DISCUSSION PAPER SERIES

No. 11138

RELATIONAL KNOWLEDGE TRANSFERS

Luis Garicano and Luis Rayo

INDUSTRIAL ORGANIZATION and LABOUR ECONOMICS

ISSN 0265-8003

RELATIONAL KNOWLEDGE TRANSFERS

Luis Garicano and Luis Rayo

Discussion Paper No. 11138
March 2016
Submitted 18 February 2016

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801
www.cepr.org

RELATIONAL KNOWLEDGE TRANSFERS

Abstract

An expert with general knowledge trains a cash-constrained novice. Faster training increases the novice's productivity and his ability to compensate the expert; it also shrinks the stock of knowledge yet to be transferred, reducing the expert's ability to retain the novice. The profit-maximizing agreement is a multi-period apprenticeship in which knowledge is transferred gradually over time. The expert adopts a " $\frac{1}{e}$ rule" whereby, at the beginning of the relationship, the novice is trained just enough to produce a fraction $\frac{1}{e}$ of the efficient output. This rule causes inefficiently lengthy relationships that grow longer the more patient the players. We discuss policy interventions.

JEL Classification: D86, J24, L24 and M53
Keywords: general human capital, international joint ventures and relational contracts

Luis Garicano l.garicano@lse.ac.uk
London School of Economics and CEPR

Luis Rayo Luis.Rayo@utah.edu
University of Utah and CEPR

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

Relational Knowledge Transfers*

Luis Garicano
LSE and CEPR
L.Garicano@lse.ac.uk

Luis Rayo
University of Utah and CEPR
Luis.Rayo@utah.edu

February 9, 2016

Abstract

An expert with general knowledge trains a cash-constrained novice. Faster training increases the novice's productivity and his ability to compensate the expert; it also shrinks the stock of knowledge yet to be transferred, reducing the expert's ability to retain the novice. The profit-maximizing agreement is a multi-period apprenticeship in which knowledge is transferred gradually over time. The expert adopts a "$\frac{1}{e}$ rule" whereby, at the beginning of the relationship, the novice is trained just enough to produce a fraction $\frac{1}{e}$ of the efficient output. 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, seminar participants at the 2013 NBER Working Group in Organizational Economics, the 2015 European Summer Symposium in Economic Theory, and the 2015 CEPR IMO Workshop, and workshop particpants at Amsterdam, CEMFI, Chicago, Kellogg, LBS, Milan, and USC for valuable suggestions.

1 Introduction

As noted by Becker (1994[1964]), training involves a transfer of human capital; thus, unlike other market transactions, it does not create an asset that can be used as collateral. Moreover, the novice (apprentice) often does not have the means to pay the expert (master), up front, for the knowledge he wishes to acquire. In this case, the rate of knowledge transfer is constrained by the need to ensure that, over time, the novice has both the means and the incentives to compensate the expert for his training. Unlike in the first-best allocation, where knowledge is transferred as fast as technologically feasible, the rate of knowledge transfer is the result of a trade-off: the faster it is, the sooner output can be generated; but if it is too fast, the novice will leave without fully compensating the expert. For example, in professional partnerships, novices, usually called "associates," are rewarded for their work not only through their current salary, but also through the promise of future learning. Such an arrangement allows partners to profit while associates are being trained. Similarly, in international joint ventures involving technological transfers between a "northern" firm (the expert) and a "southern" firm (the novice), the latter can potentially ignore formal agreements and establish its own operations. As a result, the rate of knowledge transfer is constrained by the need to ensure a high enough continuation value for both partners.

The object of this paper is to study, in an environment with limited contractability, the form and duration of the optimal dynamic relationship between expert and novice. The goal, in particular, is to explore the speed at which knowledge is transferred, together with its distributional and efficiency implications.

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), in contrast, has no knowledge, and therefore no ability to produce output; he also has no cash, and therefore no ability to purchase knowledge from the expert. By transferring knowledge, the expert can raise the novice's productivity and, with it, his ability to transfer capital back to the expert. 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 walk away from the expert, players rely on a self-enforcing agreement in which knowledge is transferred gradually over time. Such an agreement is capable of sustaining knowledge transfers from the expert by providing an expectation of future payments from the novice; these payments are sustained, in turn, via the expectation of further knowledge transfers from the expert.

As a benchmark, we first consider relational contracts in which, by assumption, all knowledge transfers must end after two periods. The optimal arrangement is forward-looking: the initial knowledge transfer is used as a knowledge "gift" (rather than a loan) that allows the novice to produce output before the overall knowledge transfer is over; and the final knowledge transfer is used as an incentive for the novice to surrender such output. In this two-period arrangement, the expert's only source of revenue is the output produced in the first period, using the knowledge gift, which means that the expert is under pressure to make this gift as large as possible. Because of this pressure, for realistic values of the players' discount factor, the arrangement is close to first best (and converges to first best when the discount factor converges to one).

We next turn to the general case in which the expert selects a multi-period arrangement of any desired length. The optimal contract is an "apprenticeship" arrangement in which, until the knowledge transfer is complete, players trade labor for training. The structure of this contract generalizes the two-period arrangement: after an initial knowledge gift, the novice is asked to work for the expert in exchange for additional knowledge that has a (discounted) value just high enough to compensate the novice for this work. Equivalently, this arrangement can be structured as an initial knowledge gift followed by a series of sales contracts in which, until all knowledge has been sold, the novice devotes 100% of his output to gradually purchase new knowledge. By delaying consumption until training is complete, the expert more quickly extract rents from the novice.

The overall length of this apprenticeship is controlled by the size of the initial knowledge gift, with a smaller gift leading to a more distant graduation. When selecting the optimal gift, the expert faces the following trade-off: a larger gift raises the novice's productivity, allowing the expert to extract higher revenues during each period of the apprenticeship; but a larger gift also reduces the amount of knowledge yet to be transferred, reducing the number of apprenticeship periods that the novice is willing to withstand. Because of the multi-period nature of the agreement, the expert is in much less of a rush, relative to the two-period setting, to raise the novice's productivity early in the relationship.

We find that, no matter how patient players are, and regardless of the details of the output technology, the optimal 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" leads to long apprenticeships in which significant output is wasted. For example, when the annual interest rate is 10% (resp. 5%), training takes 9.5 years (resp. 19.5 years) to complete. This rule also implies that, in the absence of other factors affecting the relationship, novices with different skill levels, and novices working in different professions, take equally long to train.

The optimal apprenticeship is longer, and knowledge is transferred more slowly, the more patient the players (thus, unlike in the two-period contract, more surplus is wasted when players are more patient). 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 a smaller amount of new knowledge; a fact that the expert exploits by (inefficiently) slowing down the training speed and keeping the novice in her employment for longer.

Motivated by the expert's preference for artificially lengthy apprenticeships, as well as by general commentary on real-world masters "exploiting" their apprentices via contracts with slow training and low consumption (as discussed in Section 2), we consider two policy experiments. First, we force the expert to pay the novice a minimum wage during training. We show that while this policy leaves the contract length unaffected (an implication of the expert's $\frac{1}{e}$ rule), it raises surplus by uniformly accelerating the novice's' training and, with it, his output. The reason for this efficiency gain is that raising the novice's productivity allows the expert to partially offset the expense caused by the minimum wage. 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 end our policy discussion by illustrating how both of these policies may backfire when the expert does not enjoy rents to begin with.

Finally, we consider several extensions of the model. First, the case in which the novice has concave utility, giving him a preference for smooth consumption. In this case, the expert grants the novice an increasing consumption path that is initially close to zero; representing 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). Second, we consider brief straightforward extensions of practical interest: the expert facing training costs, the novice arriving with capital, and the case in which training causes externalities on the expert (such as an expert partner in a law firm benefitting when a novice associate becomes a more effective problem solver, e.g. Garicano 2000, or an expert firm losing profits when training a novice firm who then becomes a more effective competitor). All of these modifications alter the optimal contract exclusively via the ratio of knowledge gifted to sold. Finally, we show that the set of Pareto-efficient contracts is a family of apprenticeships in which training is accelerated as the novice's Pareto weight grows. 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 will in principle not pay for general human capital acquisition of their workers – if they were to do so, they would not recoup their investment, as workers can always move to another firm. A large literature has aimed to explain, under these circumstances, firms' incentives to train their workers by relying on market imperfections. 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.[1]

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. 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 similar productivity in all jobs.

[1] Alternatively, in the 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.

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.[2] (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, in which, 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).[3] This literature focuses on eliciting a costly, productive 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.[4]

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 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.[5] In our model, in

[2] In 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.

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

[4] Owing to the liquidity constraint faced by our novice, the "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 money transfers 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).

[5] Anton 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

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 sequence of spot sales contracts, and the reason they remain in the relationship is to benefit from future sales, rather than to avoid punishments.[6] 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 stylized facts in expert-novice relationships. Sections 3 and 4 present the baseline model and derive the optimal contracts. Section 5 considers policy experiments and Section 6 considers various extensions of the baseline model. Section 7 concludes. All proofs are in the Appendix.

purchase the remaining knowledge in a one-shot transaction.

[6]In 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" (Department for Business, Innovation and Skills, 2013).[7]

An important instance of knowledge transfers occurs in professional service firms. These firms provide a wide range of general skills to junior consultants, usually called associates (see, for example, Maister, 1993, and Richter et al., 2008). Once again, training appears to be slowed down, artificially, while associates "pay their dues." In the process,

[7] "Apprentices' Pay, Training and Working Hours." BIS, UK Government, 2013. (https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/49987/bis-13-532-follow-up-research-apprentices-pay-training-and-working-hours.pdf.)

rents are extracted from their work in exchange for the promise of further training and eventual promotion.[8]

At law firms, for example, the training of associates seems, anecdotally, to proceed at a glacial pace; as a result, associates are forced into menial tasks. Numerous blog posts and articles are dedicated to describing this feature. For instance, 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" (New York Times, 2009).[9] Similarly, 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."[10] Perhaps not surprisingly, the response of the American Bar Association (ABA) is to tough it out. The ABA admonishes, in its advice to young lawyers: "No task is too menial that you can't learn from it."[11]

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 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. In this case, the relationship only lasts for as long as the parties consider it in their common interest to continue their "knowledge-for-cash" exchange.

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.)[12] For Danone, the venture was a way to profit from the growing Chinese market; for Wahaha, it was a means to learn

[8] Levin 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.

[9] "Another View: In Praise of Law Firm Layoffs." NYT. Dealbook. July 1, 2009. By Dan Slater.

[10] http://www.theaustralian.com.au/business/legal-affairs/firms-exploiting-young-lawyers-says-chief-justice-marilyn-warren/story-e6frg97x-1226717910085.

[11] http://apps.americanbar.org/litigation/committees/trialpractice/younglawyers.html.

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

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, its 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).[13] 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.[14] 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 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).[15] 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

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

[14] 2000 U.S. Dist. LEXIS 18449.

[15] "Balancing the Give and Take in GM's Chinese Partnership." WSJ. August 19, 2012. By Sharon Terlep.

the relationship, GM has not yet transferred some of its most valuable 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 the American car maker down the road."

3 Baseline model

There are two risk-neutral players: an expert E (she) and a novice N (he). Players interact over infinite periods $t = 1, 2, ...$ 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 fraction of knowledge transferred during period t and let X_t denote the novice's total knowledge, inclusive of x_t, during period t:

$$X_t = x_t + X_{t-1},$$

with $X_0 = 0$.

During period t, the novice produces output $f(X_t) \in \mathbb{R}_+$, with f continuous and increasing, and $f(0) = 0$. One interpretation is that the novice's output originates from a variety of tasks, with less valuable tasks requiring less knowledge. In this case, as the novice acquires knowledge, he efficiently spends less time on menial tasks and more on advanced ones.[16] 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.[17] Let $P(X', X) = \frac{1}{1-\delta}[f(X') - f(X)]$ denote the marginal value of a package of knowledge containing $X' - X$ units, in present discounted terms, starting from an initial knowledge stock X.

[16] For example, a cook may either engage in the menial task of chopping vegetables or in the advanced task of making sauces (e.g. Gergaud et al., 2014). Similarly, a consultant may either perform lengthy spreadsheet calculations or simply advise others on what objects to calculate (e.g. Garicano, 2000). In either case, only after appropriate training is the worker successful at the advanced task.

[17] This assumption rules out both direct training costs and externalities experienced by the expert as the novice becomes more knowledgeable. In Section 6.2, we extend the model to illustrate the impact of such costs.

The novice has a bank account that pays interest r between periods. Let $(1+r)B_{t-1}$ and B_t denote, respectively, the balance of this account at the beginning and end of period t. The novice cannot borrow from his bank ($B_t \geq 0$) and, for now, has no initial capital ($B_0 = 0$). The expert, in contrast, has deep pockets.

Each period t has three stages. First, there is a *knowledge-transfer stage* in which players swap knowledge for cash: the expert makes a knowledge transfer x_t and the novice makes a money transfer $m_t \in \mathbb{R}$ in return. Players rely on a spot contract for this swap. Each player $i = E, N$ simultaneously proposes a pair (x_t^i, m_t^i) subject to two feasibility constraints: (1) x_t^i is no larger than the expert's remaining knowledge stock $1 - X_{t-1}$; (2) m_t^i is no larger than the novice's current savings $(1+r)B_{t-1}$. If $(x_t^E, m_t^E) = (x_t^N, m_t^N)$ an agreement is reached and the corresponding transfers take place: $x_t = x_t^E$ and $m_t = m_t^E$. Otherwise, no agreement is reached and $x_t = m_t = 0$.

Second, there is a *production stage* in which the novice either works on her own ($d_t = 0$) or works for the expert ($d_t = 1$), and the expert pays a wage $w_t \in \mathbb{R}$ to the novice. Players rely on a spot employment contract. Each player simultaneously proposes a pair (d_t^i, w_t^i) (with $d_t^i \in \{0, 1\}$) subject to $-w_t^i$ not exceeding the novice's current bank balance ($-w_t^i \leq (1+r)B_{t-1} - m_t$). If $(d_t^E, w_t^E) = (d_t^N, w_t^N)$ an agreement is reached: $d_t = d_t^E$ and $w_t = w_t^E$. Otherwise, $d_t = w_t = 0$. Since knowledge is general, output $f(X_t)$ is independent of the employment decision.

Third, there is a *consumption stage* in which the novice splits his stage-2 earnings $w_t + (1 - d_t) f(X_t)$ and remaining balance $(1+r)B_{t-1} - m_t$ between consumption $c_t \in \mathbb{R}_+$ and savings. His resulting end-of-period balance is

$$B_t = w_t + (1 - d_t) f(X_t) + (1+r)B_{t-1} - m_t - c_t.$$

All variables are observed by both players. In addition, the only formal (court-enforced) contracts that can be written are the one-period spot contracts described above. All other agreements must be self-enforced.[18]

From the standpoint of the beginning of period t, the continuation payoffs for expert

[18] If the expert had commitment power (e.g. owing to external reputation concerns), the assumption that spot contracts are court-enforced would not be required.

and novice are, respectively,

$$\begin{aligned}\Pi_t &= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[m_\tau + d_\tau f(X_\tau) - w_\tau \right] \text{ and} \\ V_t &= \sum_{\tau=t}^{\infty} \delta^{\tau-t} c_\tau = (1+r) B_{t-1} + \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[-m_\tau + (1-d_\tau) f(X_\tau) + w_\tau \right].\end{aligned}$$

Since the players' combined surplus is $\Pi_1 + V_1 = \sum_{t=1}^{\infty} \delta^{t-1} f(X_t)$, the first-best allocation calls for a full knowledge transfer in the first period.

A self-enforcing *relational contract* (or more briefly, a *contract*) prescribes, on the path of play, a vector $\langle X_t, m_t, d_t, w_t, c_t \rangle$ for every period t; and, upon any deviation from that prescription, it prescribes a suspension of all further transactions between the players.[19] For notational simplicity we denote a contract by its prescribed actions on the path of play: $\mathcal{C} = \langle X_t, m_t, d_t, w_t, c_t \rangle_{t=1}^{\infty}$. Let $\Pi_t(\mathcal{C})$, $V_t(\mathcal{C})$, and $B_t(\mathcal{C})$ denote, respectively, the equilibrium values of Π_t, V_t, and B_t. Also let Ω denote the set of contracts $\mathcal{C}$ that satisfy the basic requirements that $x_t \leq 1 - X_{t-1}$, $B_t(\mathcal{C}) \leq 0$, and $c_t \geq 0$ for all t. We say that a contract $\mathcal{C}$ is feasible if it belongs to Ω and, in addition, it constitutes a subgame-perfect equilibrium of the dynamic game.

Feasibility of $\mathcal{C}$ requires that six constraints are met for every period t– three constraints for each of the first two stages.[20] As we show momentarily, however, all stage-2 constraints are redundant.

Stage-1 constraints

From the standpoint of the beginning of stage 1, both players must weakly prefer to follow their prescribed actions over reneging on the agreement:

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

$$V_t(\mathcal{C}) \;\geq\; \underbrace{\frac{1}{1-\delta} f(X_{t-1})}_{\text{value of knowledge}} \;+\; \underbrace{(1+r) B_{t-1}(\mathcal{C})}_{\text{current balance}} \;\text{ for all } t. \tag{2}$$

[19] These trigger punishments are in principle not immune to renegotiation. As we shall see, however, the optimal relational contract is itself immune to renegotiation.

[20] In principle we may also require a constraint for stage 3, in which the novice makes his consumption/savings decision. This constraint, however, is redundant because, because being risk-neutral and earning interest r, the novice can always replicate an "over-consumption" deviation in this stage with a deviation in stage 1 of the following period in which he walks away with his savings.

The novice's constraints require that he weakly prefers his continuation payoff over walking away with his current stock of knowledge X_{t-1} (worth $\frac{1}{1-\delta}f(X_{t-1})$ in present value) together with his current cash balance. We call these two constraints the players' incentive constraints.

In addition, the novice must have sufficient liquidity to afford the transfers m_t:

$$m_t \leq (1+r)\, B_{t-1}(\mathcal{C}) \text{ for all } t. \tag{3}$$

Stage-2 constraints

From the standpoint of the beginning of stage 2, both players must again be deterred from reneging on the agreement:

$$\underbrace{d_t f(X_t) - w_t}_{\text{stage-2 revenues}} \ + \ \delta \Pi_{t+1}(\mathcal{C}) \geq 0 \text{ for all } t, \tag{4}$$

$$V_t(\mathcal{C}) \geq \underbrace{\frac{1}{1-\delta} f(X_{t-1})}_{\text{value of knowledge}} \ + \ \underbrace{(1+r)\, B_{t-1}(\mathcal{C}) - m_t}_{\text{current balance}} \ \text{ for all } t. \tag{5}$$

The novice must also be able to afford the transfers $-w_t$:

$$-w_t \leq (1+r)\, B_{t-1}(\mathcal{C}) - m_t \text{ for all } t. \tag{6}$$

Lemma 1 tells us that, without loss, the residual claim of output can be granted to the novice and wages can be set to zero. As a result, all stage-2 constraints are redundant:

Lemma 1 *Let $\mathcal{C} = \langle X_t, m_t, d_t, w_t, c_t \rangle_{t=1}^{\infty}$ be a feasible contract. There exists a feasible contract $\mathcal{C}' = \langle X_t, m_t', d_t', w_t', c_t \rangle_{t=1}^{\infty}$ that prescribes the same knowledge transfers and equilibrium payoffs Π_1, V_1 as contract $\mathcal{C}$ and under which, in every period, the novice is the residual claimant of output ($d_t' = 0$) and receives zero wages ($w_t' = 0$). Moreover, under contract $\mathcal{C}'$ all stage-2 constraints ((4), (5), and (6)) are redundant.*

Intuitively, since knowledge is general, it is immaterial whether the residual claimant of output is the novice or the expert. In the former case, the novice's most tempting deviation consists in walking away from the relationship during the knowledge-transfer phase of a given period t, before making any money transfer to the expert. In the latter case, the novice's most tempting deviation is identical except for the fact that it is initiated during the production phase of period $t-1$, with the novice producing output $f(X_{t-1})$ on his own rather than for the expert. In present-discounted terms, both deviations deliver the same payoff.

In what follows, we set $d_t = w_t = 0$ for all t and refer to a relational contract simply as $\mathcal{C} = \langle X_t, m_t, c_t \rangle_{t=1}^{\infty}$. We define a "training period" of $\mathcal{C}$ as a period in which the novice receives knowledge ($x_t > 0$) and the "training phase" of $\mathcal{C}$ as the set of all its training periods. We also let $X_{\sup}(\mathcal{C}) = \lim_{t\to\infty} X_t$ denote the total knowledge transferred under $\mathcal{C}$.

4 Optimal contracts

An optimal contract solves the expert's profit-maximization problem:

$$\max_{\mathcal{C}\in\Omega} \ \Pi_1(\mathcal{C}) \qquad \text{(I)}$$
$$\text{s.t. (1), (2), and (3).}$$

To build intuition, we begin with a simple benchmark in which, by assumption, all transactions end after two periods. We then turn to contracts of optimal length. Throughout, 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 contracts in which, by assumption, knowledge and money transfers are restricted to occur in the first two periods of the game (i.e. $x_t = m_t = 0$ for all $t > 2$). The expert's problem is

$$\max_{\mathcal{C}\in\Omega} \ m_1 + \delta m_2$$
$$\text{s.t. } m_1 \leq 0,$$
$$m_2 \leq \min\left\{ \underbrace{P(X_2, X_1)}_{\text{Value of } X_2 - X_1}, \underbrace{(1+r)B_1(\mathcal{C})}_{\text{Available cash}} \right\}, \qquad (IC+L)$$

Constraint $(IC+L)$ tells us that m_2 must be weakly below both the marginal value of the knowledge transferred in the second period and the novice's cash balance at beginning of that period. A higher initial knowledge level X_1 increases the novice's subsequent liquidly, but also reduces his willingness to pay for the remaining knowledge $X_2 - X_1$, as there is less of it. (In this two-period benchmark, the expert's own incentive constraints (1) are redundant.) The optimal contract belongs to a simple class:[21]

Remark 1 *Any optimal two-period contract $\mathcal{C} = (X_t, m_t, c_t)_{t=1,2}$ consists of a knowledge gift $X_1 > 0$ and a money loan $m_1 \leq 0$ in period 1 followed by a spot contract in period 2 in which the novice devotes the full balance of his bank account to purchase additional knowledge $X_2 - X_1 > 0$ at a price equal to marginal value $P(X_2, X_1)$. Namely,*

$$m_2 = (1+r)B_1(\mathcal{C}) = P(X_2, X_1).$$

When selecting the total knowledge transfer X_2 and the first-period consumption level, the expert's only goal is to maximize the magnitude of the second-period knowledge sale. For this reason, she opts for a full knowledge transfer (which maximizes the amount of knowledge that is sold) and zero first-period consumption (which maximizes the novice's purchasing power).

The initial knowledge gift and money loan must now solve:

$$\begin{aligned} \max_{X_1 > 0, m_1 \leq 0} \quad & m_1 + \delta P(1, X_1) \\ \text{s.t. } \; & P(1, X_1) = (1+r)\left[f(X_1) - m_1\right]. \end{aligned}$$

We learn that offering a loan is counterproductive: while it increases the novice's purchasing power during the subsequent knowledge sale, it can only be recovered if the expert simultaneously increases the price $P(1, X_1)$ by shrinking X_1, which in turn hurts the expert as the first-period output is her only net source of revenues.

Finally, the optimal knowledge gift, obtained by setting $P(1, X_1) = (1+r)f(X_1)$, satisfies

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

[21] An equivalent interpretation is that the principal sells X_1 at an arbitrary price $p > 0$ to be paid in period 2. Then, in period 2, the principal collects p and simultaneously sells $X_2 - X_1$ at a discounted price $P(X_2, X_1) - p$. Such arrangement, however, is economically identical to that described in the remark (in which $p = 0$).

This gift is increasing in δ and approaches 100% of knowledge as δ approaches 1. The reason is that as players become more patient, the present value of any incremental knowledge grows without bound, which in turn allows the expert to extract all of the novice's first-period output in exchange for an ever smaller second-period knowledge transfer. Note also that the arrangement creates a deadweight loss equal to $f(1) - f(X_1) = (1-\delta)f(1)$. As δ approaches 1, this loss vanishes.

As we shall see next, once the expert is free to select the overall length of the contract, allowing her to capture output produced over multiple periods, she is in much less of a rush to increase the novice's productivity. As a result, she opts for lengthy contracts in which the above comparative statics are overturned.

4.2 Optimal contracts of unrestricted length

Here we return to the general setting in which contracts are allowed to have an arbitrary length. This setting allows for a rich set of contracts in which the expert may spread knowledge gifts and sales over arbitrarily many periods and may also use monetary rewards to keep the novice in the relationship. We begin by focusing on the relaxed problem in which the expert's incentive constraints are ignored:

$$\max_{\mathcal{C}\in\Omega} \Pi_1(\mathcal{C}) = \max_{\mathcal{C}\in\Omega} \sum_{t=1}^{\infty} \delta^{t-1} m_t \qquad \text{(II)}$$
$$\text{s.t. (2) and (3).}$$

Lemma 2 rules out contracts of infinite length in which training is carried out over infinitely many periods:

Lemma 2 *Every contract $\mathcal{C}$ that solves the expert's relaxed problem (II) has a finite training phase, namely, $X_t = X_{\sup}(\mathcal{C})$ for some t.*

Intuition is as follows. Consider a contract of infinite length. From the standpoint of any given date, the total (discounted) transfer that the expert can hope to extract from the novice is 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, the expert can instead end the contract early by selling all remaining knowledge at once, which the novice can eventually afford, and benefit from earlier revenues.

We now define a class of contracts, which we call "delayed-reward" contracts, and note that the expert can restrict attention to this class.[22]

Definition 1 *A finite contract $\mathcal{C}$ is a **delayed-reward contract** if it requires that 100% of output produced before a given period T is transferred to the expert and, in return, the expert grants the novice a one-time cash bonus M during period T, after which all transactions end and the novice keeps 100% of output. The net money transfers under such contract satisfy:*

$$m_1 = 0,\ m_t = (1+r) f(X_{t-1}) \text{ for } t \in \{2, ..., T-1\},\ \text{ and } m_T = (1+r) f(X_{T-1}) - M.$$

In such arrangements, the novice is willing to endure the training phase, during which he receives zero output, in exchange for both additional training and the eventual cash bonus M. Since in the present baseline model the novice cares only about his total (discounted) consumption and not its timing, these arrangements are efficient and also maximize the novice's incentives to remain in the relationship.

Remark 2 *Let $\mathcal{C} \in \Omega$ be a finite contract that satisfies constraints (2) and (3). There exists a delayed-reward contract $\mathcal{C}' \in \Omega$ that satisfies both constraints and delivers the same equilibrium payoffs Π_1, V_1 as contract $\mathcal{C}$.*

In a delayed-reward contract, the novice's incentive constraint for period $t < T$ takes the following form:

$$\underbrace{(1+r)\sum_{\tau=t}^{T-1}\delta^{\tau-t}f(X_\tau)}_{\text{Output surrendered to the expert}} \quad \leq \quad \underbrace{\sum_{\tau=t}^{T-1}\delta^{\tau-t}P(X_{\tau+1}, X_\tau) + \delta^{T-t}M}_{\text{Future knowledge and bonus rewards}}, \tag{7}$$

which captures the fact that, knowledge being general, the novice is only willing to work for the expert in exchange for future knowledge and money transfers, not for past ones. In present discounted terms, the L.H.S. is the output the novice is asked to surrender during the remaining periods of training and the R.H.S. is the benefit the novice receives

[22] An economically equivalent arrangement is one in which the bonus M is paid over multiple periods $t \geq T$ (while keeping its present discounted value constant) rather than being concentrated on period T alone. We also refer to such an arrangement as a delayed reward contract.

in return, namely, the value of the knowledge $X_T - X_t$ he is yet to acquire, plus the bonus.[23]

This constraint places an upper bound on the knowledge that the novice can be trusted with at any point during training: a higher X_t means that the novice is asked to surrender more output during training and, simultaneously, it lowers the value of the additional knowledge he is offered as compensation. It also shows that the expert can increase X_t, and therefore the novice's productivity, in three ways: a larger total knowledge transfer X_T, a larger bonus, and an earlier graduation date.

Lemma 3 shows that every optimal delayed-reward contract prescribes a sufficiently fast knowledge transfer so that, in any given period, the novice is able to produce no less than the incremental value of the knowledge he acquires (in other words, the expert keeps only the minimal amount of knowledge needed to prevent the novice from leaving):

Lemma 3 *Suppose a delayed-reward contract $\mathcal{C}$ (prescribing a one-time bonus M during period T) solves problem (II). Then, $\mathcal{C}$ has the property that, in every period $t > 1$, the novice pays a gross price no lower than the marginal value of the knowledge he acquires:*

$$m_t + M \cdot \mathbf{1}_{\{t=T\}} \geq P(X_t, X_{t-1}). \tag{8}$$

This result tells us that the novice gains surplus, at most, at the very beginning of the relationship (through a knowledge gift) and at the very end (through the bonus). In every other period, the novice surrenders all his output in exchange for knowledge that he values in no more than this output. Intuitively, if the expert is to gift any amount of knowledge (for instance, by selling a fraction of knowledge at a price lower than marginal value), it is most efficient to concentrate this gift at the beginning of the relationship as it allows the novice to raise his productivity sooner; and if the expert is to offer a money bonus, for incentive reasons it is best to defer this bonus until all other transactions are over. Note also that, when the novice is expecting a bonus, he may be willing to pay a price larger than $P(X_t, X_{t-1})$ for the right to remain in the relationship.

[23] In the parlance of relational contracts, this incentive constraint represents a "dynamic enforcement" constraint that indicates the extent to which future surplus can be recruited to enforce current money transfers from the novice. In the present setting, only a fraction of future surplus (the fraction pocketed by the novice) can be used for this purpose. In contrast, in the standard principal-agent setting, in which both players have deep pockets and therefore surplus can be freely redistributed across them (e.g. Levin, 2003), the players' entire continuation surplus can be recruited to enforce money transfers in any direction.

We now return to the expert's original problem. Proposition 1 shows that the bonus can be dispensed with and describes the overall structure of an optimal contract:

Proposition 1 *In the baseline model, every optimal contract – solving problem (I) – consists of a knowledge gift in period 1 followed by a sequence of spot sales contracts in which the novice devotes all period t output (plus interest) to purchase knowledge $X_{t+1}-X_t$ in period $t+1$ at a price equal to marginal value $P(X_{t+1}, X_t)$. This process continues until 100% of knowledge has been transferred, after which all transactions end.*

An optimal contract is equivalent to an "apprenticeship" (work-for-training) arrangement in which, until the knowledge transfer is complete, the novice works for the expert and is compensated exclusively through additional knowledge. Specifically, during each period before the knowledge transfer is complete, the novice is asked to produce output $f(X_t)$ for the expert, rather than for himself; in return, at the beginning of the following period, the novice is granted a package of knowledge $[X_t, X_{t+1}]$ that has a (discounted) value $\frac{\delta}{1-\delta}P(X_{t+1}, X_t)$ just high enough to offset this output loss. The overall length of this apprenticeship is controlled by the size of the initial knowledge gift, with a smaller gift leading to a more distant graduation.

Intuition for Proposition 1 is as follows. First, recall that for a given total knowledge transfer, the expert has two instruments for raising the novice's productivity during training: a bonus and an earlier graduation date. Of the two, the bonus is a finer instrument as, unlike an earlier graduation date, it allows the expert to raise output continuously. However, since the incentive constraints (7) are linear in both output and the bonus, whenever a bonus is used the expert adopts a "corner" solution in which the bonus is raised to the point that output hits its upper bound $f(1)$ during some period before the contract is over. This large bonus can then be replaced, without loss, by a shorter contract with the same output path but that allows the novice to graduate as soon as output reaches $f(1)$.[24] Second, the price of knowledge is derived from Lemma 3: on the one hand, any knowledge sold at a price lower than marginal value amounts to a fraction of knowledge being gifted, but the novice's productivity can be raised by speeding up all such

[24] The optimal output path can also be implemented with bonuses of varying size. Every such implementation calls for a bonus that is a multiple of the efficient output $f(1)$ (plus appropriate interest), which results in the expert retaining the novice for a number of periods after the knowledge transfer is complete. All such implementations, however, call for a net transfer of cash from the expert in the last period of interaction, which is not feasible once the expert's participation constraints are imposed. Uniqueness of the optimal arrangement follows from this fact.

gifts and concentrating them in the first period of the relationship, before the knowledge sale begins; on the other hand, once the bonus is set to zero, and given that the price of knowledge never falls below marginal value, in no period is the novice willing to pay more than marginal value. Finally, given that the expert captures all surplus during the knowledge sale, she finds it optimal to sell all of her knowledge.

Before we further characterize the optimal contract, some remarks are in order:

- When selecting the optimal length of the apprenticeship, the expert faces the following trade-off: by raising the novice's productivity, a larger knowledge gift allows the expert to extract higher revenues during each training period; but since it also reduces the amount of knowledge left for sale, it reduces the number of training periods that the novice is willing to withstand. As we shall see, this trade-off, which is central to the remainder of the paper, results in lengthy contracts in which significant output is wasted.

- The above apprenticeship calls for an extreme form of training in which the novice is forced to consume zero until training is complete, at which point he graduates and his consumption experiences a sudden increase. Below, we also consider less extreme arrangements in which, owing to either government intervention or the novice being risk averse, consumption is positive (though still low) during training.

- We have assumed that any deviation from a contract's prescriptions is punished through permanent separation. Since separation typically results in a destruction of output, it need not be immune to renegotiation. Fortunately, because of its forward-looking nature, the above apprenticeship can be recast to avoid this pitfall, as follows. After receiving the knowledge gift, the novice is invited, every period, to purchase any amount of knowledge he desires at a price equal to the marginal value P of the knowledge he wishes to acquire. Even if, in a given period, the expert deviates by not selling knowledge under these terms, or the novice deviates by failing to make such a purchase (or by devoting only part of his output to the purchase), the players remain invited to meet again the following period, under the same terms as before.

4.3 Optimal contracts in closed form: a $\frac{1}{e}$ rule

Here we derive the optimal knowledge gift and apprenticeship length. We begin by expressing the expert's problem in reduced form, as a function of X_1. The apprenticeship's overall length, denoted $\mathcal{T}(X_1)$, is determined by the number of periods it takes for the novice to afford the remaining knowledge $1 - X_1$ at a price equal to marginal value. A consequence of knowledge being sold at this price, and the novice exhausting all his cash holdings when paying such price, is that both output and the expert's per-period revenues grow at the interest rate r until training is complete.[25] As a result, $\mathcal{T}(X_1) - 1$ corresponds to the number of periods of compound growth at this rate required for output to reach $f(1)$ starting from level $f(X_1)$, namely,

$$\mathcal{T}(X_1) - 1 = \frac{1}{|\log \delta|} \log \left[\frac{f(1)}{f(X_1)} \right]. \tag{9}$$

Moreover, since the expert's per-period revenues also experience compound growth at rate r, the expert's discounted revenue in period $t \leq \mathcal{T}(X_1)$ is $\delta^{t-1} m_t = f(X_1)$. Combining these two observations, the expert's reduced-form problem is

$$\max_{X_1 \geq 0} \quad \underbrace{[\mathcal{T}(X_1) - 1]}_{\#\text{ of transfer periods}} \quad \cdot \quad \underbrace{f(X_1)}_{\text{discounted transfer per period}}. \tag{10}$$

The objective captures the expert's trade-off: a higher gift means that less knowledge is left for sale and so the expert enjoys fewer periods of revenues (first term), but it also means that the remaining knowledge is sold more quickly and so the expert enjoys more revenues during each spot transaction (second term). Proposition 2 describes the solution:

Proposition 2 *Up to an integer constraint for an apprenticeship's length, the optimal knowledge gift X_1 satisfies*

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

where e is Euler's number. As a result, the optimal apprenticeship has length $\mathcal{T}(X_1) = \frac{1}{|\log \delta|} + 1$, which is increasing in δ and independent of the production technology.

This result tells us that no matter how patient players are, and regardless of the details of f, the expert optimally balances her conflicting goals by allowing the novice to

[25] Namely, $m_t = (1+r)f(X_{t-1}) = P(X_t, X_{t-1})$ implies $(1+r)f(X_{t-1}) = f(X_t)$ and $(1+r)m_t = m_{t+1}$.

produce, at the start of the apprenticeship, a share $\frac{1}{e}$ of the efficient output level. Indeed, upon combining (9) and (10), the expert's problem is equivalent to that of maximizing the average logarithm $\log\left[\frac{f(1)}{f(X_1)}\right]\cdot\left[\frac{f(1)}{f(X_1)}\right]^{-1}$ of the output ratio $\frac{f(1)}{f(X_1)}$. The maximum is achieved when this ratio is e. When each period consumes an arbitrarily small amount of time (as formalized in Section 4.3.B below), this solution also has the property of equating the present value of knowledge gifted ($\frac{1}{e}P(1,0)$) to the present value of knowledge sold ($\sum_{t=2}^{\mathcal{T}(X_1)}\delta^{t-1}P(X_t,X_{t-1})$). As a result, surplus is split 50-50 between the two players.

The present solution coincides with the solution to the "secretary problem" in which a recruiter who faces a queue of job applicants of unknown quality must decide what fraction 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"). While the two problems have the same solution, they 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.[26]

Notice that the optimal apprenticeship is longer, and knowledge is transferred more slowly, the more patient the players. Intuitively, when the interest rate r falls, knowledge becomes more valuable, in the margin, to the novice (as he can use this knowledge during every subsequent period of his life). As a result, a lower interest rate means that, each period, the novice is willing to work for the expert in exchange for a smaller amount of new knowledge. The expert takes advantage of this fact by stretching out the training phase and keeping the novice with her longer. (Notice also the contrast with the two-period contract derived in Section 4.1. There, the knowledge gift grows as r falls, and converges to 100% of knowledge as r converges to 0.)

Consistent with real-world practices noted in Section 2, the expert's $\frac{1}{e}$ rule causes lengthy apprenticeships. For instance, training takes approximately 9.5 years when the annual interest rate is 10% and approximately 19.5 years when the interest rate is 5%. This rule also means that the details of f do not affect the length of the apprenticeship. The implication is that, in the absence of other factors affecting the relationship, novices of different skill levels, and novices working on different tasks, take equally long to train.

Before turning to policy implications, we discuss the inefficiencies caused by these long training phases and well as the impact of increasing the players' frequency of interaction (while holding constant their underlying discount rate):

[26] We are grateful to Thomas Bruss for providing insights on the secretary problem.

A. Efficiency and the role of δ

Here we consider the deadweight loss associated to the optimal contract. Recall that the contract is socially inefficient as it artificially spreads the knowledge transfer over multiple periods. The resulting loss in surplus (namely, the present value of the loss in output relative to first best) is

$$\sum_{t=1}^{|\log \delta|^{-1}} \left[\delta^{t-1} - \frac{1}{e}\right] f(1).$$

This loss increases with δ for two reasons. First, a higher δ leads to a longer training phase and, with it, a slower knowledge transfer. Second, a higher δ implies that the inefficiencies caused by the slower transfer loom larger from the perspective of period 1.

When measured as a fraction of first-best surplus, $P(1,0)$, the loss is $1 - \frac{1}{e} - \frac{1-\delta}{|\ln \delta|}\frac{1}{e}$. In contrast to the absolute loss, this relative loss falls with δ because first-best surplus grows with δ faster than the deadweight loss. However, as δ converges to 1 this loss remains bounded away from zero: it is no smaller than $\frac{e-2}{e} \simeq 25\%$ of first-best surplus.

B. Frequency of interaction

Here we ask whether allowing players to interact with higher frequency (multiple times during each date t) alters the length of the optimal apprenticeship, measured in calendar time. One may wonder, for instance, if the apprenticeship consumes an arbitrary short amount of time when players are allowed to interact arbitrarily often.[27] We show that, modulo the time it takes to exchange the initial knowledge gift (which is the length of time consumed by stage 1 in a given period), the overall time consumed by the apprenticeship is invariant in the frequency of interaction.

We assume that date t consumes one year and consists of $n \geq 1$ periods. Accordingly, δ is the annual discount factor and $\delta^{1/n}$ the per-period factor. Denote a period by τ and let per-period output be $f(X_\tau, n)$, with $f_1 > 0$, $f_2 < 0$, and $\lim_{n\to\infty} n \cdot f(1,n) < \infty$.[28] A given period τ within date t consumes $\frac{1}{n}$ years. Note that this setting is equivalent to the baseline model with $\delta^{1/n}$ taking the place of δ and $f(X_\tau, n)$ taking the place of $f(X_t)$.

[27] We are grateful to Larry Samuelson for suggesting this problem.

[28] A special case of interest, with slight abuse of notation, is $f(X_\tau, n) \equiv \frac{1}{n} f(X_\tau)$.

From Proposition 2, the optimal knowledge gift X_1 satisfies

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

which in turn leads the apprenticeship to last $N = \frac{1}{|\log \delta^{1/n}|} + 1$ periods. Accordingly, the actual time T consumed by the apprenticeship, measured from the beginning of the first period to the end of stage 1 of the last period, is obtained by adding the time consumed by the knowledge gift and the time consumed by the subsequent $N - 1$ periods over which the remaining knowledge is sold:

$$T = \frac{1}{n}[\lambda + N - 1] = \underbrace{\frac{\lambda}{n}}_{\text{knowledge gift}} + \underbrace{\frac{1}{|\log \delta|}}_{\text{knowledge sale}},$$

where $\lambda \in [0, 1]$ is the fraction of time consumed by stage 1 alone in any given date.

Therefore, excluding the initial gift stage, the contract's duration is invariant in n. Intuitively, the overall value of the knowledge $1 - X_1$ sold, measured as a fraction of the efficient output $f(1, n)$, depends only on the underlying discount factor δ and not on the frequency of interaction. As a result, the novice must spend an invariant amount of time in training, working for expert, in order to afford this knowledge.

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).[29] Given the difficulties we discussed in Section 2, good policy is crucial in this area. As the OECD put it, "Quality apprenticeships require good governance to prevent misuse as a form of cheap labour."[30]

[29] 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.)

[30] Governments have long been interested in regulating apprenticeships. See, for example, Malcomson et al. (2003) for a discussion and Elbaum (1989) for a historical perspective. As noted in the introduction, Malcomson et al. consider an example of regulation. There, rather than being forced to shorten their training (which is one of the experiments we study here), firms are forced to pay low wages over a minimum time period after training is over. The authors show that this seemingly counter-intuitive

Motivated by such concerns, we consider two policy experiments: a minimum wage (or minimum consumption level) during training and a limit on the contract'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 requires that the expert pay the novice a wage of at least $w_{\min} > 0$ during each period of the relationship. In the present formulation of the model, with the novice acting as the residual claimant of output, this policy is captured by setting all wages to zero and instead requiring that the novice keeps at least $w_{\min}$ each period for consumption – namely, $c_t \geq w_{\min}$ for all $t \leq T$.[31] So that the expert is willing to contract with the novice, we assume $w_{\min} < f(1)$.

Corollary 1 tells us that the optimal contract retains the basic properties of the apprenticeship characterized in Proposition 1:

Corollary 1 *Under a minimum wage policy, every optimal contract consists of a knowledge gift in period 1 followed by a sequence of spot contracts in which the novice devotes all period t output, minus the minimum wage, to purchase knowledge $X_{t+1} - X_t$ in period $t+1$ at a price equal to marginal value $P(X_{t+1}, X_t)$. This process continues until 100% of knowledge has been transferred.*

As in the baseline model, the expert concentrates the knowledge gift in the first period, after which she extracts all of the novice's surplus via prices equal to marginal value. The difference is that the novice must now split his output between consuming $w_{\min}$ and purchasing new knowledge. Consequently, for any fixed knowledge gift, the policy has the effect of reducing the magnitude of each spot transaction, which in turn slows down the speed at which knowledge is transferred. As we shall see, this fact biases the expert's trade-off (between the amount of knowledge sold and the speed of this sale) in favor of a smaller overall sale: only by lowering the amount of knowledge sold can the expert counteract the lower speed of sale caused by the minimum wage.

regulation may be beneficial when information asymmetries prevent workers from leaving the firm, and the firm is capable of committing to future wages.

[31] Since the novice is risk neutral, he is indifferent between literally consuming $w_{\min}$ in a given period and saving this money for future consumption. We assume that the regulation forbids any such savings from being transferred back to the expert.

The corollary allows us to express the expert's problem in reduced form, with X_1 serving as her choice variable. Recall that in the baseline contract, after the initial gift, output grows at rate r until training is complete. Under the present policy, it is net output $f(X_t) - w_{\min}$ that grows at rate r.[32] The length of the apprenticeship, which we now denote $\mathcal{T}(X_1, w_{\min})$, is in turn pinned down by the number of periods of compound growth required for net output to reach its final value $f(1) - w_{\min}$.

Moreover, since the novice's net output during training is transformed into revenues for the expert, such revenues also grow at rate r. As a result, the expert's discounted revenues remain constant throughout the sales process: $\delta^{t-1} m_t = f(X_1) - w_{\min}$. The expert's reduced problem becomes

$$\max_{X_1 \geq 0} \underbrace{[\mathcal{T}(X_1, w_{\min}) - 1]}_{\#\text{ of sales periods}} \cdot \underbrace{[f(X_1) - w_{\min}]}_{\text{discounted transfer per period}}. \tag{11}$$

Proposition 3 describes the solution:

Proposition 3 *Consider a minimum wage policy with wage $w_{\min} > 0$. Up to an integer constraint for an apprenticeship's length, the optimal knowledge gift X_1 satisfies*

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

Consequently, the optimal apprenticeship has a length that is independent of $w_{\min}$ and, during training, it prescribes an output path that is uniformly increasing in $w_{\min}$.

This result tells us that the expert confronts the policy by transferring additional knowledge to the novice while holding constant the length of the relationship. Specifically, up to the moment of graduation, the policy shifts the entire output path upward and, at the same time, reduces its slope.[33] The implication is that the policy is surplus enhancing. Intuitively, transferring additional knowledge helps the expert partially offset the expense caused by the minimum wage, which is attractive to the expert because it allows her to

[32] To see why, notice that the spot transaction in period $t+1$ satisfies,

$$\underbrace{(1+r)\,[f(X_t) - w_{\min}]}_{\text{cash devoted to purchase}} = \underbrace{P(X_{t+1}, X_t)}_{\text{price of purchase}},$$

which upon rearranging terms delivers $(1+r)\,[f(X_t) - w_{\min}] = f(X_{t+1}) - w_{\min}$.

[33] In addition, $f(X_1)$ converges to $f(1)$, and the deadweight loss converges to zero, as $w_{\min}$ converges to $f(1)$.

counteract the fact that the minimum wage, before the output path is adjusted, forces the knowledge sale to slow down. The policy also increases the novice's payoff: it does not affect his graduation date (or the prize received upon graduation) and yet provides him with positive consumption before that date.

That the length of the apprenticeship is invariant in $w_{\min}$ can be understood as follows. Recall, from the baseline model, that the optimal balance between the expert's conflicting goals (selling more knowledge vs. selling it faster) is achieved by selling knowledge over $\frac{1}{|\log \delta|}$ periods, regardless of the details of f. Given that the minimum wage affects the expert's per period revenues in an analogous way as constant drop in f, it also leaves the contract's length unaffected.

Limit on contract duration. An alternative intervention is a requirement that the expert's interactions with the novice end after some number $T_{\max}$ of periods. When binding, this requirement forces the expert to sell her knowledge faster and therefore to sell less of it. The result is a higher knowledge gift and a uniformly higher level of output.

It is worth noting that the above policies might backfire when the expert does not enjoy rents to begin with. For a simple example, suppose two identical experts compete face to face by simultaneously offering contracts to a single novice, and suppose each expert must pay a fixed cost $F > 0$ whenever contracting for the first time with the novice (regardless of the novice's level of knowledge).[34] The equilibrium contract is an apprenticeship with the properties in Proposition 1 and a duration that is just long enough for the expert in question to recover F. If experts are required to pay a minimum wage, they must increase the fraction of knowledge sold in order to recover this cost. The result is a loss of surplus in the form of a lower knowledge gift and a uniformly lower level of output after that. Even worse, if experts are required to reduce the contract's duration below its original equilibrium level, it is impossible for them to recover the cost. As a result, they exit the market.

6 Extensions

Here we extend the baseline model in several ways. First, we consider the case in which the novice has concave utility. Second, we return to the baseline case and consider a few

[34] In a richer model, which we do not presently attempt, experts would face varying degrees of competition as well as other types of costs.

simple modifications of practical interest – training costs, the novice arriving with capital, and externalities on the expert – all of which alter the length of the apprenticeship via the ratio of knowledge gifted to sold. Finally, we describe the set of Pareto-efficient contracts.

6.1 Consumption smoothing

Consider the case in which the novice derives instantaneous utility $u(c_t)$ from consumption. The novice's period-t continuation payoff is now

$$V_t(\mathcal{C}) = \sum_{\tau \geq t} \delta^{\tau - t} u\left(c_\tau\right),$$

with associated budget constraint $\sum_{\tau \geq t} \delta^{\tau-t} c_\tau = (1+r)\, B_{t-1} + \sum_{\tau \geq t} \delta^{\tau - t}\left[f(X_\tau) - m_\tau\right]$. We assume a constant intertemporal elasticity of substitution (CIES), namely, $u\left(c\right) = \frac{c^{1-\sigma}}{1-\sigma}$ for some $\sigma > 0$. This restriction makes it possible to derive a partial analytical characterization of the optimal contract, which we then supplement with a numerical solution. (A degree of consumption smoothing also arises if the novice requires a minimum "subsistence" level of consumption $c_{\min} > 0$ each period. In this case, the optimal contract is the same as under a minimum wage policy – see Section 5 for the case in which, beyond the subsistence level, utility is linear and see below for the case in which, beyond this level, utility is CIES.) To simplify notation, denote period-t output by $y_t = f(X_t)$ and let $\overline{y} = f(1)$.

The only constraints that differ relative to the baseline model are the novice's incentive constraints. These constraints are derived as follows. Since the novice has a preference for smooth consumption, her most tempting deviations arise after output y_t is produced but before consumption takes place (namely, at the beginning of stage 3 of the model). In such a deviation, the novice walks away from the expert and perfectly smooths consumption by setting $c_\tau = y_t + b_t$ for all $\tau \geq t$, where b_t is the interest on the novice's current cash balance. As a result, the new incentive constraints are $\frac{1}{1-\delta} u\left(y_t + b_t\right) \leq V_t(\mathcal{C})$ for all t.

As in the baseline model, the expert may restrict attention to contracts in which, at the beginning of each period, the novice transfers all his available cash to the expert in exchange for additional knowledge. In addition, the expert promises the novice a consumption profile during training $c_1, ..., c_{T(\mathcal{C})-1}$, where $T(\mathcal{C})$ is the contract's final training period, together with the graduation "prize" $c_t = f(1)$ for all $t \geq T(\mathcal{C})$.[35] The expert's

[35] The expert may in principle also promise the novice a bonus M after training is complete. As in the

problem is therefore

$$\max_{\mathcal{C}\in\Omega} \sum_{t=1}^{T(\mathcal{C})-1} \delta^{t-1}\left[y_t - c_t\right]$$

$$s.t.\ \frac{1}{1-\delta}u\left(y_t\right) \leq V_t(\mathcal{C}) \text{ for all } t < T(\mathcal{C}) \text{ and } V_{T(\mathcal{C})}(\mathcal{C}) = \frac{1}{1-\delta}u(\overline{y}).$$

Proposition 4 links the novice's optimal consumption and output profiles:

Proposition 4 *When the novice has concave utility $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ ($\sigma > 0$), an optimal contract $\mathcal{C}$ prescribes, during training, the increasing consumption path*

$$c_t = (1-\delta)^{\frac{1}{\sigma}} Y_t \textit{ for all } t < T(\mathcal{C}),$$

where $Y_t = \left(\sum_{\tau\leq t} y_\tau^\sigma\right)^{\frac{1}{\sigma}}$.

During each training period before the knowledge transfer is complete, the novice is asked to sacrifice utility $u(y_t) - u(c_t)$. As compensation, at the beginning of the following period, the novice is granted a package of knowledge $[X_t, X_{t+1}]$ with a (discounted) value $\frac{\delta}{1-\delta}\left[u(y_{t+1}) - u(y_t)\right]$ just high enough to offset this utility loss. An increasing consumption path represents a compromise between delaying consumption (which has the benefit of increasing the fraction of output that the novice devotes to buying further knowledge) and smoothing consumption (which helps the novice endure the training phase). Since knowledge purchases are most critical early on (as they expand output in every subsequent period), 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 and the resulting output path is both increasing and convex.

We derive the remaining details of the arrangement numerically.[36] The results, which we illustrate in Figure 1 (at the end of this document), are as follows:

1. As the novice's elasticity of intertemporal substitution $\frac{1}{\sigma}$ falls, he consumes a higher fraction of output during training. Consequently, the overall knowledge sale slows

baseline model, however, this bonus can be dispensed with by instead promising the novice an earlier graduation date.

[36]For any given contract length T, the profile $(c_t, y_t)_{t=1}^{T-1}$ solves $2(T-1)$ equations: $c_t = (1-\delta)^{\frac{1}{\sigma}} Y_t$ and $u(y_t) - u(c_t) = \frac{\delta}{1-\delta}\left[u(y_{t+1}) - u(y_t)\right]$ for all $t < T$. The optimal contract is then obtained by optimizing over T.

down and the apprenticeship becomes lengthier – and is always lengthier than in the baseline model (Figure 1.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 1.B).

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 1.C).[37]

6.2 Other motives for altering the apprenticeship length

The baseline model can be readily extended in several other ways. Here we describe some examples. In all of them, the optimal contracts are apprenticeships with the properties in Proposition 1.[38] As a result, the modifications that follow alter only the length of the apprenticeship via the fraction of knowledge that is gifted rather than sold.

Training costs. Up to now, in order to emphasize the expert's desire for slow training, we have assumed that knowledge transfers are free. When training costs are introduced, training may further slow down. As an example, suppose that raising the novice's knowledge from X to X' costs the expert $\beta P(X', X)$ (namely, a fraction of the marginal value of $X' - X$) for some constant $\beta \in (0, 1)$. (Equivalently, the cost could be a constant fraction of the incremental output $f(X') - f(X)$.)[39] Once this cost is considered, the optimal knowledge gift satisfies

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

where $A = \frac{\beta}{1-\beta}\frac{|\ln \delta|}{1-\delta}$ (this constant is approximately equal to $\frac{\beta}{1-\beta}$ when δ is close to 1). As one might expect, a higher cost results in a smaller gift and a longer training phase.

[37] In contrast to an increase in σ, the policy has a pronounced impact on output early in the relationship. Consequently, the overall training phase is (slightly) reduced despite the fact that the policy uniformly raises consumption.

[38] The proof of this claim is a straightforward extension of the proof of Proposition 1. From brevity, we leave this proof to the reader. (It is also available upon request.)

[39] As a function of X_1, the expert's payoff is $[\mathcal{T}(X_1) - 1]\,(1-\beta) f(X_1) - \beta P(X_1, 0)$, which shows that the transmission cost reduces the net per-period revenues from the second period onwards (first term) and also makes the gift itself costly (second term).

Novice's liquidity. In the baseline model, the novice has no capital to begin with. Suppose instead that the novice arrives with capital $L > 0$. Suppose also that L does not exceed the value of all knowledge $P(1,0)$ – otherwise, the expert would simply sell all knowledge in the first period and keep all surplus to herself. In an optimal arrangement, the expert asks the novice to surrender all capital up front and, in return, offers the novice an apprenticeship that is worth, in present discounted terms, no less than L.[40] Since an apprenticeship with knowledge gift X_1 is worth $P(X_1, 0)$ to the novice, the expert faces the new participation constraint $L \leq P(X_1, 0)$. The optimal value of X_1 is obtained by maximizing the expert's revenues during training, above and beyond L, subject to this new constraint. The result is that the novice's access to capital weakly accelerates training while also weakly *decreasing* the novice's rents. In fact, when L is insufficient to cover the value of the baseline knowledge gift ($L < \frac{1}{e}P(1,0)$) the expert merely extracts the novice's capital without altering any other feature of the apprenticeship. Otherwise, X_1 satisfies $L = P(X_1, 0)$, which tells us that a more liquid novice is trained faster, but obtains zero rents.

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 might 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 $\gamma f(X_t)$ each period, with $\gamma > -1$ capturing the magnitude and sign of the externality.[41] In this case, the optimal knowledge gift satisfies

$$\frac{f(X_1)}{f(1)} = \frac{1}{e} \cdot \delta^{-\frac{\gamma}{(1+\gamma)(1-\delta)}},$$

which, as expected, is increasing in γ: a larger externality accelerates training. When γ is sufficiently large, training ends in one period; when γ approaches -1, training is stretched to infinity.

[40] As shown in Section 6.3, using an apprenticeship with the properties in Proposition 1 guarantees that the knowledge transfer is structured efficiently.

[41] As a function of X_1, the expert's payoff is $[\mathcal{T}(X_1) - 1]\,(1+\gamma)f(X_1) + \gamma\delta^{[\mathcal{T}(X_1)-1]}P(1,0)$, which captures the impact of the externality during training (embedded in the first term) and during every period after that (second term).

6.3 Pareto-optimal contracts

Here we characterize the broader class of Pareto-efficient contracts that maximize a weighted average of the players' payoffs. The exercise captures, in reduced form, instances in which the novice has bargaining power ex-ante – for example, owing to the existence of alternative occupations or alternative experts. (This simple setting does not capture the potential impact of a competing expert who might attempt to poach the novice while he is still being trained. We leave this possibility to future work.) Letting $\lambda \in (0,1)$ be the novice's Pareto weight, an efficient contract solves[42]

$$\max_{\mathcal{C}\in\Omega} \ \lambda V_1(\mathcal{C}) + (1-\lambda)\Pi_1(\mathcal{C}) \qquad \text{(III)}$$
$$s.t. \ (1), \ (2), \ \text{and} \ (3).$$

Corollary 2 shows that the family of efficient contracts is a family of apprenticeships. The length of these apprenticeships ranges from the baseline level $\frac{1}{|\log\delta|} - 1$ to a single period, depending on the novice's bargaining power:[43]

Corollary 2 *Every efficient contract – solving problem (III) – satisfies the properties in Proposition 1 and prescribes a knowledge gift such that*

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

This gift is increasing in the novice's Pareto weight λ up to the point in which all knowledge is gifted.

The above contracts have the property that all surplus that is granted to the novice is granted exclusively through a knowledge gift, rather than consumption, and this gift is concentrated in the first period. The expert, for instance, never offers the novice a discount during the spot knowledge sales. This feature is shared by every efficient contract because it guarantees that the novice's productivity is raised as early as possible.[44]

---

[42] To avoid knife-edge cases in which problem (III) has multiple solutions, we assume that δ and λ are "generic" in the sense that $\frac{1-2\lambda}{1-\delta} \neq n$ for all $n \in \mathbb{N}$.

[43] When $\lambda > \frac{1}{2}$ all knowledge is transferred in period 1. In addition, since players are risk neutral, problem (III) calls for an unbounded money transfer from expert to novice.

[44] The Pareto frontier is given by the equation $V_1(\Pi_1) = -\frac{\Pi_1}{W(-z)}$, where $z = \frac{\Pi_1|\ln\delta|}{f(1)} \in [0, \frac{1}{e}]$ and $W(\cdot)$ is the Lambert W (or Product Log) function. When players interact infinitely often, the allocation of surplus ranges from a 50-50 split (when $\lambda = 0$) to all surplus going to the agent (when $\lambda \geq \frac{1}{2}$).

7 Conclusion

We study how to structure the dynamic relationship between an expert and a novice to accomplish knowledge transfers under limited enforceability. The optimal contract takes the form of an apprenticeship in which labor is traded for knowledge. At the start of the relationship, the novice is trained just enough to produce a fraction $\frac{1}{e}$ of the efficient output; after that, and until training is complete, the novice uses all of his output to gradually pay for additional knowledge. In this arrangement, it is the promise of future knowledge sales that induces players to remain in the relationship.

Since the expert has a motive to keep knowledge in reserve, to retain the novice, she artificially delays the knowledge transfer, leading in turn to an inefficiently low level of output. (When the agent is risk averse, output is split between knowledge purchases and consumption, and the knowledge transfer further slows down.) A likely instance of this inefficiency can be found in careers within professional services. It appears, anecdotally, that juniors spend years "paying their dues" to the 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, the expert can keep the novice around with smaller knowledge transfers, which are now more valuable, as compensation. Thus, features that are traditionally considered to affect the discount factor, such as having more reliable partners, lead to slower transfers and lower productivity during training.

We also find that, provided the expert enjoys sufficient rents, simple policy interventions are surplus enhancing. In particular, both minimum wages for the apprentice and exogenous limits on the apprenticeship's length, accelerate training and therefore raise surplus. In both cases, to partially counteract the cost caused by the policy, the expert uniformly raises the novice's productivity throughout the apprenticeship.

Beyond apprenticeships, our model has implications for knowledge transfers in international alliances and joint ventures. The imperfect contractability in many developing countries, resulting from weak institutions, means that contracts between companies may exhibit the same lack of commitment that is characteristic of training relationships between individuals. As a result, the "expert" partner transferring knowledge may benefit from slowing down the rate of transfer to ensure incentive compatibility while extracting maximum rents.

The model is highly tractable and can be used as a building block for other models in which human capital acquisition is relevant. 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 would be to study, in the context of knowledge transfers between firms, how other dimensions of the firms' interaction can be used to strengthen their relationship, and therefore accelerate training. 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 training is inefficiently lengthy and whether, as we suggest, patience slows down the rate at which knowledge is transferred.

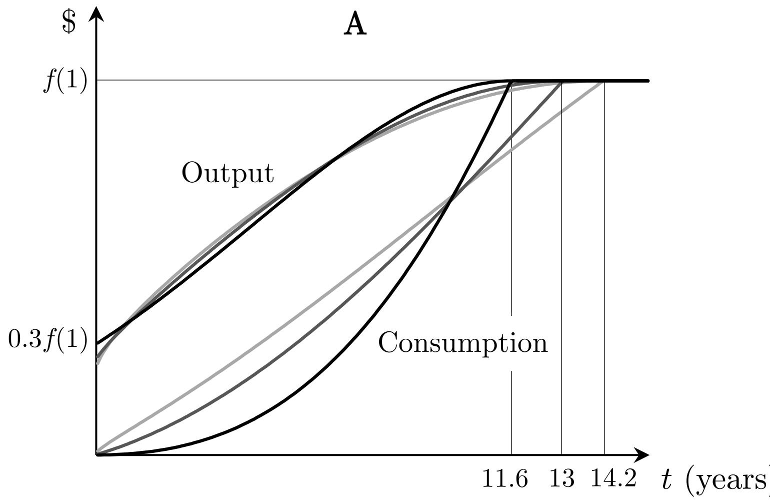

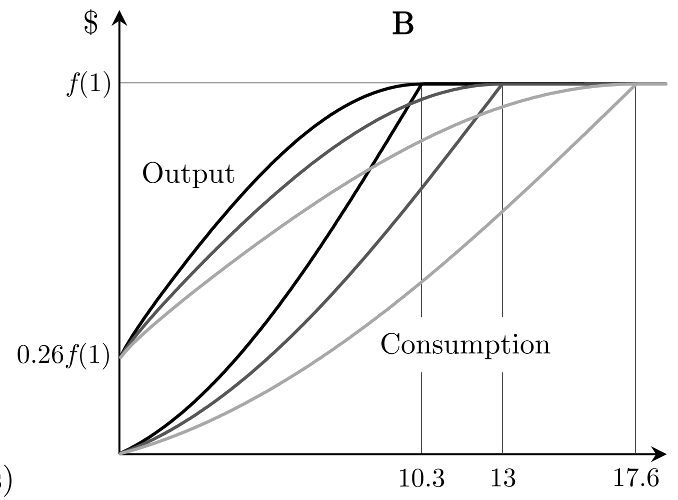

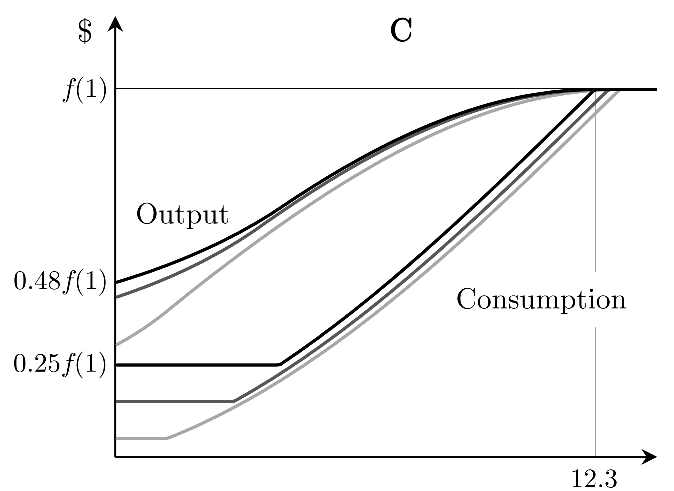

Figure 1. Apprenticeships with concave utility

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 $w_{\min}$

(δ is annual discount factor; players meet monthly)

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. Elbaum, Bernard (1989). "Why Apprenticeship Persisted in Britain But Not in the United States." *Journal of Economic History* 49(2): 337-349.

20. Garicano, L. (2000). "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy* 108(5): 874-904.

21. 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.

22. Goloboy, J. L. (2008). *Industrial Revolution: People and Perspectives*. ABC-CLIO.

23. Halac, M. (2012). "Relational contracts and the value of relationships." *American Economic Review* 102(2): 750-779.

24. Heckman, James (1971). "Three Essays on Household Labor Supply and the Demand For Market Goods." Chapter 1, Ph.D. Dissertation. Princeton University.

25. Hörner, J., and A. Skrzypacz (2010). "Selling Information." Cowles Foundation Discussion Paper No. 1743.

26. Kahn, C., and G. Huberman (1988). "Two-Sided Uncertainty and "Up-or-Out" Contracts." *Journal of Labor Economics* 6(4): 423-444.

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

28. 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.

29. Levin, J. (2003). "Relational Incentive Contracts." *American Economic Review* 93(3): 835-857.

30. Levin, J., and S. Tadelis (2005). "Profit sharing and the role of professional partnerships." *Quarterly Journal of Economics* 120(1): 131-171.

31. Li, J., and N. Matouschek (2013). "Managing conflicts in relational contracts." *American Economic Review* 103(6): 2328-2351.

32. Loewenstein, M. A., and J. R. Spletzer (1998). "Dividing the Costs and Returns to General Training." *Journal of Labor Economics* 16(1): 142-171.

33. MacLeod, W. B., and J. Malcomson (1989). "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57(2): 447-480.

34. MacLeod, W. B., and J. Malcomson (1998). "Motivation and Markets." *American Economic Review* 88(3): 388-411.

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

36. 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.

37. McVeigh, S. (2006). *Concert Life in London from Mozart to Haydn.* Cambridge University Press.

38. Prendergast, C. (1993), "The Role of Promotion in Inducing Specific Human Capital Acquisition," *Quarterly Journal of Economics* 108(2): 523-534.

39. Rayo, L. (2007). "Relational incentives and moral hazard in teams." *Review of Economic Studies* 74(3): 937-963.

40. Richter, A., M. Dickmann, and M. Graubner (2008). "Patterns of human resource management in consulting firms." *Personnel Review* 37(2): 184-202.

41. Rosen, S. (1972). "Learning and Experience in the Labor Market." *Journal of Human Resources* 7(3): 326-342.

42. Shaw, K. (1989). "Lifecycle Labor Supply with Human Capital Accumulation." *International Economic Review* 30(2): 431-456.

43. Smith, A. (1863). *An Inquiry into the nature and causes of the Wealth of Nations.* J. R. M'Culloch., Ad. and Ch. Black.

44. Stevens, M. (1994). "Labour Contracts and Efficiency in On-the-Job Training." *Economic Journal* 104(423): 408-419.

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

46. Waldman, M. (1990). "Up-or-Out" Contracts: A Signaling Perspective." *Journal of Labor Economics* 8(2): 230-250.

47. Weiss, Y. (1972). "On the Optimal Lifetime Pattern of Labour Supply." *Economic Journal* 82(328): 1293-1315.

8 Appendix

Proof of Lemma 1. Let $\mathcal{C}' = \langle x_t, m'_t, d'_t, w'_t, c_t \rangle_{t=1}^{\infty}$ be such that, for all t, $d'_t = w'_t = 0$ and

$$m'_t = m_t + (1+r)\max\{d_{t-1}f(X_{t-1}) - w_{t-1}, 0\} + \min\{d_t f(X_t) - w_t, 0\},$$

where $d_0 = w_0 = X_0 = 0$.[45] By construction, $\mathcal{C}'$ and $\mathcal{C}$ prescribe the same knowledge transfers and overall payoffs Π_1 and V_1. Moreover, for all t,

$$\Pi_t(\mathcal{C}') = \Pi_t(\mathcal{C}) + (1+r)\max\{d_{t-1}f(X_{t-1}) - w_{t-1}, 0\} \geq \Pi_t(\mathcal{C}),$$
$$V_t(\mathcal{C}') = \sum_{\tau=t}^{\infty} \delta^{\tau-t} c_t = V_t(\mathcal{C}), \text{ and } B_t(\mathcal{C}') = B_t(\mathcal{C}) + \max\{d_t f(X_t) - w_t, 0\}.$$

We now show that $\mathcal{C}'$ satisfies the stage-1 constraints (1), (2), and (3). Constraint (1) is met because $\Pi_t(\mathcal{C}') \geq \Pi_t(\mathcal{C})$ for all t. Constraint (3) is met because, for all t, $m_t \leq (1+r)B_{t-1}(\mathcal{C})$ implies

$$m'_t \leq (1+r)B_{t-1}(\mathcal{C}) + (1+r)\max\{d_{t-1}f(X_{t-1}) - w_{t-1}, 0\} + \min\{d_t f(X_t) - w_t, 0\},$$

which in no greater than $(1+r)B_{t-1}(\mathcal{C}')$. Constraint (2) is met because, for all t, either $d_{t-1}f(X_{t-1}) - w_{t-1} \leq 0$ and therefore the fact that $\mathcal{C}$ satisfies (2) in period t implies

$$V_t(\mathcal{C}) \geq \frac{1}{1-\delta} f(X_{t-1}) + \frac{1}{\delta} B_{t-1}(\mathcal{C}) = \frac{1}{1-\delta} f(X_{t-1}) + \frac{1}{\delta} B_{t-1}(\mathcal{C}'),$$

or $d_{t-1}f(X_{t-1}) - w_{t-1} > 0$ and therefore the fact that $\mathcal{C}$ satisfies (5) in period $t-1$ implies

$$\underbrace{c_{t-1} + \delta V_t(\mathcal{C})}_{V_{t-1}(\mathcal{C})} \geq \frac{1}{1-\delta} f(X_{t-1}) + (1+r) B_{t-2}(\mathcal{C}) - m_{t-1}$$

which is equivalent to $V_t(\mathcal{C}) \geq \frac{1}{1-\delta} f(X_{t-1}) + (1+r)B_{t-1}(\mathcal{C}')$.

Finally, we show that the stage-2 constraints (4), (5), and (6) are redundant. Constraint (4) is redundant because, given that $d'_t = w'_t = 0$ for all t, it is implied by constraint

[45] Namely, if in stage 2 the principal receives positive revenues ($d_t f(X_t) - w_t > 0$), add these revenues, with interest, to m'_{t+1}; if in stage 2 the principal receives negative revenues ($d_t f(X_t) - w_t < 0$), subtract these revenues from m'_t. As a result, both contracts give the expert the same total revenues in present discounted terms.

(1). Constraint (6) is redundant because, given that $w_t' = 0$ for all t, it is implied by constraint (3). Constraint (5) is redundant because constraint (2) for period $t+1$ implies

$$\delta V_{t+1}(\mathcal{C}') = V_t(\mathcal{C}') - c_t \geq \frac{\delta}{1-\delta} f(X_t) + B_t(\mathcal{C}'),$$

which in turn implies $V_t(\mathcal{C}') \geq \frac{1}{1-\delta} f(X_t) + (1+r)B_{t-1}(\mathcal{C}') - m_t$, as desired. ■

Proof of Remark 1. That $X_1 > 0$ and $X_2 - X_1 > 0$ follows from the fact that both inequalities are necessary for the expert to make positive profits. We now show that $(1+r)B_1(\mathcal{C}) = P(X_2, X_1)$. If $(1+r)B_1(\mathcal{C}) > P(X_2, X_1)$, the expert can relax $(IC+L)$ by reducing X_1 by a small amount such that $(1+r)B_1(\mathcal{C})$ remains above $P(X_2, X_1)$; if instead $(1+r)B_1(\mathcal{C}) < P(X_2, X_1)$, the expert can relax $(IC+L)$ by increasing X_1 by a small amount such that $(1+r)B_1(\mathcal{C})$ remains below $P(X_2, X_1)$. Finally, that X_1 is transferred for free follows noting that the only positive money transfer from expert to novice, m_2, is fully devoted to purchasing knowledge $X_2 - X_1$ rather than being used to pay for X_1. ■

Proof of Lemma 2. Suppose toward a contradiction that contract $\mathcal{C} = \langle X_t, m_t, c_t \rangle_{t=1}^{\infty}$ solves problem (II) and yet has an infinite training phase – i.e. $X_t < X_{\sup}(\mathcal{C})$ for all t. The novice's incentive constraints (2) require that, for all t,

$$\frac{1}{1-\delta} f\left(X_{t-1}\right) + (1+r)B_{t-1}(\mathcal{C}) \leq V_t(\mathcal{C}) = (1+r)B_{t-1}(\mathcal{C}) + \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[f(X_\tau) - m_\tau\right]$$

and therefore

$$\sum_{\tau=t}^{\infty} \delta^{\tau-t} m_\tau \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left[f(X_\tau) - f\left(X_{t-1}\right)\right]. \tag{12}$$

Now select a distant enough period k such that $(1+r)f\left(X_{k-1}\right) \geq P\left(X_{\sup}(\mathcal{C}), X_{k-1}\right)$ (which exists because X_{k-1} must converge to $X_{\sup}(\mathcal{C})$ as k grows) and consider a new contract $\mathcal{C}' = \langle X_t', m_t', c_t' \rangle_{t=1}^{\infty}$ that is identical to $\mathcal{C}$ except for the following variables (in words, $\mathcal{C}'$ asks the novice to save all period $k-1$ output and use these savings to purchase at once, in period k, all knowledge $X_{\sup}(\mathcal{C}) - X_{k-1}$ that was eventually to be transferred):

1. $X_t' = X_{\sup}(\mathcal{C})$ for all $t \geq k$.
2. $c_k' = 0$ and $c_t' = f\left(X_{\sup}(\mathcal{C})\right) + (1-\delta)\left[(1+r) f\left(X_{k-1}\right) - m_k'\right]$ for all $t > k$ (i.e. $\mathcal{C}'$ fully smooths consumption from period k onward).
3. $m_k' = \sum_{\tau=k}^{\infty} \delta^{\tau-k} m_\tau + \sum_{\tau=k}^{\infty} \delta^{\tau-k} \left[f\left(X_{\sup}(\mathcal{C})\right) - f(X_\tau)\right]$ (i.e. in period k, the novice

is asked to pay a sum equal to all future discounted transfers in the original contract, plus the present discounted value of the additional knowledge he receives in the new contract) and $m'_t = 0$ for all $t > k$.

We now show that $\mathcal{C}'$ meets all constraints. On the one hand, inequality (12) implies that

$$m'_k \leq \sum_{\tau=k}^{\infty} \delta^{\tau-k} \left[f(X_\tau) - f(X_{t-1}) \right] + \sum_{\tau=k}^{\infty} \delta^{\tau-k} \left[f(X_{\sup}(\mathcal{C})) - f(X_\tau) \right] = P\left(X_{\sup}(\mathcal{C}), X_{k-1}\right),$$

which in turn implies that $m'_k \leq (1+r) f(X_{k-1})$ and therefore all liquidity constraints (3) are met.

On the other hand, from the definition of m'_k we obtain $-m'_k + \sum_{\tau=k}^{\infty} \delta^{\tau-k} f(X_{\sup}(\mathcal{C})) = \sum_{\tau=k}^{\infty} \delta^{\tau-k} \left[-m_\tau + f(X_\tau) \right]$, which implies that

$$V_k(\mathcal{C}') = (1+r) \underbrace{\left[B_{k-2}(\mathcal{C}) + f(X_{k-1}) \right]}_{B_{k-1}(\mathcal{C}')} - m'_k + \sum_{\tau=k}^{\infty} \delta^{\tau-k} f(X_{\sup}(\mathcal{C})) = V_k(\mathcal{C}) + (1+r) c_{k-1} \geq V_k(\mathcal{C}),$$

and that, for all $t < k$,

$$V_t(\mathcal{C}') = (1+r) B_{t-1}(\mathcal{C}) + \sum_{\tau=t}^{k-1} \delta^{\tau-t} \left[f(X_\tau) - m_\tau \right] + \delta^{k-t} \left[\sum_{\tau=k}^{\infty} \delta^{\tau-k} f(X_{\sup}(\mathcal{C})) - m'_k \right] = V_t(\mathcal{C}),$$

Consequently, the incentive constraints (2) are met for all $t \leq k$. Moreover, since there are no transactions after period k, all remaining incentive constraints are met as well.

Note, finally, that $\mathcal{C}'$ delivers a strictly higher payoff than $\mathcal{C}$ (i.e. $\Pi_1(\mathcal{C}') = \sum_{\tau=1}^{k-1} \delta^{\tau-1} m_\tau + m'_k < \sum_{\tau=1}^{k-1} \delta^{\tau-1} m_\tau + \sum_{\tau=t}^{\infty} \delta^{\tau-t} m_\tau = \Pi_1(\mathcal{C}')$), a contradiction. ■

Proof of Remark 2. Let $\mathcal{C} = \langle X_t, m_t, c_t \rangle_{t=1}^{\infty} \in \Omega$ be an arbitrary finite contract, with final training period T, satisfying constraints (2) and (3). Select a bonus M, to be paid in period T, such that

$$M = \sum_{t=1}^{T-1} (1+r)^{T-t} \left[f(X_t) - m_t \right] - \sum_{t=T}^{\infty} \delta^{t-T} m_t$$

(where the R.H.S. equals all output produced before period T net of all original money transfers, expressed in period T dollars) and let contract $\mathcal{C}' = \langle X_t, m'_t, c'_t \rangle_{t=1}^{\infty}$ satisfy:

1. $c'_t = 0$ for all $t < T$ and $c'_t = f(X_T) + (1-\delta) M$ for all $t \geq T$ (i.e. $\mathcal{C}'$ fully smooths consumption from period k onward).

2. $m'_t = (1+r) f(X_{t-1})$ for all $t = 2, ..., T-1$, $m'_t = (1+r) f(X_{t-1}) - M$ for $t = T$, and $m'_t = 0$ for all other t.

Notice that $\mathcal{C}'$ is a delayed-reward contract (and $\mathcal{C}' \in \Omega$). Moreover, from the definitions of m'_t and M,

$$\sum_{t=1}^{\infty} \delta^{t-1} m'_t = \sum_{t=1}^{T-1} \delta^{t-1} f(X_t) - \delta^{T-1} M = \sum_{t=1}^{\infty} \delta^{t-1} m_t,$$

and so both contracts prescribe the same total money transfers in present value.

That $\mathcal{C}$ and $\mathcal{C}'$ deliver the same payoffs Π_1, V_1 follows from the above observation combined with the fact that both contracts prescribe the same output levels. That $\mathcal{C}'$ satisfies the liquidity constraints (3) follows directly from the definitions of c'_t and m'_t. Finally, that $\mathcal{C}'$ satisfies the incentive constraints (2) follows from the fact that $V_t(\mathcal{C}') \geq V_t(\mathcal{C})$ for all $t \leq T$ (which in turn follows from the fact that $c'_t = 0$ for all $t < T$) and the fact that all incentive constraints after period T are trivially met. ■

Proof of Lemma 3. Being a delayed-reward contract, $\mathcal{C}$ has the property that, for all $t \leq T$,

$$V_t(\mathcal{C}) = \delta^{T-t} V_T(\mathcal{C}) = \delta^{T-t} \left[\frac{1}{1-\delta} f(X_T) + M \right]. \tag{13}$$

We now argue that, being a solution to (II), $\mathcal{C}$ has the property that in every period $t > 1$ in which knowledge is transferred ($X_{t-1} < X_t$), the novice's incentive constraint holds with equality:

$$\frac{1}{1-\delta} f(X_{t-1}) + (1+r) \underbrace{f(X_{t-1})}_{B_{t-1}(\mathcal{C})} = \delta^{T-t} V_T(\mathcal{C}). \tag{14}$$

If instead the period-t constraint was met with strict inequality, the expert could raise X_{t-1} by a small amount while still satisfying the constraint (together with the requirement that $X_{t-1} \leq X_t$), allowing the novice to produce additional output in period $t-1$. Since $m_t = (1+r) f(X_{t-1})$, this additional output would raise the expert's revenue during t, contradicting the optimality of $\mathcal{C}$. (Notice that this change leaves all remaining incentive constraints unaffected because the novice's continuation values $V_t(\mathcal{C})$ depend only on X_T

and M.)

We are now ready to show that, for any given $t > 1$, inequality (8) holds:

First, assume that $X_{t-1} = X_t$. It follows that the R.H.S. of (8) is $P(X_t, X_{t-1}) = 0$ and the L.H.S. is either $m_t + M \cdot \mathbf{1}\{t = T\} = (1+r)f(X_{t-1}) \geq 0$ (if $t \leq T$) or $m_t = 0$ (if $t > T$), as desired.

Second, assume that $X_{t-1} < X_t$ and further assume, toward a contradiction, that inequality (8) fails. Suppose initially that $t < T$. Combining the incentive constraint for period t with eq. (13) we obtain $\frac{1}{1-\delta} f(X_t) \leq \delta^{T-t} V_T(\mathcal{C})$. In addition, since (8) fails we must have $(1+r) f(X_{t-1}) < f(X_t)$. From these two inequalities we obtain

$$\frac{1}{\delta(1-\delta)} f(X_{t-1}) < \frac{1}{(1-\delta)} f(X_t) \leq \delta^{T-t} V_T(\mathcal{C}),$$

which contradicts eq. (14). Suppose finally that $t = T$. Since $m_T = (1+r) f(X_{T-1}) - M$, inequality $m_T + M < P(X_T, X_{T-1})$ is equivalent to

$$\frac{1}{\delta(1-\delta)} f(X_{T-1}) < \frac{1}{1-\delta} f(X_T) \leq V_T(\mathcal{C}),$$

which again contradicts eq. (14). ■

Proof of Proposition 1. Let $\mathcal{D}$ be the set of delayed-reward contracts. We proceed in two steps. First, we show that every optimal contract in $\mathcal{D}$ has the features in the proposition. Second, we show that no contract outside $\mathcal{D}$ is optimal.

Step 1. Consider an arbitrary contract $\mathcal{C} \in \mathcal{D}$ (with bonus M and final period T) that solves the relaxed problem (II). From the proof of Lemma 3, eq. (14), in every period $t+1$ such that $X_t < X_{t+1}$ we have

$$f(X_t) = (1-\delta)\delta^{T-t} V_T(\mathcal{C}). \tag{15}$$

We now claim that

$$f(X_t) = \min\left\{(1-\delta)\delta^{T-t} V_T(\mathcal{C}), f(X_T)\right\} \text{ for all } t < T. \tag{16}$$

We prove this claim by induction. That the claim is true for $t = T-1$ follows from noting that the novice's incentive constraint for period $T-1$ requires that $f(X_{T-1}) \leq (1-\delta)\delta V_T(\mathcal{C})$, and whenever $f(X_{T-1}) < f(X_T)$ this constraint holds with equality (eq.

(15)). Now suppose the claim holds for $t = T - n$ ($n \in \mathbb{N}$). That the claim holds for $t = T - (n+1)$ follows from noting that the novice's incentive constraint for period t requires that $f(X_t) \le (1-\delta)\delta^{T-t}V_T(\mathcal{C})$, and whenever $f(X_t) < f(X_T)$ we also have $f(X_t) < f(X_{t+1})$ (since, by hypothesis, $f(X_{t+1}) = \min\left\{(1-\delta)\delta^{T-(t+1)}V_T(\mathcal{C}), f(X_T)\right\}$) and therefore the constraint must again hold with equality (eq. (15)).

From (16), and the fact that $\mathcal{C} \in D$, we obtain

$$\Pi_1(\mathcal{C}) = \sum_{t=1}^{T-1} \underbrace{\delta^{t-1} \min\left\{(1-\delta)\delta^{T-t}V_T(\mathcal{C}), f(X_T)\right\}}_{\delta^{t-1}f(X_t)} - \delta^{T-1}M, \tag{17}$$

where $V_T(\mathcal{C}) = \frac{1}{1-\delta} f(X_T) + M$. The optimality of $\mathcal{C}$ requires that the above payoff is maximized w.r.t. X_T and M, subject to $X_T \le 1$ and $M \ge 0$. It follows that $X_T = 1$.

We now observe that if $M > (1+r)f(1)$ there exists an alternative contract $\mathcal{C}' = \langle X_t, m_t', c_t \rangle_{t=1}^{\infty}$ in $\mathcal{D}$ that also solves problem (II) and that prescribes, in period $T-1$, a bonus $M' = \delta\left[M - (1+r)f(1)\right]$. Contract $\mathcal{C}'$ is identical to $\mathcal{C}$ except for the lower bonus and the fact that it ends one period sooner: $m_T' = 0$ and $m_{T-1}' = (1+r)f(X_{t-1}) - M'$. Notice also that $\delta^{T-t}V_T(\mathcal{C}) = \delta^{T-1-t}V_{T-1}(\mathcal{C}')$ and that $M > (1+r)f(1)$ implies $f(X_{T-1}) = f(1)$. Therefore, after manipulation of (17), we obtain:

$$\begin{aligned} \Pi_1(\mathcal{C}) &= \sum_{t=1}^{T-1} \delta^{t-1} \min\left\{(1-\delta)\delta^{T-t}V_T(\mathcal{C}), f(1)\right\} - \delta^{T-1}M \\ &= \sum_{t=1}^{T-2} \delta^{t-1} \min\left\{(1-\delta)\delta^{T-1-t}V_{T-1}, f(1)\right\} - \delta^{T-2}M' = \Pi_1(\mathcal{C}'). \end{aligned}$$

In addition, we observe that $M \le (1+r)f(1)$ if and only if $m_T \ge 0$, which follows from combining the definition of m_T (namely, $m_T = (1+r)f(X_{T-1}) - M$) with the fact that $f(X_{T-1}) = \min\{(1-\delta)\delta V_T(\mathcal{C}), f(1)\}$.

Consider now the expert's original problem (I) in which the expert's incentive constraints (1) are included. For contracts in $\mathcal{D}$ this constraint imposes the single additional restriction that $m_T \ge 0$ (otherwise the expert would walk away before paying $-m_T$), which, as noted, is equivalent $M \le (1+r)f(1)$. It follows from the first observation in the previous paragraph that problems (I) and (II) have the same value (as M can be repeatedly reduced without lowering the expert's payoff). In addition, from eq. (17), any contract $\mathcal{C} \in \mathcal{D}$ that solves problem (I) (for which $M \le (1+r)f(1)$ and so $(1-\delta)\delta^{T-t}V_T(\mathcal{C}) \le f(1)$

for all $t < T$) delivers payoff

$$\Pi_1(\mathcal{C}) = (T-1)\delta^{T-1}\left[f(1) + (1-\delta)M\right] - \delta^{T-1}M.$$

Since this payoff is linear in M, and δ takes a generic value (namely, $(1-\delta)(T-1) \neq 1$), it follows that M is either equal to zero (its lowest possible value) or equal to $(1+r)f(1)$ (its highest possible one).

Finally, suppose contract $\mathcal{C} \in \mathcal{D}$ solves problem (I). If $M = 0$, the prices stated in the proposition are obtained by noting, from eq. (16), that $f(X_t) = \delta f(X_{t+1})$, and therefore $m_{t+1} = (1+r)f(X_t) = P(X_{t+1}, X_t)$, for all $t < T$. If instead $M = (1+r)f(1)$, in the last period of $\mathcal{C}$ no training takes place (as $f(X_{T-1}) = (1-\delta)\delta V_T(\mathcal{C}) = f(1)$) and the net money transfer is zero (as $m_T = (1+r)f(1) - M = 0$). As a result, this contract is identical to a contract with $M = 0$, the same output and money transfers for $t < T$, and that ends in period $T-1$.

Step 2. Consider a contract $\mathcal{C} \notin \mathcal{D}$ and assume toward a contradiction that $\mathcal{C}$ solves the original problem (I). From the proof of Remark 2, there is a contract $\mathcal{C}' \in \mathcal{D}$ that solves the relaxed problem (II) (as problems (I) and (II) have the same value) and that prescribes the same output levels as $\mathcal{C}$. Given the output levels prescribed by these two contracts, during the training phase $\mathcal{C}'$ simultaneously maximizes each of the novice's continuation payoffs V_t (because consumption is delayed) and minimizes each of the novice's cash balances $(1+r)B_{t-1}$ (because no savings beyond a single period are allowed). Moreover, since $\mathcal{C}$ does not share this feature, it follows that under $\mathcal{C}'$ the novice's incentive constraint is slack during at least one period $t \leq T$. Consequently, $\mathcal{C}'$ fails to satisfy eq. (16), a contradiction. ■

Proof of Proposition 2. By substituting for the value of $\mathcal{T}(X_1)$ in the objective of problem (10), and multiplying this objective by the constant $\frac{|\log\delta|}{f(1)}$, problem (10) simplfies to

$$\max_{X_1 \geq 0} \quad \log\left[\frac{f(1)}{f(X_1)}\right] \cdot \frac{f(X_1)}{f(1)}.$$

Since the average logarithm $\frac{\log(z)}{z}$ is uniquely maximized at $z = e$, the optimal gift X_1 uniquely satisfies $\frac{f(X_1)}{f(1)} = \frac{1}{e}$. ■

Proof of Corollary 1. Remark 2 and Lemma 3 remain valid under a minimum wage after two modifications.[46] First, a delayed-reward contract $\mathcal{C}$ with length T is now defined as

[46] It is also straightforward to verify that every optimal contract is finite.

one that, until graduation, grants the novice the smallest allowable consumption $w_{\min}$ each period. As a result, the expert obtains gross revenues equal to $(1+r)\left[f(X_{t-1})-w_{\min}\right]$, rather than $(1+r)f(X_{t-1})$, during each period $t=2,...,T$. Second, in such a delayed-reward contract, the novice's incentive constraint for period $t=2,...,T$ is $\frac{1}{1-\delta}f(X_{t-1})+(1+r)\left[f(X_{t-1})-w_{\min}\right]\leq V_t(\mathcal{C})$, where $V_T(\mathcal{C})=\frac{1}{1-\delta}f(X_T)+M$ and $V_t(\mathcal{C})$ is obtained recursively from

$$V_t(\mathcal{C})=w_{\min}+V_{t+1}(\mathcal{C})\text{ for }t<T.$$

As a result, equation (14) in the proof of Lemma 3 is replaced by $f(X_t)-w_{\min}=(1-\delta)\delta^{T-t}V_T(\mathcal{C})-\delta^{T-t}w_{\min}$.

Other than these modifications, the proof of the present Corollary is identical to the proof of Proposition 1, with equation (16) replaced by

$$f(X_t)-w_{\min}=\min\left\{(1-\delta)\delta^{T-t}V_T(\mathcal{C})-\delta^{T-t}w_{\min},f(X_T)-w_{\min}\right\}\text{ for all }t<T.$$

■

Proof of Proposition 3. After substituting for $\mathcal{T}(X_1,w_{\min})$ and multiplying the objective by $\frac{\ln(1+r)}{f(1)-w_{\min}}$, problem (11) simplifies to

$$\max_{X_1\geq 0}\ \ln\left[\frac{f(1)-w_{\min}}{f(X_1)-w_{\min}}\right]\cdot\frac{f(X_1)-w_{\min}}{f(1)-w_{\min}},$$

which is uniquely solved when $\frac{f(X_1)-w_{\min}}{f(1)-w_{\min}}=\frac{1}{e}$. Consequently, the length of the apprenticeship satisfies

$$\mathcal{T}(X_1,w_{\min})-1=\ln\left[\frac{f(1)-w_{\min}}{f(X_1)-w_{\min}}\right]\cdot\left[\ln(1+r)\right]^{-1}=\frac{1}{|\log\delta|}.$$

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

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

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

Proof of Proposition 4. For a fixed contract duration T, the expert's problem consists in finding consumption and output paths $(c_t,y_t)_{t=1}^{T-1}$ (from period T onward the novice produces and consumes $f(1)$). These paths solve:[47]

$$\max_{(c_t,y_t)_{t=1}^{T-1}} \quad \sum_{t=1}^{T-1} \delta^{t-1}\left[y_t - c_t\right] + \sum_{t=1}^{T-1} \lambda_t \left[V_t - \frac{1}{1-\delta} u\left(y_t\right)\right],$$

where λ_t is the Lagrange multiplier for the period t incentive constraint, $V_t = \sum_{\tau \geq t} \delta^{\tau - t} u\left(c_\tau\right)$, and the output path is constrained to be nondecreasing and weakly below $f(1)$. We begin by ignoring these latter two constraints. In the resulting relaxed problem, all incentive constraints bind (otherwise profits can be raised by raising y_t) and the first-order order conditions for c_t and y_t deliver, respectively: $\sum_{\tau=1}^{t} \lambda_\tau \delta^{t-\tau} u'\left(c_t\right) = \delta^{t-1}$ and $\lambda_t = \frac{(1-\delta)\delta^{t-1}}{u'(y_t)}$.[48] After combining these two equalities, and rearranging terms, we obtain $c_t = (1-\delta)^{\frac{1}{\sigma}} Y_t$. Moreover, since the Inada condition $u'(0) = \infty$ implies that $y_t > 0$ for all $t < T$, the resulting consumption path is increasing.

It remains to show that focusing on the above relaxed problem is appropriate. On the one hand, that $(y_t)_{t=1}^{T-1}$ is nondecreasing follows from the fact that $\frac{1}{1-\delta} u\left(y_t\right) = V_t$ and, since $(c_t)_{t=1}^{T-1}$ is increasing, so is $(V_t)_{t=1}^{T-1}$. On the other hand, that $(y_t)_{t=1}^{T-1}$ is weakly below $f(1)$ follows from the fact that setting $y_t = f(1)$ for some $t < T$ is equivalent to reducing the contract's duration T from the outset. ■

Proof of Corollary 2. Define a relaxed problem (IV) identical to (III) except that the expert's incentive constraints (1) are ignored. We begin by observing that the claims in Lemmas 2, 3 and Proposition 1 remain valid for every efficient contract (with problem (III) replacing problem (I) and problem (IV) replacing problem (II)). To see why, notice that the proof of each of these results is constructed by showing that any contract lacking the desired properties is Pareto-dominated by a contract (delivering a strictly higher payoff for the expert and a weakly higher payoff for the novice) that does satisfy them.

Once we restrict to contracts satisfying Proposition 1, the objective in (III), as a function of X_1, becomes

$$\lambda \underbrace{P(X_1, 0)}_{V_1} \;+\; (1-\lambda) \underbrace{\left[\mathcal{T}(X_1) - 1\right] f(X_1)}_{\Pi_1},$$

where $\mathcal{T}(X_1) - 1$ satisfies eq. (9). Ignoring the integer constraint for $\mathcal{T}(X_1)$, this objective

[47] We have assumed that the expert does not use a bonus ($M = 0$). The proof that doing so is optimal is omitted for brevity and is available upon request.

[48] It also follows from the strict concavity of u that the first-order conditions are sufficient for optimality of the two paths.

is uniquely maximized when $\frac{f(X_1)}{f(1)} = \frac{1}{e} \cdot \min\left\{\delta^{-\frac{\lambda}{(1-\lambda)(1-\delta)}}, e\right\}$. ■