel Review* 37(2): 184-202.
46. Rosen, S. (1972). "Learning and Experience in the Labor Market." *Journal of Human Resources* 7(3): 326-342.
47. Shaw, K. (1989). "Lifecycle Labor Supply with Human Capital Accumulation." *International Economic Review* 30(2): 431-456.
48. Slater, D. (2009). "Another View: In Praise of Law Firm Layoffs." *The New York Times Dealbook*.
49. Smith, A. (1863). *An Inquiry into the nature and causes of the Wealth of Nations*. J. R. M'Culloch., Ad. and Ch. Black.
50. Stevens, M. (1994). "Labour Contracts and Efficiency in On-the-Job Training." *Economic Journal* 104(423): 408-419.
51. Thomas, J., and T. Worrall (1994). "Foreign direct investment and the risk of expropriation." *Review of Economic Studies* 61(1): 81-108.

52. Waldman, M. (1990). “Up-or-Out” Contracts: A Signaling Perspective.” *Journal of Labor Economics* 8(2): 230-250.
53. Weiss, Y. (1972). “On the Optimal Lifetime Pattern of Labour Supply.” *Economic Journal* 82(328): 1293-1315.

8 Appendix: Proofs

Proof of Lemma 1. For the first part of the Lemma (T is finite), suppose toward a contradiction that contract $\mathcal{C} = (y_t, w_t)_{t=0}^{\infty}$ solves problem (I) and yet training takes infinitely long – i.e. $y_t < y_{\text{sup}}$ for all t , where $y_{\text{sup}} = \lim_{t \rightarrow \infty} y_t$.

Now select a distant enough period k such that $y_k \geq \frac{1}{r} [y_{\text{sup}} - y_k]$ (which exists because y_k must converge to y_{sup} as k grows) and consider a new contract $\mathcal{C}' = (y'_t, w'_t)_{t=1}^{\infty}$ that is identical to \mathcal{C} except for the following variables: $y'_t = w'_t = y_{\text{sup}}$ for all $t > k$; and $w'_k = V_k(\mathcal{C}) - \frac{1}{r} y_{\text{sup}} \geq 0$.³⁷ Namely, in period k , the novice is asked to surrender net output $y_k - w'_k$ in exchange for all additional knowledge $y_{\text{sup}} - y_k$, with w'_k chosen such that the novice’s continuation value is unaffected ($V_k(\mathcal{C}') = w'_k + \frac{1}{r} y_{\text{sup}} = V_k(\mathcal{C})$).

We now show that \mathcal{C}' meets all constraints. Constraint (2) is met because, for all $t \leq k$, $\frac{1}{1-\delta} y'_t = \frac{1}{1-\delta} y_t \leq V_t(\mathcal{C}) = V_t(\mathcal{C}')$; and, for all $t > k$, $\frac{1}{1-\delta} y'_t = \frac{1}{1-\delta} y_{\text{sup}} = V_t(\mathcal{C}')$. Constraint (3) is met because, for all $t < k$, $w_t = w'_t$; and, for all $t \geq k$, $w'_t \geq 0$. Constraint (1) is met because, for all $t > k$, $\Pi_t(\mathcal{C}') = 0$; for period k ,

$$\Pi_k(\mathcal{C}') = y_k - w'_k = y_k + \frac{1}{r} y_{\text{sup}} - V_k(\mathcal{C}) > \sum_{\tau=k}^{\infty} \delta^{\tau-k} y_{\tau} - V_k(\mathcal{C}) = \Pi_k(\mathcal{C})$$

and therefore $\Pi_k(\mathcal{C}') > \Pi_k(\mathcal{C}) \geq 0$; and, for all $t < k$, $y'_t - w'_t = y_t - w_t$ and therefore

$$\Pi_t(\mathcal{C}') - \Pi_t(\mathcal{C}) = \delta^{k-t} [\Pi_k(\mathcal{C}') - \Pi_k(\mathcal{C})] > 0.$$

Finally, note from the above expression that \mathcal{C}' delivers a strictly higher profit than \mathcal{C} (namely $\Pi_0(\mathcal{C}') > \Pi_0(\mathcal{C})$), a contradiction.

For the second part of the Lemma ($X_T = 1$), note that if a given contract \mathcal{C} (satisfying all constraints) prescribes a finite graduation date T and prescribes $X_T < 1$, then the

³⁷Since $V_k(\mathcal{C}) \geq \frac{1}{1-\delta} y_k$ (from (2)) we have $w'_k \geq y_k - \frac{1}{r} [y_{\text{sup}} - y_k] \geq 0$.

expert's profits can be raised by scaling up all variables in \mathcal{C} by $\frac{1}{X_T}$ (while respecting all constraints). ■

Proof of Lemma 2. We say that a contract \mathcal{C} with graduation date T is a *delayed-reward* contract if $w_t = 0$ for all $t < T$; and is a *quasi-delayed-reward* contract if $w_{T-1} \in [0, f(1))$ and $w_t = 0$ for all $t < T - 1$. Let \mathcal{D} denote the set of delayed-reward contracts and let \mathcal{Q} denote the set of quasi-delayed-reward contracts (note that $\mathcal{Q} \supset \mathcal{D}$).

Step 1. Every optimal contract belongs to \mathcal{Q} . Let $\mathcal{C} = (y_t, w_t)_{t=0}^{\infty}$ be an arbitrary feasible contract (satisfying (1)-(3)) with a finite graduation date T and a full knowledge transfer. Let $\mathcal{C}' = (y'_t, w'_t)_{t=0}^{\infty}$ be the (unique) contract in \mathcal{Q} such that:

- (a) The novice's overall payoff is equal under \mathcal{C} and \mathcal{C}' , namely,

$$V_0(\mathcal{C}) = \sum_{t=0}^{\infty} \delta^t w_t = \sum_{t=0}^{\infty} \delta^t w'_t = V_0(\mathcal{C}').$$

- (b) Wages are $w'_t = f(1)$ for all $t \geq S$, $w'_{S-1} \in [0, f(1))$, and $w'_t = 0$ for all $t < S - 1$, where S is defined as the smallest period t such that $\frac{\delta^t}{1-\delta} f(1) \leq V_0(\mathcal{C})$.³⁸

- (c) All of the novice's incentive constraints (2) hold with equality: $V_t(\mathcal{C}') = \frac{1}{1-\delta} y'_t$.

Since $V_t(\mathcal{C}') = \frac{1}{1-\delta} f(1)$ for all $t \geq S$, and $V_t(\mathcal{C}') = \delta^{S-1-t} w'_{S-1} + \frac{\delta^{S-t}}{1-\delta} f(1)$ for all $t < S$, we have

$$y'_t = f(1) \text{ for all } t \geq S; \text{ and } y'_t < f(1) \text{ for all } t < S.$$

As a result, the novice's graduation date is S .

Notice that \mathcal{C}' is feasible.³⁹ In addition, \mathcal{C}' has the property that

$$V_t(\mathcal{C}) \leq V_t(\mathcal{C}') \text{ for all } t. \tag{A1}$$

For $t < S$, (A1) follows from the fact that $\sum_{\tau=0}^{t-1} \delta^\tau w'_\tau = 0$ and $\sum_{\tau=0}^{t-1} \delta^\tau w_\tau \geq 0$.⁴⁰ And,

³⁸Since $V_0(\mathcal{C}') = V_0(\mathcal{C})$, w'_{S-1} satisfies $\delta^{S-1} w'_{S-1} = V_0(\mathcal{C}) - \frac{\delta^S}{1-\delta} f(1)$, and therefore $w'_{S-1} \in [0, f(1))$.

³⁹Constraints (2) and (3) (as well as the monotonicity constraint for y'_t) hold by construction. Constraint (1) holds because $\Pi_t(\mathcal{C}') = \sum_{\tau=t}^{\infty} \delta^{\tau-t} y'_\tau - V_t(\mathcal{C}') \geq \frac{1}{1-\delta} y'_t - V_t(\mathcal{C}') = 0$.

⁴⁰Indeed, for all $t < S$,

$$\delta^t V_t(\mathcal{C}') = \sum_{\tau=t}^{\infty} \delta^\tau w'_\tau = \sum_{\tau=0}^{t-1} \delta^\tau w_\tau + \sum_{\tau=t}^{\infty} \delta^\tau w_\tau \geq \delta^t V_t(\mathcal{C}),$$

where the second equality follows from the fact that $V_0(\mathcal{C}') = V_0(\mathcal{C})$.

for $t \geq S$, (A1) follows from the fact that, owing to the expert's incentive constraint (1), $V_t(\mathcal{C}) \leq \frac{1}{1-\delta}f(1) = V_t(\mathcal{C}')$.

Notice, finally, that properties (c) and (A1) together imply that $y_t \leq y'_t$ for all t . Therefore, since \mathcal{C} and \mathcal{C}' deliver the same payoff for the novice, we have

$$\Pi_0(\mathcal{C}') - \Pi_0(\mathcal{C}) = \sum_{t=0}^{\infty} \delta^t [y'_t - y_t] \geq 0. \quad (\text{A2})$$

We now show that if \mathcal{C} is optimal, it must belong to \mathcal{Q} . Suppose instead that \mathcal{C} is optimal and yet $\mathcal{C} \notin \mathcal{Q}$. As a result, we must have $\sum_{t=0}^{t^*-1} \delta^t w_t > 0$ for some $t^* < T$.

Since S may be smaller than T , there are two cases to consider: $t^* < S$ and $t^* \geq S$. When $t^* < S$, we have $\sum_{t=0}^{t^*-1} \delta^t w'_t = 0$ and therefore

$$\delta^{t^*} V_{t^*}(\mathcal{C}') = \sum_{t=t^*}^{\infty} \delta^t w'_t = \sum_{t=0}^{t^*-1} \delta^t w_t + \sum_{t=t^*}^{\infty} \delta^t w_t > \delta^{t^*} V_{t^*}(\mathcal{C}).$$

When instead $t^* \geq S$, we have $V_{t^*}(\mathcal{C}') = \frac{1}{1-\delta}f(1) > V_{t^*}(\mathcal{C})$ (where the inequality follows from the fact that (1) requires that $\frac{1}{1-\delta}f(1) > V_t(\mathcal{C})$ for all $t < T$, and the fact that $t^* < T$). Either way, $V_{t^*}(\mathcal{C}') > V_{t^*}(\mathcal{C})$ and therefore $y'_{t^*} > y_{t^*}$. It follows from (A2) that $\Pi_0(\mathcal{C}') - \Pi_0(\mathcal{C}) \geq \delta^{t^*} [y'_{t^*} - y_{t^*}] > 0$, a contradiction.

Step 2. Every optimal contract belongs to \mathcal{D} (as claimed in the Lemma). Let \mathcal{C} be an optimal contract. Since \mathcal{C} belongs to \mathcal{Q} (from Step 1), there exists a period s such that $w_t = 0$ for all $t < s$, $w_s \in [0, f(1)]$, and $w_t = y_t = f(1)$ for all $t > s$. It follows that the expert's profits are

$$\Pi_0(\mathcal{C}) = \sum_{t=1}^s \delta^t y_t - \delta^s w_s.$$

Moreover, the novice's incentive constraints (2) up to period s are

$$V_t(\mathcal{C}) = \delta^{s-t} V_s(\mathcal{C}) \geq \frac{1}{1-\delta} y_t,$$

where $V_s(\mathcal{C}) = w_s + \frac{\delta}{1-\delta}f(1)$. Since \mathcal{C} is optimal, and $V_t(\mathcal{C})$ is nondecreasing, all such constraints must hold with equality.

Solving for y_t and rearranging terms, we obtain

$$\Pi_0(\mathcal{C}) = \delta^s w_s [(1-\delta)s - 1] + \text{constant},$$

where $constant = s\delta^{s+1}f(1)$. Recall that, by assumption, $(1 - \delta)n \neq 1$ for all $n \in \mathbb{N}$, and therefore $[(1 - \delta)s - 1] \neq 0$. Since the expert is free to vary w_s in the range $[0, f(1)]$, the optimality of \mathcal{C} requires that $w_s \in \{0, f(1)\}$. As a result, \mathcal{C} also belongs to \mathcal{D} . ■

Proof of Proposition 1. Let $\mathcal{C} = (y_t, w_t)_{t=0}^{\infty}$ be an optimal contract. From Lemmas 1 and 2, \mathcal{C} has a finite graduation date T , $y_t = f(1)$ for all $t \geq T$, and $w_t = 0$ for all $t < T$. In addition, as noted in the text, (1) and (2) jointly require that $w_t = f(1)$ for all $t \geq T$. It follows that the expert's profits are

$$\Pi_0(\mathcal{C}) = \sum_{t=0}^{\infty} \delta^t [y_t - w_t] = \sum_{t=1}^{T-1} \delta^t y_t.$$

Moreover, for all $t = 1, \dots, T - 1$, the novice's incentive constraints are

$$V_t(\mathcal{C}) = \frac{\delta^{T-t}}{1 - \delta} f(1) \geq \frac{1}{1 - \delta} y_t.$$

Since \mathcal{C} is optimal – and $V_t(\mathcal{C})$ is nondecreasing in t – all incentive constraints above must hold with equality. Namely, $y_t = \delta^{T-t} f(1)$. As a result, for all $t = 1, \dots, T - 1$, we have $\frac{y_{t+1}}{y_t} = 1 + r$ and therefore

$$\frac{1}{r} [y_{t+1} - y_t] = y_t,$$

as claimed in the proposition. Finally, that the knowledge gift X_1 is positive follows from the fact that $f(X_1) = y_1 = \delta^{T-1} f(1) > 0$. ■

Proof of Proposition 2. Consider problem (II). After setting $\mathcal{T}(X_1) - 1 = \frac{1}{r_0} \log \left[\frac{f(X_1)}{f(1)} \right]$, and multiplying the objective by the constant $\frac{r_0}{f(1)}$, the problem simplifies to

$$\begin{aligned} \max_{X_1 \geq 0} \quad & \frac{f(X_1)}{f(1)} \cdot \log \left[\frac{f(X_1)}{f(1)} \right] \\ \text{s.t.} \quad & \mathcal{T}(X_1) \in \mathbb{N}. \end{aligned}$$

Notice that the average logarithm $\frac{1}{z} \log z$ is uniquely maximized at $z = e$. As a result, when the integer constraint is ignored, the optimal X_1 satisfies $\frac{f(X_1)}{f(1)} = \frac{1}{e}$ and therefore $\mathcal{T}(X_1) - 1 = \frac{1}{r_0}$.

When the integer constraint is introduced, since the objective is single-peaked, the optimal X_1 satisfies either $\mathcal{T}(X_1) - 1 = \left\lfloor \frac{1}{r_0} \right\rfloor$ (the largest integer weakly smaller than $\frac{1}{r_0}$)

or $\mathcal{T}(X_1) - 1 = \left\lceil \frac{1}{r_0} \right\rceil$ (the smallest integer weakly larger than $\frac{1}{r_0}$), depending on which option delivers the highest profits.

Finally, the optimal gift satisfies $\frac{f(1)}{f(X_1)} = \alpha$, where α is either $(1+r)^{\lfloor \frac{1}{r_0} \rfloor}$ or $(1+r)^{\lceil \frac{1}{r_0} \rceil}$. Since $(1+r)^{\frac{1}{r_0}} = e$, the ratio $\frac{\alpha}{e}$, which is either $(1+r)^{\lfloor \frac{1}{r_0} \rfloor - \frac{1}{r_0}}$ or $(1+r)^{\lceil \frac{1}{r_0} \rceil - \frac{1}{r_0}}$, belongs to the interval $(\frac{1}{1+r}, 1+r)$. Therefore, as r converges to 0, $\frac{\alpha}{e}$ converges to 1. ■

Proof of Corollary 1. Every Pareto-efficient contract satisfies Lemmas 1-2 and Proposition 1. Other than the modification below, the proofs are identical to before. (Indeed, the proofs of Lemmas 1-2 and Proposition 1 show that any contract lacking the property in question is Pareto-dominated by a contract that has it – and that delivers a strictly higher Π_0 and a weakly higher V_0).

The modification is in the proof of Lemma 2, Step 2. There, one needs to replace the objective $\Pi_0(\mathcal{C}) = \sum_{t=1}^s \delta^t y_t - \delta^s w_s$ with the new objective

$$\lambda V_0(\mathcal{C}) + \Pi_0(\mathcal{C}) = \lambda \left[\delta^s w_s + \frac{\delta^{s+1}}{1-\delta} f(1) \right] + \sum_{t=1}^s \delta^t y_t - \delta^s w_s.$$

After solving for y_t , this objective simplifies to $\delta^s w_s [(1-\delta)s - (1-\lambda)]$ (plus a constant). And, by assumption, $[(1-\delta)s - (1-\lambda)] \neq 0$.

Once we restrict to contracts satisfying Proposition 1, the planner's objective (expressed in date-0 dollars) is

$$\underbrace{\lambda \frac{1}{r} f(X_1)}_{V_0} + \underbrace{\delta f(X_1) [\mathcal{T}(X_1) - 1]}_{\Pi_0},$$

where $\mathcal{T}(X_1) - 1 = \frac{1}{r_0} \log \left[\frac{f(1)}{f(X_1)} \right]$. When the integer constraint is ignored, the optimal X_1 satisfies $\frac{f(X_1)}{f(1)} = \min \left\{ \frac{1}{e^{1-\lambda A}}, 1 \right\}$ (where $A = \frac{r_0}{1-e^{-r_0}}$). Therefore, $\mathcal{T}(X_1) - 1 = D$, where $D = \max \left\{ \frac{1-\lambda A}{r_0}, 0 \right\}$. When the integer constraint is introduced, since the objective is single-peaked, $\mathcal{T}(X_1) - 1$ is either $\lfloor D \rfloor$ or $\lceil D \rceil$ and X_1 satisfies $\frac{f(X_1)}{f(1)} = \frac{1}{\alpha}$, where α is either $(1+r)^{\lfloor D \rfloor}$ or $(1+r)^{\lceil D \rceil}$. (Since $\min \left\{ \frac{1}{e^{1-\lambda A}}, 1 \right\} = (1+r)^{-D}$, and $\frac{1}{\alpha}$ converges to $(1+r)^{-D}$ when r converges to zero, it follows that $\frac{f(X_1)}{f(1)}$ converges to $\min \left\{ \frac{1}{e^{1-\lambda A}}, 1 \right\}$ when r converges to zero.) ■

Proof of Corollary 2. Lemmas 1 and 2 remain valid under the minimum wage policy, with the modification that, before graduation, the wage earned over and above the minimum wage $w_t - w_{\min}$ takes the place formerly occupied by the wage w_t . After

this modification, the proof of Lemma 1 remains valid, with period k chosen so that $y_k - w_{\min} \geq \frac{1}{r} [y_{\sup} - y_k]$. And the proof of Lemma 2 remains valid, with a *delayed-reward* contract now requiring that $w_t = w_{\min}$ for all $t < T$; and a *quasi-delayed-reward* contract now requiring that $w_{T-1} \in [w_{\min}, f(1))$ and $w_t = w_{\min}$ for all $t < T - 1$.

Finally, the proof of the present Corollary is identical to the proof of Proposition 1, with the exception that profits are now $\sum_{t=1}^{T-1} \delta^t [y_t - w_{\min}]$ and the novice's incentive constraints for $t = 1, \dots, T - 1$ are now

$$\frac{1 - \delta^{T-t}}{1 - \delta} w_{\min} + \frac{\delta^{T-t}}{1 - \delta} f(1) \geq \frac{1}{1 - \delta} y_t.$$

Once these constraints bind, we obtain $y_t - w_{\min} = \delta^{T-t} [f(1) - w_{\min}]$ and therefore $\frac{1}{r} [y_{t+1} - y_t] = y_t - w_{\min}$. ■

Proof of Corollary 3. Problem (III) is identical to problem (II) in the baseline model, with $f(X_1) - w_{\min}$ in the place of $f(X_1)$, $f(1) - w_{\min}$ in the place of $f(1)$, and $\mathcal{T}(X_1, w_{\min})$ in the place of $\mathcal{T}(X_1)$.

As a result, when the integer constraint is ignored, the solution satisfies $\frac{f(1) - w_{\min}}{f(X_1) - w_{\min}} = e$ and $\mathcal{T}(X_1, w_{\min}) - 1 = \frac{1}{r_0}$. And when the integer constraint is introduced, the solution satisfies $\frac{f(1) - w_{\min}}{f(X_1) - w_{\min}} = \alpha$, where α is either $(1 + r)^{\lfloor \frac{1}{r_0} \rfloor}$ or $(1 + r)^{\lceil \frac{1}{r_0} \rceil}$, and $\mathcal{T}(X_1, w_{\min}) - 1$ is either $\lfloor \frac{1}{r_0} \rfloor$ or $\lceil \frac{1}{r_0} \rceil$. (As before, when r converges to 0, α converges to e .)

Moreover, since net output $f(X_t) - w_{\min}$ grows at rate r during training, we obtain

$$\frac{f(X_t) - w_{\min}}{f(1) - w_{\min}} = \frac{1}{\alpha} (1 + r)^{t-1} \text{ for all } t \leq \mathcal{T}(X_1, w_{\min}),$$

which implies that $f(X_t)$ is increasing in w_{\min} for all $t < \mathcal{T}(X_1, w_{\min})$. ■

Proof of Proposition 3. For any given graduation T , the expert's problem is

$$\begin{aligned} & \max_{(y_t, w_t)_{t=1}^{T-1}} \sum_{t=1}^{T-1} \delta^{t-1} [y_t - w_t] \\ & \text{s.t. } V_t \geq \frac{1}{1 - \delta} u(y_t), \end{aligned}$$

$$w_t \geq 0, y_t \in [0, f(1)] \text{ and nondecreasing,}$$

where $V_t = \sum_{\tau=t}^{T-1} \delta^{\tau-t} u(w_\tau) + \frac{\delta^{T-t}}{1-\delta} u(f(1))$.⁴¹ Let λ_t denote the Lagrange multiplier for

⁴¹That there exists an optimal contract follows from noting that, without loss, we can bound w_t above

the period t incentive constraint.

Now suppose T is the optimal graduation date. We begin by ignoring the monotonicity constraint for y_t . In the resulting relaxed problem, all incentive constraints bind (otherwise profits can be raised by raising y_t). Moreover, the Inada condition $u'(0) = \infty$ implies that $w_t > 0$ (and therefore $y_t > 0$), and since by hypothesis the novice's graduation is at time T , we also have $y_t < 1$ for all $t < T$. As a result, the solution to the relaxed problem is interior and therefore characterized by the first-order conditions $u'(w_t) \sum_{\tau=1}^t \lambda_\tau \delta^{t-\tau} = \delta^{t-1}$ and $\lambda_t = \frac{(1-\delta)\delta^{t-1}}{u'(y_t)}$, together with $\frac{1}{1-\delta}u(y_t) = V_t$.⁴²

By combining the first-order conditions and rearranging terms, we obtain

$$w_t = (1 - \delta)^{\frac{1}{\sigma}} Y_t,$$

where $Y_t = (\sum_{\tau=1}^t y_\tau^\sigma)^{\frac{1}{\sigma}}$. (Given that Y_t is increasing in t , so are w_t and V_t .) Moreover, since we can write $u(w_t) + \delta V_{t+1} = V_t$, and $\frac{1}{1-\delta}u(y_t) = V_t$, we obtain $u(w_t) + \frac{1}{r}u(y_{t+1}) = \frac{1}{1-\delta}u(y_t)$. After rearranging terms, this equality is

$$\frac{1}{r} [u(y_{t+1}) - u(y_t)] = u(y_t) - u(w_t).$$

Finally, focusing on the above relaxed problem is appropriate (namely, the monotonicity constraint for y_t is redundant) because the solution to the relaxed problem satisfies $\frac{1}{1-\delta}u(y_t) = V_t$, with V_t increasing. ■

by some finite number W . As a result, for any arbitrary T , the above problem has a solution: it consists of maximizing a continuous function over a compact set. Moreover, that an optimal T exists follows from noting that every contract in which training never ends is dominated by one in which training ends in finite time (which follows from an argument analogous to that in Lemma 1).

⁴²That these conditions describe a maximum follows from the fact that there is a unique contract $(y_t, w_t)_{t=1}^{T-1}$ that satisfies them (and the fact that the expert's problem has an interior solution).