The Foerder Institute for Economic Research
at Tel Aviv University

מכון למחקר כלכלי על שם
ד"ר ישעיהו פורדר
על יד אוניברסיטת תל אביב

The Eitan Berglas School of Economics

בית-הספר לכלכלה ע"ש איתן ברגלס

עמותה רשומה

What to Study and When: A Dynamic Roy Model of Specialization

Titan Alon, Daniel Fershtman

Discussion Paper No. 8-2023

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

מכון פורדר, אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409255 | פקס: 03-6405815
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409255 | Fax 972-3-6405815 | email: foerder@tauex.tau.ac.il

What to Study and When: A Dynamic Roy Model of Specialization*

Titan Alon† Daniel Fershtman‡

January 2023

Abstract

This paper generalizes the canonical model of human capital accumulation through schooling to endogenize the process of academic specialization. It provides the solution to a class of dynamic investment problems with switching and stopping under sequential uncertainty. Under mild assumptions, we show that the model's optimal policy has a particularly simple form that can be reduced to the comparison of independent indices. The optimal policy implies that schooling should begin with a period of general education, common to all students, followed by a period of gradual academic specialization before graduation. At the microeconomic level, it is consistent with the dynamics of student course taking observed in the data and the outcomes of educational interventions studied in the literature. At the macroeconomic level, its predictions are consistent with models of how education should adapt to changes in the speed and scope of technological change in labor markets.

Keywords: education, human capital, specialization, learning, optimal switching, optimal stopping, dynamic Roy model.

*We thank Ezra Oberfield, Chris Tonetti, Theodore Papageorgiou, Tomasso Porzio, Benjamin Moll, and seminar participants at Mannheim, NYU, Princeton, UCSD, VCU, Barcelona Summer Forum: Economic Fluctuations and Growth, ESEM Manchester, NASMES Seattle, and SITE Macroeconomics and Inequality workshop for helpful comments. The paper supersedes previous versions circulated under the titles "Gradual Specialization and the Allocation of Talent" and "A Theory of Human Capital Specialization". Fershtman gratefully acknowledges financial support from the Foerder Institute.

†Corresponding author. Department of Economics, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093. Email: talon@ucsd.edu

‡Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 6997801, Israel. Email: danielfer@tauex.tau.ac.il

1 Introduction

A nearly universal feature of modern education systems is that their curriculum becomes more specialized as education progresses. Schooling typically begins with a period of general education common to all, followed by a gradual process of academic specialization that differentiates students by field-of-study. While a common feature of many systems, the length of general education and the timing and scope of academic specialization varies considerably between countries and over time. For instance, today over 40% of upper secondary students in the European Union are enrolled in vocational degree programs, and students in countries like Germany and Austria are tracked starting from age 10. In the United States and Canada, tracking of students begins at age 16 and over 90% of upper secondary students are enrolled in general academic programs.[1] What should be the length and scope of general education? When should academic specialization begin, and how quickly and narrowly should it progress? The answers to these questions have important implications for the design of education systems and for the wider economy, as the allocation of students across distinct fields-of-study determines, in part, the future skill composition of the workforce.

To begin answering these questions, this paper extends the canonical model of human capital accumulation through schooling to endogenize the process of academic specialization. Agents dynamically allocate time to different fields-of-study before graduating at a time and with a specialization of their choosing to enter the labor market and begin work. Students do not initially know their true skill-specific abilities, but can learn their comparative advantage over time through academic experiences.[2] The dynamics of education are shaped by the interaction of this learning process and the human capital accumulation technology. In the absence of uncertainty, the model collapses to the Mincerian model of optimal schooling, nesting it as a special case.

The paper shows that, in this extended environment, optimal schooling begins with a period of general education followed by a gradual process of academic specialization before graduation. During general education, students exploit complementarities between skills through broad-based study, which is beneficial regardless of their ultimate field-of-study. During the specialization phase, students experiment to learn their underlying talents, progressively narrowing the set of skills they study in order to maximize the returns to schooling by focusing on their (perceived) comparative advantages.

[1] See Figures 3 and 4 and their accompanying source notes.

[2] The model also allows for *tastes* or other skill-specific attributes to be endogenously learned over time.

Methodologically, the paper presents the solution to a dynamic stochastic Roy (1951) model with learning from endogenous experience and human capital accumulation through schooling, as in Mincer (1974). The model belongs to a wider class of dynamic investment problems with optimal switching and stopping under sequential uncertainty. The technical challenge in solving this class of models arises from the combination of an *optimal stopping* problem of how long to invest, with an *optimal switching* problem of what to invest in and when.[3] An important insight of this paper is that these challenges can be overcome in canonical models of human capital by applying results on monotonic multidimensional stopping problems from Glazebrook (1979). Under mild assumptions, we show that the optimal policy takes a particularly simple form that can be reduced to the comparison of suitably defined indices, one for each field. Each index is independent of information about other fields and the optimal investment in each period is determined by the greatest prevailing index. Its tractability provides both analytical and computational advantages. In the context of schooling, the policy characterizes the *process* of academic specialization, including the optimal sequence of course taking, field switching, major choice, and educational attainment.

By endogenizing the process of academic specialization, this paper helps close the gap between economic models of optimal schooling and a burgeoning empirical literature on the dynamics of educational attainment and college major choice (Altonji, Arcidiacono, and Maurel 2016; Patnaik, Wiswall, and Zafar 2020). The tractability of the optimal policy – based on a comparison of skill specific indices – allows us to show that it generates behavior consistent with the new facts documented by this literature. For instance, every skill's index is associated with an *educational markup* that captures the option value of continuing to study that field. It reflects the value of both the human capital and new information that individuals expect to acquire. We show that the markups are always greater than one, reflecting the value of information; the marginal change in such markups declines as education progresses and uncertainty is resolved; and they respond asymmetrically to good and bad academic outcomes. As observed in the data, these dynamics imply that field switching should be concentrated in early stages of schooling, just after general education, and is driven by unexpectedly bad academic performance.

[3] Incorporating stopping decisions into models of optimal switching typically precludes the possibility of tractable solutions. This is the main feature distinguishing the model here from the classic multi-armed bandit problem. Our environment is a generalization of the latter in which, beyond "pulling" an arm, agents can additionally decide whether to "stop" or "continue" the evolution of each arm's state. The human-capital investment problem satisfies a key monotonicity property that allows us to tractably characterize the optimal investment policy.

The optimal policy is also consistent with the results of educational interventions and natural experiments studied in the literature. It correctly predicts how student course taking and major choice respond to exogenous curricular interventions. For example, as in the natural experiment analyzed by Fricke, Grogger, and Steinmayr (2018), forcing students to study any topic for a given period will increase the probability that they ultimately specialize in that field. As conjectured by the authors, the result depends on both the accumulation of field-specific human capital and new information received. As in the field study conducted by Conlon (2021), simply providing students with information that resolves uncertainty or corrects biases about field-specific payoffs can have large effects on subsequent course-taking and major choice. Moreover, the impact of these curricular and informational interventions is greater the earlier they are implemented, consistent with the experimental findings of Patterson, Pope, and Feudo (2019).

The model also makes predictions for how education systems should adapt in response to macroeconomic changes in the labor market. The literature emphasizes how the pace and direction of technological change can lead societies to optimally adopt more general or more specialized education systems (Krueger and Kumar 2004a, 2004b). To capture these changes in labor market demand, the final section extends the model to include complementarities between skills and shows that a suitably defined index policy is still optimal. The complementarities are modelled in reduced form using a skill-weights approach similar to Lazear (2009) and Cavounidis and Lang (2019), except here the weights are stochastic and use geometric aggregation. Consistent with the literature, the optimal policy with complementarities shows that the scope and duration of general education depends on prevailing labor market conditions. Skill-biased technical change which increases the demand for some skills relative to others leads to longer periods of general education, with more focus on the effected skills, and a delay in specialization. General education also increases after a rise in labor market risk, as students attempt to self insure against the uncertain composition of future skill demand.

In summary, the paper makes two contributions. Methodologically, it provides a useful characterization of the optimal policy in a class of dynamic investment problems with sequential uncertainty, stopping, and switching. The results hold under relatively mild assumptions on the investment and learning technologies, and so may be applicable in other economic settings. As an application, we show the model nests the classic Mincerian model of optimal schooling, generalizing the theory to accommodate recent empirical findings on the dynamics of academic specialization.

Related Literature. The investment problem studied here is closely related to those of Jovanovic (1979), Miller (1984), and Papageorgiou (2014) who study optimal switching that takes place on the job, rather than prior to entry into the labor market. As a result, the model here includes an endogenous stopping problem alongside endogenous switching decisions. Education dynamics under the optimal policy are driven by an exploitation versus experimentation trade-off similar to Perla and Tonetti (2014) and Lucas and Moll (2014). More broadly, the paper contributes to the literature developing models of multidimensional skill formation. Recent contributions cover a variety of economic contexts, including the sorting of workers (Lindenlaub 2017), the evolution of lifecycle earnings (Cavounidis and Lang 2019), business cycle volatility (Grigsby 2020), technology adoption (Adão, Beraja, and Pandalai-Nayar 2020), and more. This paper focuses in particular on formal schooling and the optimal design of curricula, with important implications for the productivity of education systems and their interaction with the real economy (Krueger and Kumar 2004a, 2004b; Martellini, Schoellman, and Sockin 2022).

The remainder of the paper is organized as follows. Section 2 lays out the model and presents the optimal policy. Section 3 characterizes learning dynamics under the optimal policy. Section 4 demonstrates the consistency of the optimal policy with experimental and empirical evidence from education interventions studied in the literature. Section 5 presents the extended model with skill complementarities, derives its solution, and discusses its implications. Section 6 concludes. All proofs are relegated to Appendix A.

2 The Benchmark Model

We study an economy in which production requires the execution of differentiated tasks, each requiring a particular skill. Faced with uncertainty about their true abilities, individuals sequentially decide how much time to spend studying each skill before entering the labor market at a time, and with a specialization, of their choosing.

Human Capital Accumulation. In each period $t = 0, ...\infty$, an agent's human capital consists of a vector $\mathbf{h}_t = (h_{1,t}, ..., h_{K,t})$ of skill-specific components. The agent sequentially chooses which skill to study. When studying a skill k, agents make a stochastic amount of progress $a_{k,t}$, drawn from distribution F_{θ_k} which depends on their unknown skill-specific ability θ_k. Consistent with the idea that higher ability students understand new material more quickly, we assume that F_{θ_k} is stochastically increasing in θ_k.

Progress in studying a skill k contributes to human capital accumulation according to the education technology,

$$h_{k,t+1} = H_k(h_{k,t}, a_{k,t}). \tag{1}$$

The amount of human capital accumulated within a given period is therefore a function of the agent's current level of human capital, as well as their underlying ability, which affects the amount of progress they make while studying. While the formulation imposes no parametric restrictions and allows the education technology to vary by skill, tractability requires that the human capital accumulation within a period be bounded so that there exists $\bar{H} > 0$ such that $H_k(h_{k,t}, a_{k,t}) - h_{k,t} \leq \bar{H}$ for all k and t. Beyond this, the only restriction we place on the human capital technology is that it satisfies a monotonicity condition summarized in Assumption 1.

Assumption 1. For any skill k in any period t, $H_k(h_{k,t}, a_{k,t}) \geq h_{k,t}$.

The monotonicity assumption arises naturally in the context of education technologies. It states simply that individual's cannot end up with less human capital after attempting to study a skill (though they may learn nothing, leaving human capital unchanged).

Endogenous Learning Dynamics. Since individuals cannot observe their true abilities, they must rely on their beliefs when making decisions about what to study or how much time to spend in school. Let $\mathbf{P}_t = (P_{1,t}, ..., P_{K,t})$ denote an agent's vector of period-t beliefs about their ability in each skill. After each period of studying, agents observe the progress $a_{k,t}$ they have made in skill k and update their skill-k period-t belief $P_{k,t}$ to $P_{k,t+1}$ according to the belief-updating rule Γ,

$$P_{k,t+1}(\theta_k | a_{k,t}) = \Gamma(P_{k,t}(\theta_k), a_{k,t}). \tag{2}$$

As with the human capital technology, the formulation in (2) is general and imposes no parametric or updating restrictions on the learning process. It allows for the possibility that agent's are myopic and do not learn from experience, so that $P_{k,t+1} = P_{k,0}$ for all k and t, as in the case of exogenous switching models. It also allows for the possibility of

Bayesian updating, in which case

$$\Gamma(P_{k,t}(\theta_k), a_{k,t}) = \frac{P_{k,t}(\theta_k) f_{\theta_k}(a_{k,t})}{\int P_{k,t}(x) f_x(a_{k,t}) dx},$$

where f_{θ_k} denotes the density of the distribution F_{θ_k}. It also accommodates non-Bayesian learning processes, time-inconsistent beliefs, and many other belief updating rules.[4] The only restriction we impose is that the learning process must be endogenous, so that agents learn their abilities only through experience. Assumption 2 formalizes this restriction on the learning process.

Assumption 2. If skill j is not studied in period t, then $P_{j,t+1} = P_{j,t}$.

The assumption captures a notion of learning-by-doing. It states that agents do not learn about their ability in skills that they do not use. For example, one cannot discover a talent for playing the piano without first investing time in practice. The realism of this assumption depends on the scope with which skills are defined in the model. When skills are broadly defined, the assumption of local learning appears appropriate. One is not likely to learn about their ability in mathematics while studying literature. On the other hand, if the skills being modelled are closely related–say applied math and physics–then the assumption may no longer be without loss of generality.[5] In Section 5 we discuss an extension of the model that is more suited for a definition of skills in which complementarities play a prominent role.

Educational Attainment and the Labor Market. At any point during their education, agents can choose to graduate from school and enter the labor market with a specialization (e.g. college major) of their choice. Entering the labor market after s years of schooling as a skill-k specialist yields lifetime utility,

$$\sum_{t=s}^{\infty} \delta^t U_k(w_k, h_{k,s})$$

[4]The model can also accommodate heterogeneous, non-stationary belief updating rules $\Gamma_{k,t}(\cdot)$, though we do not pursue this formulation here.

[5]For instance, one could learn about their ability to study math while studying physics. For instance, Stinebrickner and Stinebrickner (2014) shows that outside of Science, beliefs about a particular major may be shaped by experiences in other majors. Zafar (2011) provides suggestive evidence that learning is not entirely local and that beliefs updating in non-pursued majors can also respond to the arrival of new information. Our results continue to hold even in the presence of such spillovers, as long as they are small relative to the information that is learned directly about a skill.

where $\delta \in (0,1)$ denotes the discount factor and w_k is the wage rate for skill k, which individuals take as given. Utility function $U_k(w_k, h_{k,s})$ represents the expected *average* per-period payoff to being a skill-k specialist with human capital $h_{k,s}$. The dependence of the utility function on k allows for skill-specific preferences or amenities. Defining utility in this way also allows for the possibility that household income and human capital continues to evolve after graduation (e.g., due to on-the-job training).[6] We assume U_k is bounded for tractability and that it is strictly increasing in both arguments. The pay-off function can be extended to allow for additional uncertainty over skill-specific wages and preferences while in the labor market.

At each point in time, the state space of the agent's problem consists of $(\mathbf{h}_t, \mathbf{P}_t)$, which includes their K skill-specific human capitals and the K distributions encoding their prevailing beliefs. A solution to the individual's problem is a *policy* π which maps the state space $\mathbf{h}_t \times \mathbf{P}_t$ into decisions about what to study, when to graduate, and in what field. The optimal policy π^* is that which maximizes expected lifetime utility, formally

$$\mathbb{E}^{\pi}\left[\sum_{t=s}^{\infty} \delta^t \, U_k(w_k, h_{k,s}) \,|\, (h_{1t}, P_{1t}), \ldots, (h_{Kt}, P_{Kt})\right], \tag{3}$$

subject to the human capital technology in (1) and the belief-updating rule in (2), where the expectation operator $\mathbb{E}^{\pi}$ is taken with respect to the endogenous stochastic process induced by the policy π. The expectation operator represents the fact that agents make education decisions sequentially, and do not know precisely when they will graduate (s^*) or in what field-of-specialization (k^*) while proceeding through their education. The problem is difficult because these expectations depend on information which arises endogenously from the agents past decisions. Agents internalize that their decisions about what to study will affect both their accumulated human capital as well as the information they acquire about their underlying abilities. Consequently, the policy π that agent's pursue both shapes, and is shaped by, the skills they acquire and the information they receive while progressing through formal schooling.

[6] $U_k(w_k, h_{k,s})$ can also simply represent a constant flow utility payoff to the agent in each period after graduation. All of the results can be extended to allow payoffs, positive or negative, and even skill-specific, to accrue during study.

2.1 Optimal Human Capital Accumulation

This section characterizes the optimal policy π^* which maximizes household expected lifetime utility in (3). It describes the optimal sequence of course taking, educational attainment, and field of specialization for each individual. The key step in this characterization is to show that the complex decision problem defined over $(\mathbf{h}_t, \mathbf{P}_t)$ can be decomposed into K lower-dimensional problems, one for each skill, and a static choice between them. In particular, define an index v_k for each skill $(h_{k,t}, P_{k,t})$ given by

$$v_k(h_{k,t}, P_{k,t}) = \sup_{\lambda \geq 0} \mathbb{E}^{\lambda} \left(\frac{\delta^{\lambda}}{1-\delta} U_k(w_k, h_{k,t+\lambda}) \,|\, h_{k,t}, P_{k,t} \right), \tag{4}$$

where the supremum is over all realization dependent stopping times.[7] The index v_k measures the expected lifetime utility an individual would receive if they optimally invested their time studying and working with skill-k exclusively, *ignoring the presence of other skills*. In other words, it is the value function of a one-dimensional ($K = 1$) human capital accumulation problem with endogenous learning. In the absence of learning or stochastic human capital accumulation, the index corresponds to a discrete time version of the Mincer (1974) model of optimal schooling.

For each skill k, define the *graduation region* $\mathcal{G}_k$ as the set of skill-k states in which – ignoring the existence of all skills other than k – it is optimal to graduate immediately and enter the labor market as a skill-k specialist. Formally,

$$\mathcal{G}_k = \left\{ (h_k, P_k) | \arg\max_{\lambda \geq 0} \mathbb{E}^{\lambda} \left(\frac{\delta^{\lambda}}{1-\delta} U_k(w_k, h_{k,\lambda}) \,|\, h_k, P_k \right) = 0 \right\}. \tag{5}$$

The graduation region $\mathcal{G}_k$ admits a natural interpretation. It contains the set of all states $(h_{k,t}, P_{k,t})$ such that – if no other skills $j \neq k$ existed – the opportunity cost of further schooling would be greater than the expected value of continued education. As we

[7] Formally, denote by $\hat{a}_{k,1}, \hat{a}_{k,2}, \ldots$ the independent random variables with common distribution F_{θ_k}, representing the successive progress made studying skill k. Stopping times are then defined on $\{\hat{a}_{k,1}, \hat{a}_{k,2}, \ldots\}$, and the expectation is taken with respect to the stochastic process induced by λ. Note that allowing the stopping time to be infinite, a solution to the stopping problem is guaranteed to exist.

demonstrate in Section 3, whether or not a skill k is in its graduation region can typically be verified by a simple one-step look-ahead rule (1-SLA) comparing the expected lifetime utility of entering the labor market as a skill-k specialist today, $U_k(w_k, h_k)/(1-\delta)$, to the expected value, given current beliefs $P_{k,t}(\theta_k)$, of studying skill k for an additional period and then entering the labor market.

The following proposition defines the optimal policy π^* solving the agent's investment problem (3) subject to (1) and (2).

Proposition 1. *If Assumptions 1 and 2 hold, the agents' optimal policy π^* is as follows:*

1. *In each period t, select the skill $k^* \in arg\ max_{i\in\{1,\ldots,K\}}\ v_i$ with the highest index. If multiple skills have the highest index, select among them at random.*
2. *If $(h_{k^*}, P_{k^*}) \in \mathcal{G}_{k^*}$, enter the labor market as a k^* specialist. Otherwise, study k^* for an additional period and then return to step 1.*

Proposition 1 provides the solution to a class of sequential investment problems with heterogeneous skills and endogenous uncertainty. It shows that the agent's optimal policy can be reduced to a comparison among K independent indices, and verifying whether a skill is in its graduation region or not. The key property behind these features of the solution is that the optimal investment in any skill is independent of information pertaining to other skills. In other words, the optimal amount of time someone pursuing an economics degree should spend in school depends only on their ability in economics and is independent of their ability in other fields. This feature of the optimal policy constitutes a dynamic *Independence of Irrelevant Alternatives* (IIA). Specifically, a policy satisfies IIA if at any stage, conditional on not selecting skill i, the probability of studying any of the $K-1$ skills $k \neq i$ is independent of i's state.

The characterization of the optimal policy in Proposition 1 has both theoretical and computational advantages. First, as we show in the next section, it can yield tractable analytical characterizations of the optimal policy that can be computed directly or used for comparative statics and theoretical analysis. Second, even without analytical solutions, numerically solving the model using the characterization in Proposition 1 can substantially reduce computational time, especially when the number of skills K is large. The source of these computational savings is the decomposability of the optimal policy in the space of $(\mathbf{h}_t, \mathbf{P}_t)$ into K smaller problems, alleviating the curse of dimensionality.[8]

From a technical standpoint, the proof of Proposition 1 applies a result by Glazebrook (1979) which provides the solution to a multi-armed bandit problem in which players can additionally decide whether to "stop" or "continue" the evolution of each arm's state.[9] The necessary feature required for the optimal policy to satisfy the IIA property described above is that the continuation payoff upon stopping be non-decreasing, so that the option value of stopping weakly improves over time. This monotonicity property preserves the decomposability of the multi-armed bandit solution by ensuring that the irreversibility of stopping has no bite. An insight of this paper is that this property arises naturally in models of human capital accumulation and is guaranteed by Assumption 1 on the human capital technology.

The proposition also provides a link to the canonical models of human capital accumulation, which are nested as special cases. With a single skill and in the absence of uncertainty, the model collapses to the Mincer (1974) model of optimal schooling. In this case, the optimal policy in Proposition 1 says to continue with formal schooling until the opportunity cost of further study surpasses its its expected return, as in the Mincerian model. With multiple skills and no uncertainty, the model becomes a Roy (1951) model of college major choice. An interesting feature of the solution is that the indices agents use to rank each skill correspond to the lifetime utility they would expect to receive in a one-dimensional human capital problem where no other skills were available and switching never occurs. These differ from the value function of the agent's actual decision problem, since under the optimal policy agents may switch fields often and their sequence of course taking and years of schooling will depend on their entire vector of human capital and beliefs. The fact that agents can ignore these considerations in deciding what to study demonstrates how the IIA property above allows us to separate the vertical and horizontal dimensions of the human capital investment problem.

[8] For example, if we define the learning process Γ so that beliefs reside in a two-parameter family, the resulting dynamic programming problem would have $3 \times K$ state variables. In contrast, the optimal policy in Proposition 1 would have K 3-dimensional problems, which could substantially reduce computation time when K is large.

[9] To the best of our knowledge, this is the first application of the result in the Economics literature.

3 Learning Dynamics and Gradual Specialization

In certain cases, the characterization in Proposition 1 yields closed-form expressions for the optimal policy governing the evolution of investment and learning dynamics. To demonstrate the usefulness of these expressions, and to provide further insight into the model's mechanics, this section considers the classic Beta-Bernoulli learning model from the statistics literature which gives rise to closed-form solutions.

Following the literature, suppose agents make education investments to maximize their lifetime income, so $U_k(w_k, h_{k,t}) = w_k h_{k,t}$. Human capital accumulates according to the education technology given by

$$H(h_{k,t}, a_{k,t}) = h_{k,t} + a_{k,t}, \tag{6}$$

where the stochastic amount of progress made while studying, $a_{k,t}$, is Bernoulli distributed with success probability θ_k. That is,

$$a_{kt} = \begin{cases} \nu_k, & \text{w.p. } \theta_k \\ 0, & \text{w.p. } 1 - \theta_k \end{cases},$$

where $\nu_k > 0$. No progress is made in skills that are not studied in period t.

The education technology admits a simple interpretation. In each period an individual studies skill-k, they progress in understanding the new material they studied with probability θ_k. The new material learned increases their human capital in skill-k by ν_k. With probability $1 - \theta_k$, they fail to progress their understanding, leaving their human capital unchanged. An individual's skill-k ability is therefore summarized by the probability that they succeed in understanding new material, θ_k. The increment of human capital upon success, ν_k, proxies for the difficulty of acquiring human capital in different skills. While we assume ν_k is fixed for each skill, the results can also accommodate the case where it is stochastic.

Students do not know their true skill-specific abilities θ_k, but learn about what they may be through academic experience. Let an agent's initial beliefs about their ability in skill-k, $P_{k,0}$, be given by a Beta distribution $B(\alpha_{k,0}, \beta_{k,0})$. For simplicity, we assume belief updating follows Bayes' rule. As a result, posterior beliefs $P_{k,t+1}$ remain in the

Beta distribution family and evolve according to a simple process.[10] Specifically, given period t beliefs $B(\alpha_{k,t}, \beta_{k,t})$, beliefs in period $t+1$ will be

$$\Gamma(B(\alpha_{k,t}\,,\ \beta_{k,t}), a_{k,t}) = \begin{cases} B(\alpha_{k,t}+1, \beta_{k,t}), & \text{if } a_{k,t} = \nu_k \\ B(\alpha_{k,t}\,,\ \beta_{k,t}+1), & \text{if } a_{k,t} = 0 \end{cases}. \tag{7}$$

Initial human capital $\mathbf{h}_0$ is set to $h_{k,0} = \nu_k \alpha_{k,0}$, so that agents' starting skill level is consistent with their initial beliefs $P_{k,0}$ and the human capital technology. The initial period can therefore be thought of as the result of a pre-period of learning and human capital accumulation before time zero subject to the same learning and education technologies.

In each period t, the state $(h_{k,t}, P_{k,t})$ of every skill k is therefore entirely described by the triplet $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$. In line with the intuition above, a sufficient statistic for beliefs about one's own talent is a record of past academic success and failure in skill-k. The expected probability of success in skill-k at period t given current belief is given by

$$\mathbb{E}(\theta_k | P_{k,t}) = \frac{\alpha_{k,t}}{\alpha_{k,t} + \beta_{k,t}}. \tag{8}$$

Every time the individual succeeds in understanding new material, $\alpha_{k,t}$ increases by one; if they fail to understand new material, $\beta_{k,t}$ increases by one. At any period t, the individual's expectation of their own talent is simply their expected probability of success given past successes and failures. The more individuals fail when studying, the more pessimistic they become about their talent; the more times they succeed, the more optimistic they become.

It is straightforward to verify that the the human capital technology (6) and learning process (7) satisfy Assumptions 1 and 2. Consequently, Proposition 1 can be applied to characterize the agents' optimal policy.

Proposition 2 (Optimal Policy in the Beta-Bernoulli Model)**.** *In each period t, agents select the skill k^* with the highest index $v_k(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$, breaking ties according to any rule. Given state $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$, the index of skill-$k$ is given by*

$$v_k(h_{k,t}, \alpha_{k,t}, \beta_{k,t}) = \frac{w_k h_{k,t}}{1-\delta} \cdot \Omega(\alpha_{k,t}, \beta_{k,t}), \tag{9}$$

[10] This follows from the conjugacy property of the Beta and Bernoulli distributions.

where

$$\Omega_k(\alpha_{k,t}, \beta_{k,t}) \equiv \frac{\lceil \frac{\delta}{1-\delta} \rceil \delta^{\lceil \frac{\delta}{1-\delta} \rceil - (\alpha_{k,t}+\beta_{k,t})}}{\alpha_{k,t} + \beta_{k,t}}$$

during formal schooling and $\Omega_k(\alpha_{k,t}, \beta_{k,t}) = 1$ *at the moment of graduation and thereafter. Agents complete formal schooling and enter the labor market to work as a skill-*k^* *specialist if*

$$\alpha_{k,t} + \beta_{k,t} \geq \frac{\delta}{1-\delta}. \tag{10}$$

Otherwise, they study skill-k^* *for an additional period and the select amongst the skills again.*

In the proof (see Appendix A), we show that condition (10) describes the optimal years of schooling that individuals would choose in a one-dimensional model of human capital where only skill-k was available. It characterizes the set of states of skill k for which the opportunity cost of further schooling is greater than the expected value of continued education. In particular, we show that (10) is equivalent to

$$w_k h_{k,t} \geq \frac{\delta}{1-\delta} \mathbb{E}\left[w_k h_{k,t+1} | (h_{k,t}, \alpha_{k,t}, \beta_{k,t})\right], \tag{11}$$

where the left hand side is the opportunity cost of schooling and the right hand side is the expected value of further study. The equivalence shows how, in the presence of only one skill ($K = 1$), the model collapses to a stochastic version of the Mincerian framework and the optimal policy predicts the same optimal years of schooling as Mincer (1974). Importantly, however, condition (10) *does not* describe the optimal years of schooling when there are multiple skills ($K > 1$). This is because in the multi-dimensional model of human capital individuals can switch between skills and therefore have additional options to studying k or entering the labor market as a k-specialist.

3.1 The Dynamics of Academic Specialization

The indices in (9), which serve as the basis for comparison between skills, drive the dynamics of individual course taking, field switching, major choice, and ultimate years of schooling under the optimal policy. After graduating and entering the labor market, agents value each skill according to the net present value of earnings it can generate, $w_k h_{k,t}/(1-\delta)$. While in school, the value students assign to each skill is greater than the market net present value of their prevailing human capital. In particular, its market value is multiplied by a skill-specific *educational markup* $\Omega(\alpha_{k,t}, \beta_{k,t})$ reflecting the expected value of continued study, which includes the value of new information about one's ability. Proposition 3 describes how the markups evolve as education progresses.

Proposition 3 (The Educational Markups Ω)**.** *As long as a skill k has not entered its graduation region, its educational markup $\Omega_{k,t}$ satisfies the following properties:*

1. *$\Omega_{k,t}$ is strictly greater than 1 and decreases with every period the skill is studied.*
2. *The marginal change in Ω_k decreases with every period the skill is studied.*

The educational markup applied to any skill k being studied is always greater than one, except at the moment of entry into the labor market. The declining value of information about one's own abilities θ_k is reflected in the fact that $\Omega_{k,t}$ is declining in the number of study periods. The value of information about one's ability is maximal at early stages of education, when uncertainty is highest and the length of time to profitably exploit new information in subsequent human capital investments is the greatest. The marginal change in markups $\Omega_{k,t}$ also declines in the number of periods a skill is studied, reflecting in part the declining value of information as uncertainty is resolved. The fact that eventually Ω_{kt} does not change much indicates that in later stages of education it is prevailing human capital $\mathbf{h}_t$ that drives investment decisions, while at earlier stages of education information plays an important role.

The dynamics of $\Omega_{k,t}$ also embody predictions about when major switching occurs. Intuitively, students switch away from studying a skill after failing to make as much progress as they expect while studying. In the Beta-Bernoulli model, this occurs if agents become pessimistic about their ability after failing too many times in a row. However, since information about ability is more valuable early on, the effects of failure on major switching diminishes as education progresses. Formally, the marginal decrease in the skill's index as the result of failure decreases in the number of periods the skill is studied.

The prediction of the optimal policy that major switching is concentrated amongst poorly performing students at earlier stages of their education is consistent with a growing body of empirical evidence. The idea that these dynamics are driven by an endogenous learning process through which students learn their true abilities dates back at least to Schultz (1968)[11], while a more recent empirical literature provides direct evidence by documenting new empirical regularities in the dynamics of student academic histories. For instance, Zafar (2011) gathers panel data on the academic histories and subjective beliefs of Northwestern University undergrads and, comparing GPA to prior elicited beliefs, finds evidence that students revise beliefs about their ability in a manner consistent with learning from their academic performance. Using longitudinal data from the Berea Panel Study, Stinebrickner and Stinebrickner (2012, 2014) also provide evidence that students dynamically learn about their underlying abilities through grades. Similarly, structural models of student academic histories often recover estimated shock processes which suggest learning processes are important to explain the dynamics of student academic histories, major choice, and drop-out decisions (Arcidiacono 2004; Arcidiacono, Hotz, and Kang 2012; Wiswall and Zafar 2015). This paper shows that these empirical findings are indeed consistent with the optimal education choices of students in a generalized model of learning and human capital accumulation through schooling.

Numerical Illustration. Figure 1 illustrates the dynamics of academic specialization under the optimal policy. It demonstrates how the optimal allocation of study time evolves as education progresses. In particular, it reports the outcome of simulating 100,000 individual academic histories under the optimal policy in Proposition 2. To facilitate comparison across individuals with different chosen years of schooling, time allocations are reported by deciles of education completed. Within each individual, skills are ordered based on the total study time they received by the end of formal schooling. The model's time scale is calibrated to mimic post-secondary course taking at the Bachelors level; college (endogenously) lasts an average of four years, with each period representing one post-secondary course sequence. Three skills are available to study, each with the same population ability distribution F_θ. Each individual draws their (unknown) field-specific abilities from the population distributions at the start of schooling and has consistent rational beliefs, i.e. $P_{k,0} = F_{\theta_k}$. Further details on the calibration and computation are contained in Appendix B.

[11]In particular, Schultz (1968) observes that while the premise that part of the value of education is in the revelation of individual talents is a long-standing idea in economics, it is conspicuously absent from most modern theories of human capital and schooling. He describes this process of discovering one's talents as one of the "strongest features of U.S. higher education" while at the same time being the "much neglected activity" in economic research on education and human capital.

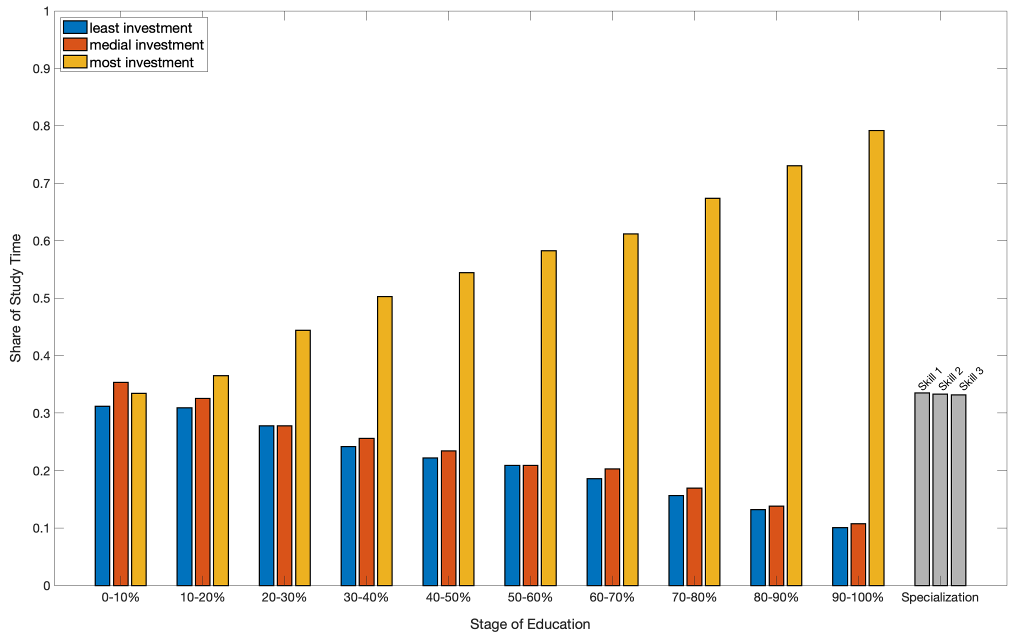

Figure 1: Gradual Specialization with Identical Skills. *Note: The panel shows the average distribution of human capital investment for different stages of education completion. Skills are ordered according to the amount of study-time they received, with skill #3 representing the skill an individual invested the most time in, and skill #1 receiving the least amount of time. Details of the underlying parametric assumptions are contained in Appendix B.*

The results in figure 1 illustrate the gradual process of academic specialization under the optimal policy. At early stages of education, individuals allocate their time relatively uniformly across all three skills; a period of schooling which resembles a general education curriculum. As individuals progress through schooling, they increasingly focus their study time on their intended field of specialization in anticipation of graduation and entering the labor market.

Since all skills are distributed ex-ante identically in the population, an equal number of individuals end up specializing in each field (the final grey columns). However, it is important to note that the cumulative allocation of each individual's investments in human capital are not uniform across the skills, nor do they reflect perfect specialization in one skill alone. In the simulation, individuals on average invest only 61% of their schooling time in their field of specialization; 39% of time is invested in skills that will never be used again after graduation. Despite never using these skills in the labor market, the investments are optimal because they help individuals learn their comparative advantage early on, which improves the efficiency of subsequent investments in human capital and ultimate major choice. These effects are most evident in the positive selection of student into fields for which they are better suited. For instance, the average ability of individuals who choose to specialize in each field, $\mathbb{E}[\theta_k \mid k = k^*]$, is 42% higher than the underlying average ability in the population, $\mathbb{E}[F_{\theta_k}]$. In other words, the endogenous learning process within schools improves the allocation of talent and can be an important component of the returns to education and aggregate labor productivity.

Noteworthy is the fact that the model endogenously produces an imperfect allocation of individuals across fields, even when everyone is behaving optimally. In the simulation of the optimal policy, only 56% of individuals selected into the field in which their true ability was highest. If all individuals selected into their best field, average ability would be 81% higher than the population average (rather than only 42%, based on equilibrium outcomes in the last paragraph). These imperfect outcomes reflect inherent uncertainties in knowing one's own true talents, and the costly trade-offs we face in attempting to discover them versus simply building on existing skills.

The investment behavior and field sorting which emerges under the model's optimal policy contrasts starkly with the predictions of other popular models of multidimensional skills. For instance, in Roy-type models without sequential uncertainty, the allocation of talent would be perfect, whereas in model's of exogenous switching, their would be no productivity premium associated with selection. Instead, the model here predicts an intermediate case where the allocation of talent lies between these extremes, and the extent to which an economy allocates individuals efficiently depends on the curricular structure of its education institutions.

4 Educational Interventions

Education policies such as curricular requirements, which limit an individual's ability to choose their own optimal sequence of course topics, can have important implications for the returns to education and the sorting of students across fields. The optimal policy provides a benchmark to analyze the extent to which curricular requirements and educational interventions may distort individual behavior. Its predictions are also consistent with a growing body of empirical and quasi-experimental evidence which show that even minor curricular requirements can have large, permanent effects on a student's subsequent course taking and major choice.

For example, Fricke, Grogger, and Steinmayr (2018) analyze a natural experiment at the University of St. Gallen in Switzerland and find that early exposure to a field can substantially alter the probability with which students major in that field. In particular, they find that students randomly assigned to write a research paper in economics were 50% more likely to choose economics as their major. The authors conclude that the lasting effect is likely the result of both the accumulation of area-specific human capital and information revealed to the student regarding their taste for and ability in economics.

The result that curricular requirements to study a skill k will increase the probability of specializing in that area is surprising. One would expect that negative realizations of required study could make students less likely to study the topic further in the future, offsetting any positive effects. Nevertheless, Proposition 4 shows that student behavior following the intervention is consistent with the model's optimal policy.

Proposition 4 (Educational Interventions I)**.** *A policy that forces an agent to study skill k in period 0 increases the probability with which the agent ends up graduating as a skill-k specialist.*

The proof follows from arguments similar to those in Gossner, Steiner, and Stewart (2021). To understand the intuition, consider the path of human capital investments for each possible sequence of outcome realizations, for every skill. Given a realization path, the agent's human-capital-accumulation strategy can be viewed as selecting in each period a skill for which to uncover one more step along the sequence. The choice of a specialization with which to enter the labor market can then be thought of as resulting from a contest among the skills: the agent continues to study until one of the skills gains sufficiently many positive realizations. The ultimate decision therefore boils down to which of the skills gains a sufficient amount of positive outcomes first. Studying any one of these skills accelerates the process of specialization for this skill while slowing it down for the other skills. Consequently, the likelihood that the target skill will ultimately win this contest increases.[12]

There is, however, one simplification in this intuition: forcing an agent to study a particular skill affects future decisions. It is possible that, for some accumulation strategies, such an educational intervention may lead to a path along which the agent studies the

[12]In other words, fixing any realization path, if in the absence of intervention the agent ultimately specializes in skill j given this realization path, then she will certainly do so when she is forced to study skill j initially. Such an intervention could also lead to specialization in j when, without it, the agent would have ultimately specialized in a different skill.

target skill less afterwards to an extent that might outweigh the direct effect of forcing the agent to study it. It is here that the IIA property again plays a crucial role. Together with the monotonicity of human capital accumulation in Assumption 1, IIA allows us to consider learning about the target skills seperately from learning about the alternatives.

While curricular requirements appear to play an important role, other studies have shown that simply providing students with information can have large effects on their subsequent course taking and major choice. Conlon (2021) conducts an experiment at a flagship state university and finds negative bias and substantial heterogeneity in beliefs about future major specific payoffs. Treating students with information about future payoffs of specific majors increased the probability of selecting that major by 16%.

Since the value of information changes as students progress through their education, the effect of interventions can decline the later they are implemented. Patterson, Pope, and Feudo (2019) document the importance of the timing of information revelation by exploiting a natural experience at the US Military Academy which randomly assigned students to take certain courses during (as opposed to after) the semester that they declared their major. The authors document big effects on major selection for students who took the course before major selection, rather than after. They conclude that the results are consistent with a learning process over uncertain ability. The same process emerges under the optimal policy, which leads the informational value of educational interventions to decline over time. Proposition 5 formalizes the result.

Proposition 5 (Educational Interventions II)**.** *Providing an agent with information about their ability in skill k has a diminishing marginal effect on their index v_k.*

Numerical Illustration. The parametric model in Section 3 can demonstrate how educational interventions influence academic specialization. To demonstrate the effect, we simulate the academic history of 100,000 college students choosing between two fields of study. For the purpose of illustration, we let the fields correspond loosely to STEM and non-STEM subject groups, and calibrate population ability distributions so that the average level of education attainment corresponds to four years of college and the optimal policy results in 18% of the population selecting into the STEM field, as in the data.[13] Mirroring experiments in the literature, we compute the effect of *curricular in-*

[13]Specifically, we match the share of total bachelor's degree awarded in STEM fields by postsecondary institutions, after adjusting for race and ethnicity differences. See `https://nces.ed.gov/programs/raceindicators/indicator_reg.asp`.

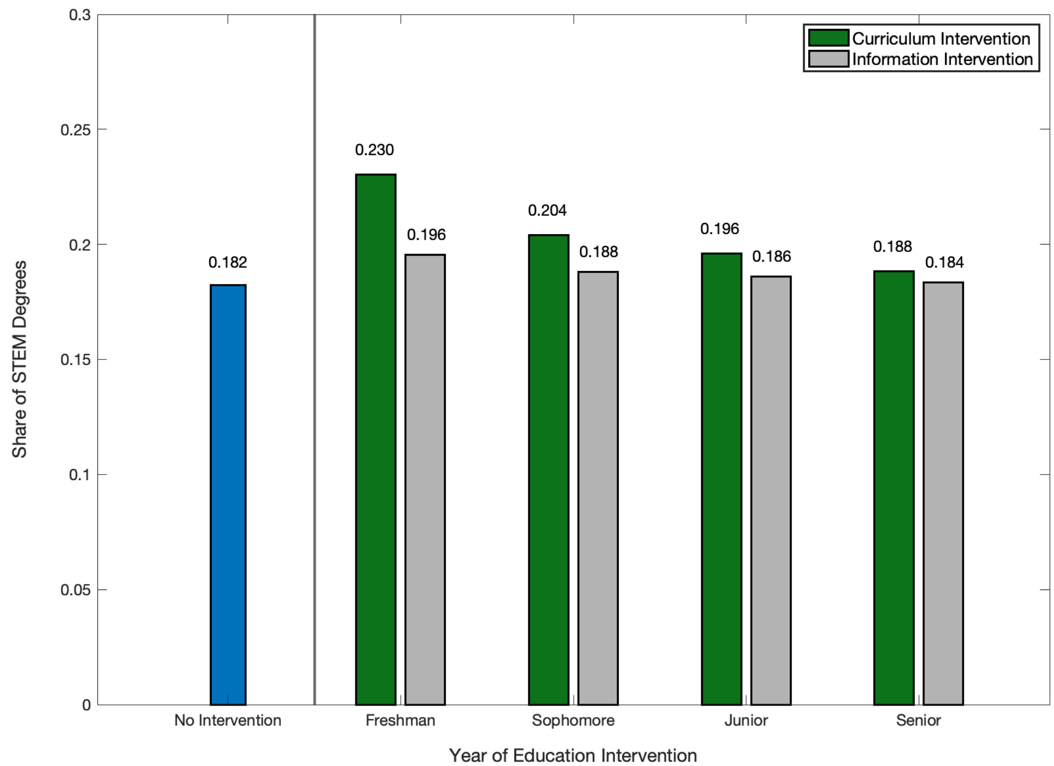

Figure 2: Educational Interventions. *Note: The figure displays the share of college graduates who choose a STEM major in the benchmark calibration, and after each educational intervention. Interventions occur at the start of Freshman, Sophomore, Junior, and Senior years. Curriculum interventions force students to take a course in STEM fields. Information interventions provide students with a signal about their abilities in STEM, but does not lead to the formation of new human capital. Additional details are contained in Appendix B.*

terventions which force students to take a STEM-field course at the beginning of their Freshman, Sophomore, Junior, and Senior years. To isolate the contribution of learning, we re-simulate the model focusing on *informational interventions* which provide students with a signal of their ability in STEM while holding constant their human capital.

One technical challenge of the simulation is that the index v_k will no longer have a closed-form solution after an information-only educational intervention. We therefore calculate the indices in these simulations using the recursive method in Sonin (2008).[14] See Appendix B for additional computational details.

Figure 2 plots the results and shows how the share of students choosing to major in STEM changes after each educational intervention. Consistent with findings in the lit-

[14]In general, a variety of methods exist to calculate the indices of the optimal policy, even when their analytical form is not known (Varaiya, Walrand, and Buyukkoc 1985; Katehakis and Veinott 1987).

erature, it shows that course requirements can have large effects on ultimate student major choice. For instance, the model predicts that requiring Freshman to take an additional STEM course at the beginning of college can increase the ultimate number of STEM majors by 4.8 percentage points, from 18.2% to 23.0%. The increase in STEM majors is the result of both accumulated skill-specific human capital and students learning about their ability in STEM fields. The simulation shows that even without accumulated human capital, the information revealed by the course requirements would increase specialization in STEM by 1.4 percentage points, from 18.2% to 19.6%. While the effect of the educational interventions are visible in each year, their impact on college major choice falls the later it occurs. Interventions at the beginning of Senior year, for instance, have only minimal impact on student major choice, with curricular interventions increasing STEM majors by 0.6 percentage points and information interventions increasing it by 0.2 percentage points. The declining effect reflects both the falling value of information and the fact that students acquire skill-specific human capital as they progress through school that endogenously raises their cost of switching fields later in their education.

5 A Model with Skill Complementarities

This section extends the model to allow for complementarity between skills. One reason for such complementarities may be knowledge hierarchies, whereby certain skills must be learned before others, for instance reading comprehension or basic numeracy. Another reason may be technological, in that every job may require using each skill with some positive probability. Complementarities may also reflect insurance motives, whereby workers invest in multiple skills to insure against obsolescence risk or job loss in frictional labor markets. This section shows that, despite the added complexity, the optimal policy can still be characterized by a suitably defined index rule.

In the presence of these complementarities, the optimal policy gives rise to periods of *purely general education*, where all individuals study the same mix of topics at the start of their education regardless of their ultimate field of specialization. The model without complementarities in the preceding sections gave rise to periods of exploration which could resemble, in terms of time allocation, periods of general education. The general education induced by complementarities precedes this period of exploration and has different economic motives. The particular mix of skills studied and the length of time

Figure 3: Age of First Educational Tracking

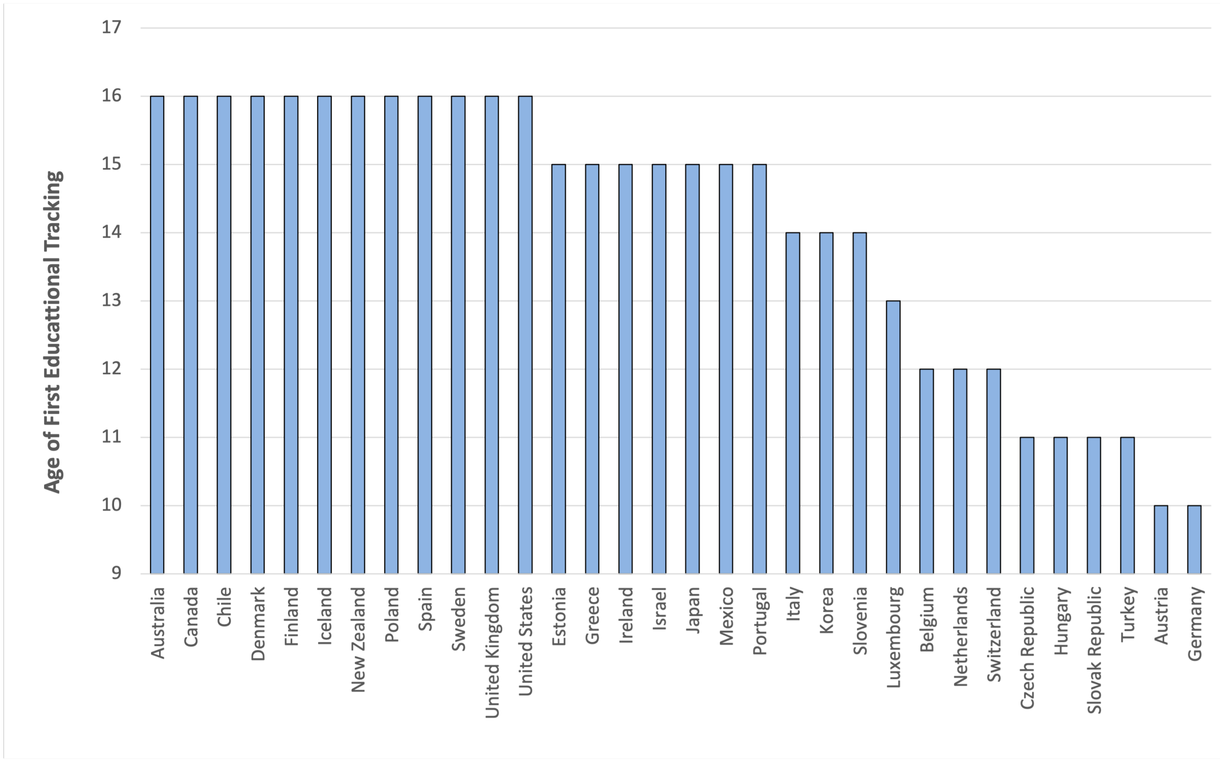

Notes: *Source data taken from the OECD's PISA 2012 Database.*

spent in general education depends on the strength of the complementarities between skills, and hence the technological and structural characteristics of labor markets.

Understanding these channels is important because there is considerable cross-country variation in both the timing and scope of academic specialization in formal schooling. Figure 3 illustrates this point by plotting cross-country differences in the age of first educational tracking, when students are separated and assigned to different curricular tracks. In countries like the United States and Canada, tracking begins at age 16, while in countries like Germany and Austria, students are put on distinct curricular tracks starting from age 10. Academic specialization is also often associated with separating students into academic and vocational tracks. Figure 4 shows that this notion of educational specialization also varies substantially across countries by plotting the share of upper secondary graduates who hold general education versus vocational degree. While vocational education is rare in the United States, accounting for less than 9% of graduates, it is very common in other countries like Austria and Germany, where the majority of students are vocationally specialized. Interestingly, there does not appear to

Figure 4: Education System Comparisons: Attainment and Curricular Specialization

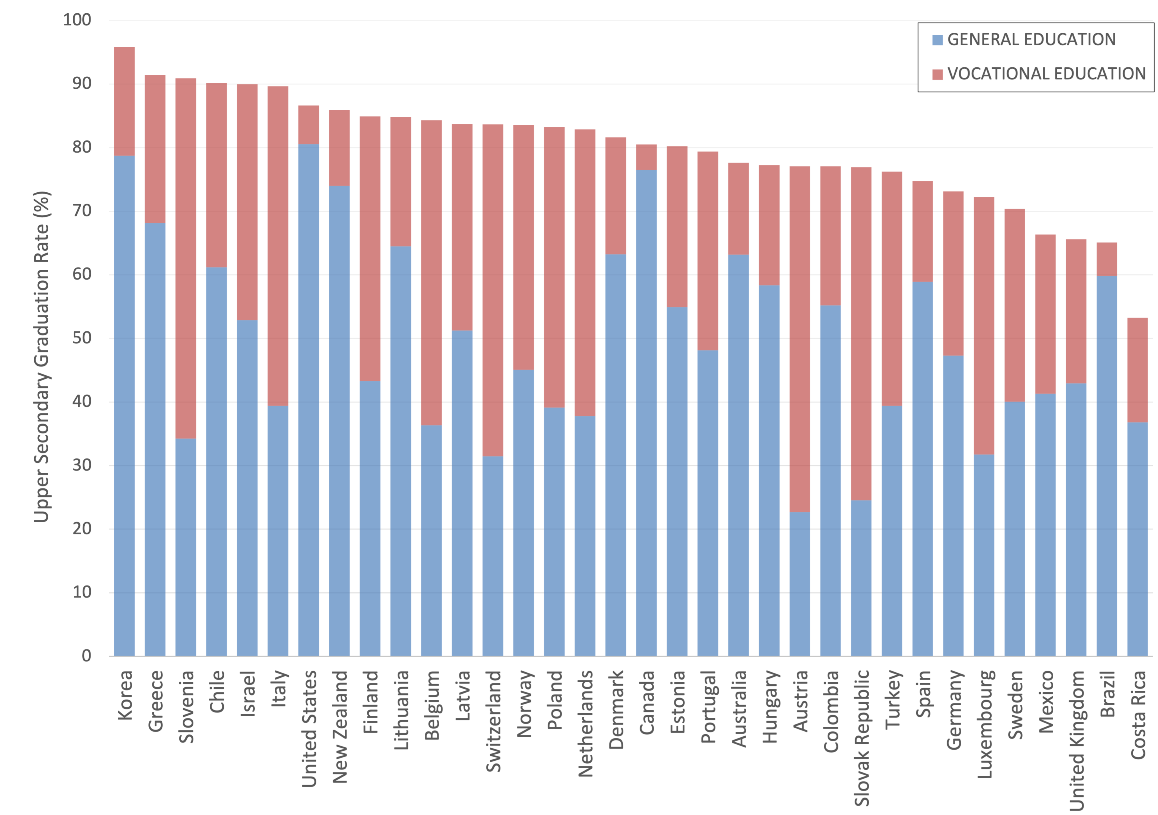

Notes: *Bar height reports upper secondary school graduation rates in 2019. Bar shading depicts upper secondary enrollment share by program type (general or vocational) for upper secondary students, ages 15-19, in 2014. Vocational shares include students in combined school and work-based programs. Canadian shares correspond to all ages 18-64. Underlying data come from OECD's Education at a Glance Database. U.S. vocational share is not included in the database and so is taken from estimates in Alon (2019).*

be a clear relationship between educational attainment and educational specialization in the cross-section of countries, as more or less specialized education systems are present in countries with both high and low levels of educational attainment.

To capture complementarities, we alter the payoff function in (3) so that it depends on the entire vector of human capital. Specifically, letting $\mathbf{h} = (h_1, \ldots, h_K)$, for each skill k the new payoff function $U_k(w_k, \mathbf{h})$ is given by

$$U_k(w_k, \mathbf{h}) = w_k \cdot \mathbb{E}\left[\prod_{j=1}^{K} h_j^{\eta_j}\right], \tag{12}$$

where $\boldsymbol{\eta} \equiv (\eta_k)_{k \in \{1,\ldots,K\}}$ are stochastic *skill weights*. We assume that $\boldsymbol{\eta}$ is independently distributed with each η_k drawn in each period from a distribution with cdf Ψ_k. The skill weights capture, in a reduced-form, the employment and technological uncertainty that may lead human capital outside one's specialization to matter in the labor market. In this sense, the model with complementarities resembles the skill-weights framework of Lazear (2009) except that here the weights are stochastic and take the form of a geometric average, rather than arithmetic average. Similarly, Cavounidis and Lang (2019) develop an earnings model where skill weights enter through a CES aggregator. Within fields, the specification is also consistent with the findings of Deming and Kahn (2018) who analyze job postings and document substantial variation in the mix of skills required in even narrowly defined jobs, as well as evidence of complementarities across skill groups, notably cognitive and social abilities.

5.1 General Education

General education emerges in the presence of complementarities because additional knowledge in some particular skill may be of such high *general value* that further study of this skill is expected to be beneficial independently of an agent's ultimate field of specialization. Formally,

Definition 1. *Studying skill k has* general value *if its state (h_k, P_k) satisfies*

$$\delta^\tau \mathbb{E}\left[h_{k,\tau}^{\eta_{k,\tau}}\right] \geq h_k^{\eta_k} \tag{13}$$

for some stopping time τ.

Condition (13) holds when a skill's impact on other skills is expected to increase at a rate that exceeds discounting, optimizing over the period of time it is studied. Real world examples include studying Mathematics to develop basic numeracy or English for literacy. For each skill k, denote by $\mathcal{E}_k$ the set of states of general value.[15]

Proposition 6 below shows that, even in the presence of complementarities between skills, the optimal policy can still be expressed as a simple comparison of appropriately

[15]Note that a skill may satisfy (13) for some states and violate it for others. Furthermore, each set $\mathcal{E}_k$ is only a function of information about skill k.

defined, independent indices. The horizontal decision (across skills) and vertical decision (study/work) of human capital investment are disentangled, as in Proposition 1, satisfying the IIA condition. To characterize the optimal policy in this extended environment with complementarities, we define the following indices v^C. For each skill k and state $(h_k, P_k) \notin \mathcal{E}_k$, let

$$v_k^C(h_{k,t}, P_{k,t}) = \sup_{\lambda,\tau} \left\{ \frac{\mathbb{E}^{\tau,\lambda}\left[\sum_{s=t+\lambda-1}^{\tau-1} \delta^s U_k(w_k, h_{k,s})|h_{k,t}, P_{k,t}\right]}{h_{k,t}^{\eta_{k,t}} - \mathbb{E}^{\tau,\lambda}\left[\delta^\tau h_{k,t+\tau}^{\eta_{k,t+\tau}}|h_{k,t}, P_{k,t}\right]} \right\}, \tag{14}$$

where, again, τ is a stopping time (possibly infinite), and λ is an additional stopping time specifying when entry into the labor market occurs. Note that for $(h_k, P_k) \notin \mathcal{E}_k$, the difference $h_{k,t}^{\eta_{k,t}} - \mathbb{E}^{\tau,\lambda}\left[\delta^\tau h_{k,t+\tau}^{\eta_{k,t+\tau}}|h_{k,t}, P_{k,t}\right]$ in the denominator is always positive. Furthermore, as in the original model, define skill k's graduation region $\mathcal{G}_k^C$ as the set of states such that – in the solution to the maximization problem in the definition of (14) – it is optimal to graduate formal schooling and enter the labor market. Given these definitions, the following proposition describes the agent's optimal policy:

Proposition 6. *In the environment with skill complementarity, the following procedure describes the agent's optimal policy:*

1. *At each period t, if there are any skills is a state with general value, the agent studies any one of them.*

2. *At each period t, if there are no skills in a state of general value, the agent selects the skill $k^* \in \arg\max_{i\in\{1,\ldots,K\}} v_i^C$ with the highest index. If multiple skills have the highest index, a skill is selected at random.*

3. *In each period t, if there are no skills in a state of general value and skill k^* is selected, then: If $(h_{k^*}, P_{k^*}) \in \mathcal{G}_{k^*}^C$, the agent enters the labor market as a k^* specialist; otherwise, she studies skill k^* for an additional period.*

Proposition 6 shows that a suitably defined index policy remains optimal in a model with complementarities.[16] The following dynamics emerge under the optimal policy:

[16] See Eliaz, Fershtman, and Frug (2021) for an analysis of scheduling problems with externalities in other contexts.

Initially, individuals invest solely in skills with general value (i.e., those satisfying Definition 1), during a period of time that may be viewed as the general education phase of schooling. During this period, human capital is accumulated in skills that enhance other skills to the extent that they overcome the opportunity cost of time and are beneficial regardless of eventual specialization. After general education ends, agents begin the process of academic specialization. Once again, a process of gradual specialization arises whereby individuals progressively narrow the range of skills they invest in as they learn their comparative advantage. The specialization process is even more gradual in the presence of complementarities since skills enhance one another, providing yet another incentive for broad investment across skills. Formal schooling ends endogenously when individuals choose a field of specialization and enter the labor market.

The optimal policy for educational investments described in Proposition 6 resembles the structure of many education systems observed in the real world. Most systems begin with a period of common general education before students become gradually differentiated and eventually specialize in narrower fields of expertise. The model shows that periods of general education, regardless of ultimate specialization, can emerge endogenously for very distinct reasons. Further periods of common education across students may also emerge because of broad-based academic exploration. While resembling a period of general education, this latter phase occurs for fundamentally different reasons than the general education induced by complementarities. An important implication of the model is that the duration and composition of general education, the speed and scope of academic specialization, and the optimal years of schooling all jointly depend on the characteristics of labor markets awaiting students after graduation.

5.2 Skill-Biased Technological Change and General Education

The optimal policy allows us to characterize how particular changes in the labor market affect the structure of education systems and their curricula. One particularly well-studied recent shift in the labor market has been skill-biased technical change, which has lead to an increase in demand for certain skills, such as cognitive and STEM-based, relative to others. The economic sources of these shifts in demand remain varied and interrelated, including the advent of new technological processes (Acemoglu 2002; Violante 2008), changes in the pattern of trade (Traiberman 2019), and structural transformation of the economy (Buera et al. 2021). A large ensuing literature examines the implications for employment dynamics, earnings, occupation choice, technology adoption, growth and more (Dvorkin and Monge-Naranjo 2019; Cavounidis and Lang 2019; Adão, Beraja, and Pandalai-Nayar 2020).

The model complements this literature by showing how changes in the relative demand of certain skills can induce changes in the education system and the process of academic specialization. Within the model, the effect of skill-biased technical change can be captured by a first-order stochastic (FOS) improvement in the skill weight distributions of a subset of skills. A FOS improvement in skill-j's skill weight means that human capital in j becomes (stochastically) more important in determining a worker's productivity at work. The following proposition summarizes the effect on the optimal provision of education.

Proposition 7 (General Education and Skill-Biased Technological Change)**.** *An improvement in demand for any subset of skills $\mathcal{J} \subseteq \{1, ..., K\}$ – represented by a first-order stochastic improvement in Ψ_j for each $j \in \mathcal{J}$ – leads to more general education and delays the start of academic specialization.*

The proposition shows that skill-biased technical change (SBTC) increases the optimal provision of general education pursued by all students. The expansion in general education focuses in particular on the skills which experience increased demand. The prediction is consistent with the observed expansion of general education, and delay in academic specialization, accompanying SBTC in the United States (Goldin and Katz 2010; Alon 2019). Moreover, as in Section 4, these curricular changes can alter the skill composition of the labor force by leading some students to change their specialization.

5.3 Labor Market Uncertainty and General Education

Technological change may also increase labor market risk by creating uncertainty about the future demand for certain skills. Advances in automation increase the risk of skill obsolescence as capital replaces labor in a growing variety of production tasks (Acemoglu and Restrepo 2022; Jones and Liu 2022). Recent studies of comprehensive online job databases provide empirical evidence of how technological innovation changes the composition of skills that firms demand (Deming and Noray 2020; Acemoglu et al. 2022).

As with skill-biased shocks to demand, the model provides predictions for how formal education should optimally respond to changes in labor market uncertainty. Representing an increase in skill-specific risk as a mean preserving spread (e.g. a second order stochastic shift) in its skill weights, the following proposition formalizes the result.

Proposition 8 (General Education and Labor Market Uncertainty)**.** *An increase in labor market risk – represented by a mean preserving spread of the distributions* $(\Psi_k)_{k\in\{1,\ldots,K\}}$ *– leads to more general education and delays the start of academic specialization.*

A mean preserving spread of the skill weight distributions increases the uncertainty individuals face in determining which skills they will use over the course of their careers. Greater variability in η_k makes agents more uncertain about the importance of human capital in skill-k for future earnings and productivity. Proposition 8 shows that agents respond by pursuing more general education to insure themselves against labor market risk. The process of academic specialization is delayed to later periods of formal schooling, when agents are also more certain of their underlying abilities.

The literature studying the link between education and labor market uncertainty has emphasized the trade off between more general and more specialized training. The prevailing belief is that more specialized training leads to higher earnings in the short-run, but more general education provides insurance against future obsolescence or other labor market risks. Hanushek et al. (2017) provide evidence consistent with this perspective, showing that gains in youth employment from specialized education may be offset by less adaptability to technological change later in life. Similarly, Krueger and Kumar (2004a, 2004b) show how more general education may have lead the United States to be more adaptable to the changing skill demands of rapid technological progress compared to Europe, where education curricula are more specialized.

The result in Proposition 8 is consistent with the literature in predicting that labor market uncertainty increases general education, though the underlying mechanisms are distinct. Moreover, while the literature typically focuses on the final allocation of workers across the two types of schooling (specialized v. general), the model here describes how the whole *process* of academic specialization adjusts. This includes the precise mix of skills constituting general education, the duration and timing of academic specialization, and how the curriculum interacts with the optimal years of schooling. In doing so, the model enables a more detailed comparison of education systems; provides a mapping to new microdata on the dynamics of students' academic histories; and, in conjunction with the results in Section 4, generates predictions for how educational interventions may be tailored to contemporary changes in the labor market. Future work should further develop these lines of inquiry, both theoretically and empirically.

6 Conclusion

This paper extends the canonical models of human capital accumulation through schooling to incorporate the process of academic specialization. Under mild assumptions, we show that the model's optimal policy has a particularly simple form that can be reduced to the comparison of independent indices. The optimal policy predicts that schooling should begin with a period of general education common to all students, following by a period of gradual academic specialization before graduation. At the microeconomic level, it is consistent with the dynamics of student course taking observed in the data and the outcomes of educational interventions studied by the literature. At the macroeconomic level, its predictions are consistent with models of how education should adapt to changes in the speed and scope of technological change in labor markets.

A potential avenue for future research is to further expand the realism and generality in models of schooling to capture additional aspects of education technologies and institutional arrangements. Such work would help connect existing theory to a proliferating body of research on the dynamics and heterogeneity in education markets. These advances are a necessary step to answering important questions regarding the optimal structure of education institutions and how they should adapt to economic conditions. The results here provide a modest step in this direction.

References

Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1): 7–72.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40 (S1): S293–S340.

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90 (5): 1973–2016.

Adão, Rodrigo, Martin Beraja, and Nitya Pandalai-Nayar. 2020. ""Technological Transitions with Skill Heterogeneity Across Generations"." Working paper 26625, National Bureau of Economic Research.

Alon, Titan. 2019. "Earning More by Doing Less: Human Capital Specialization and the College Wage Premium." *Working paper*.

Altonji, J.G., P. Arcidiacono, and A. Maurel. 2016. "Chapter 7 - The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects." In , edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, Volume 5 of *Handbook of the Economics of Education*, 305 – 396. Elsevier.

Arcidiacono, Peter. 2004. "Ability sorting and the returns to college major." *Journal of Econometrics* 121 (1-2): 343–375.

Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang. 2012. "Modeling college major choices using elicited measures of expectations and counterfactuals." *Journal of Econometrics* 166 (1): 3–16. Annals Issue on "Identification and Decisions", in Honor of Chuck Manski's 60th Birthday.

Buera, Francisco J, Joseph P Kaboski, Richard Rogerson, and Juan I Vizcaino. 2021. "Skill-Biased Structural Change." *The Review of Economic Studies* 89 (2): 592–625 (07).

Cavounidis, Costas, and Kevin Lang. 2019. ""Ben-Porath Meets Lazear: Microfoundations for Dynamic Skill Formation"." *Journal of Political Economy*.

Chow, Y. S., and Herbert Robbins. 1961. "A Martingale System Theorem and Applications." *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. Berkeley, Calif.: University of California Press, 93–104.

Conlon, John J. 2021. "Major Malfunction A Field Experiment Correcting Undergraduates' Beliefs about Salaries." *Journal of Human Resources* 56 (3): 922–939.

Deming, David, and Lisa B. Kahn. 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals." *Journal of Labor Economics* 36 (S1): S337–S369.

Deming, David J, and Kadeem Noray. 2020. "Earnings Dynamics, Changing Job Skills, and STEM Careers*." *The Quarterly Journal of Economics* 135 (4): 1965–2005 (06).

Dvorkin, Maximiliano, and Alexander Monge-Naranjo. 2019. "Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations." Working paper 2019-13, Federal Reserve Bank of St. Louis.

Eliaz, Kfir, Daniel Fershtman, and Alexander Frug. 2021. "On the Optimal Scheduling of Attention." CEPR Discussion Paper No. DP16364.

Fricke, Hans, Jeffrey Grogger, and Andreas Steinmayr. 2018. "Exposure to academic fields and college major choice." *Economics of Education Review* 64:199–213.

Glazebrook, Kevin D. 1979. ""Stoppable Families of Alternative Bandit Processes"." *Journal of Applied Probability* 16 (4): 843–854.

Goldin, Claudia, and Lawrence F. Katz. 2010. *The Race between Education and Technology*. Harvard university press.

Gossner, Olivier, Jakub Steiner, and Colin Stewart. 2021. "Attention please!" *Econometrica* 89 (4): 1717–1751.

Grigsby, John. 2020. ""Skill Heterogeneity and Aggregate Labor Market Dynamics"." Working paper, University of Chicago.

Hanushek, Eric, Guido Schwerdt, Ludger Woessmann, and Lei Zhang. 2017. "General Education, Vocational Education, and Labor-Market Outcomes over the Lifecycle." *Journal of Human Resources* 52 (1): 48–87.

Jones, Benjamin F, and Xiaojie Liu. 2022, September. "A Framework for Economic Growth with Capital-Embodied Technical Change." Working paper 30459, National Bureau of Economic Research.

Jovanovic, Boyan. 1979. ""Job Matching and the Theory of Turnover"." *Journal of Political Economy* 87 (5): 972–990.

Katehakis, Michael N, and Arthur F Veinott. 1987. ““The Multi-Armed Bandit Problem: Decomposition and Computation”.” *Mathematics of Operations Research* 12 (2): 262–268.

Krueger, Dirk, and Krishna B. Kumar. 2004a. “Skill-Specific Rather than General Education: A Reason for US-Europe Growth Differences?” *Journal of Economic Growth* 9 (2): 167–207.

———. 2004b. “US–Europe differences in technology-driven growth: quantifying the role of education.” *Journal of Monetary Economics* 51 (1): 161–190.

Lazear, Edward P. 2009. “Firm-Specific Human Capital: A Skill-Weights Approach.” *Journal of Political Economy* 117 (5): 914–940.

Lindenlaub, Ilse. 2017. “Sorting Multidimensional Types: Theory and Application.” *The Review of Economic Studies* 84 (2): 718–789 (01).

Lucas, Robert E., and Benjamin Moll. 2014. ““Knowledge Growth and the Allocation of Time”.” *Journal of Political Economy* 122 (1): 1–51.

Martellini, Paolo, Todd Schoellman, and Jason A. Sockin. 2022, March. “The Global Distribution of College Graduate Quality.” Working papers 791, Federal Reserve Bank of Minneapolis.

Miller, Robert A. 1984. “Job matching and occupational choice.” *Journal of Political economy* 92 (6): 1086–1120.

Mincer, Jacob. 1974. *Schooling, Experience, and Earnings*. National Bureau of Economic Research, Inc.

Papageorgiou, Theodore. 2014. ““Learning Your Comparative Advantages”.” *The Review of Economic Studies* 81 (3): 1263–1295.

Patnaik, Arpita, Matthew J. Wiswall, and Basit Zafar. 2020, August. “College Majors.” Nber working papers 27645, National Bureau of Economic Research, Inc.

Patterson, Richard, Nolan G. Pope, and Aaron Feudo. 2019, January. “Timing Is Everything: Evidence from College Major Decisions.” Iza discussion papers 12069, Institute of Labor Economics (IZA).

Perla, Jesse, and Christopher Tonetti. 2014. ““Equilibrium Imitation and Growth”.” *Journal of Political Economy* 122 (1): 52–76.

Roy, A. D. 1951. ““Some Thoughts on the Distribution of Earnings”.” *Oxford Economic Papers* 3 (2): 135–146.

Schultz, Theodore W. 1968. ““Resources for Higher Education: An Economist’s View”.” *Journal of Political Economy* 76:327–327.

Sonin, Isaac M. 2008. ““A Generalized Gittins Index for a Markov Chain and its Recursive Calculation”.” *Statistics & Probability Letters* 78 (12): 1526–1533.

Stinebrickner, Ralph, and Todd Stinebrickner. 2014. “A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout.” *The Review of Economic Studies* 81 (1 (286)): 426–472.

Stinebrickner, Todd, and Ralph Stinebrickner. 2012. “Learning about Academic Ability and the College Dropout Decision.” *Journal of Labor Economics* 30 (4): 707–748.

Traiberman, Sharon. 2019. ““Occupations and Import Competition: Evidence from Denmark”.” *American Economic Review* 109 (12): 4260–4301.

Varaiya, Pravin, Jean Walrand, and Cagatay Buyukkoc. 1985. ““Extensions of the Multiarmed Bandit Problem: The Discounted Case”.” *IEEE transactions on automatic control* 30 (5): 426–439.

Violante, Giovanni L. 2008. “Skill-biased technical change.” *The new Palgrave dictionary of economics* 2:1–6.

Whittle, Peter. 1980. ““Multi-Armed Bandits and the Gittins Index”.” *Journal of the Royal Statistical Society: Series B (Methodological)* 42 (2): 143–149.

Wiswall, Matthew, and Basit Zafar. 2015. “Determinants of College Major Choice: Identification using an Information Experiment.” *The Review of Economic Studies* 82 (2): 791–824 (4).

Zafar, Basit. 2011. “How Do College Students Form Expectations?” *Journal of Labor Economics* 29 (2): 301–348.

A Proofs

Proof of Proposition 1. Agents' human capital investment problem is a special case of a class of problems that consist of sequential choice among independent decision problems. This class generalizes the classical multi-armed bandit problem in that in addition to choosing which "arm" to "pull" at each stage, the decision maker must also choose *how* to pull the arm. Generally, this class of problems does not admit a simple solution. However, Assumptions 1 and 2 allow us to follow Glazebrook (1979) (see also Whittle 1980) and exploits the special structure of the human capital investment problem to show that it can indeed be decomposed and admits a simple "index characterization".

We first consider the following simplified problem. Suppose there is only one skill, and at each period either the skill is selected, or an "outside option" that yields a fixed, known payoff of $(1-\delta)M \geq 0$ whenever selected.

A policy π for this problem specifies, given each state $(\mathbf{h}, \mathbf{P})$ of the skill, whether to choose it – and if so whether to study it or enter the labor market – or to utilize the outside option and obtain a payoff of $(1-\delta)M$. Let π_M^* denote an optimal policy for this problem (i.e., a policy that maximizes the expected discounted payoff). For convenience, suppose that in cases of indifference between the outside option and the skill, π_M^* breaks ties in favor of the skill, and that in cases of indifference between study and work, π_M^* breaks ties in favor of study. For any $(\mathbf{h}, \mathbf{P})$, denote by $\mathcal{M}^*(\mathbf{h}, \mathbf{P}) \subseteq \mathbb{R}_+$ the set of outside-option payoffs M for which π_M^* selects the skill when it is in state $(\mathbf{h}, \mathbf{P})$.

Lemma 1. *Conditional on the skill being chosen under the optimal policy π^* for this problem, the decision to study or work is independent of the size of M: for any $(\mathbf{h}, \mathbf{P})$ and $\tilde{M}, \hat{M} \in \mathcal{M}^*(\mathbf{h}, \mathbf{P})$, $\pi_{\tilde{M}}^*$ prescribes* studying *the skill if and only if $\pi_{\hat{M}}^*$ does.*

Proof. Let $(\mathbf{h}, \mathbf{P})$ be the current state of the skill and let $M_1 \in \mathcal{M}^*(\mathbf{h}, \mathbf{P})$. Suppose that given M_1, studying the skill is strictly optimal. There must, therefore, exist a $\hat{\tau} > 0$ such that the payoff from continuing to study for $\hat{\tau}$ periods and then choosing whether to enter the labor market or choose the outside option yields a greater expected payoff than both choosing to enter the labor market immediately and choosing the outside option immediately:

$$\mathbb{E}\left[\delta^{\hat{\tau}}\max\left\{\frac{U(w.h_{\hat{\tau}})}{1-\delta}, M_1\right\}|(\mathbf{h}, \mathbf{P})\right] > \max\left\{\frac{U(w, h)}{1-\delta}, M_1\right\}. \quad (15)$$

Now suppose, towards a contradiction, that for some $M_2 \in \mathcal{M}^*(\mathbf{h},\mathbf{P})$ it is strictly optimal to choose the skill and enter the labor market. Then for all $\tau > 0$,

$$\frac{U(w,h)}{1-\delta} > \mathbb{E}\left[\delta^{\tau}\left\{\frac{U(w.h_{\tau})}{1-\delta}\right\}|(\mathbf{h},\mathbf{P})\right]. \tag{16}$$

Combining (15) and (16), it must be that

$$\begin{aligned}&\mathbb{E}\left[\delta^{\hat{\tau}}\left(\max\left\{\frac{U(w.h_{\hat{\tau}})}{1-\delta}, M_1\right\} - \frac{U(w.h_{\hat{\tau}})}{1-\delta}\right)|(\mathbf{h},\mathbf{P})\right] \\ &> \max\left\{\frac{U(w,h)}{1-\delta}, M_1\right\} - \frac{U(w,h)}{1-\delta},\end{aligned}$$

which is a contradiction since $U(w,h_{\hat{\tau}}) \geq U(w,h)$. □

Lemma 1 shows that the choice between study and work within each skill depends only on information about the skill, and is independent of outside information. We now use Lemma 1 to prove the Proposition.

Returning to the original problem with K skills, we again augment it with a fictitious outside option. Note that the outside option can be interpreted as an additional skill, which yields a fixed payoff $(1-\delta)M \geq 0$ in each period regardless of whether the individual studies or works. Clearly in this case the index associated with the outside option is equal to M. Consider now the policy described in Proposition 1, applied to this augmented problem. Note that since the payoff from the outside option is fixed, if this policy selects the outside option in some period, it will continue to select it in all subsequent periods (yielding a continuation payoff of M). Given the initial state $(\mathbf{h}_0,\mathbf{P}_0)$, denote by τ_k the number of periods spent on skill k, studying or working, until the index v_k falls below M. This time τ_k is of course stochastic, and may be infinite, in which case we let $\tau_k = \infty$. Denote by $V^{\overline{\tau}}$ the (stochastic) discounted payoff obtained during the $\overline{\tau} = \sum_{i=1}^{K} \tau_i$ periods spent on the K skills (with $\overline{\tau} = \infty$ if $\tau_i = \infty$ for some i). The expected discounted payoff $V((\mathbf{h}_0,\mathbf{P}_0),M)$ under the index policy described in the statement of Proposition 1 is then given by

$$V((\mathbf{h}_0,\mathbf{P}_0),M) = \mathbb{E}\left[V^{\overline{\tau}} + \delta^{\overline{\tau}}M|(\mathbf{h}_0,\mathbf{P}_0)\right].$$

Let $V_k^*((h_{k,0},P_{k,0}),M)$ be the value function for the auxiliary single-skill problem considered above, consisting of only skill k and the outside option M (note that for $M =$

0, V_k^* coincides with the index v_k). Since $\tau_1, ..., \tau_N$ are independent, and since each $V_i^*((h_{i,0}, P_{i,0}), M)$ is convex in M and differentiable almost everywhere (it is a maximum over stopping times of a linear function of M),

$$\frac{\partial}{\partial M} V((\mathbf{h}_0, \mathbf{P}_0), M) = \prod_{i=1}^{K} \frac{\partial}{\partial M} V_i^*((h_{i,0}, P_{i,0}), M).$$

Therefore, taking $\overline{M} > 0$ large enough, $V((\mathbf{h}_0, \mathbf{P}_0), \overline{M}) = \overline{M}$, and

$$V((\mathbf{h}_0, \mathbf{P}_0), 0) = \overline{M} - \int_0^{\overline{M}} \prod_{i=1}^{K} \frac{\partial}{\partial M} V_i^*((h_{i,0}, P_{i,0}), M) \mathrm{d}M. \tag{17}$$

Note that for $M = 0$, the augmented problem is equivalent to the original one (i.e., with no outside option). Hence, it remains to show that $V((\mathbf{h}_0, \mathbf{P}_0), 0)$ solves the Bellman equation, i.e., that

$$V((\mathbf{h}_0, \mathbf{P}_0), 0) = \max_i \Lambda_i V((\mathbf{h}_0, \mathbf{P}_0), 0),$$

with the operator Λ_i defined as

$$\Lambda_i G((\mathbf{h}_0, \mathbf{P}_0), 0) = \max \left\{ \underbrace{\delta \mathbb{E} \left[G((h_{i,1}, P_{i,1}), 0) | (h_{i,0}, P_{i0}) \right]}_{\text{study}}, \underbrace{\frac{U_i(w_i, h_{i,0})}{1 - \delta}}_{\text{enter labor market}} \right\}.$$

Similarly, for a function of $((h_{i,0}, P_{i,0}), 0)$ only, Λ_i will be defined in the same way by dropping the states of the other skills.

For each skill k, denote

$$V_{-k}^*((\mathbf{h}_0, \mathbf{P}_0), M) = \prod_{i \neq k} \frac{\partial}{\partial M} V_i^*((h_{i,0}, P_{i,0}), M).$$

In the remainder of the proof, to ease the notation, we drop all arguments except for M. Let $v^* \equiv \max_{i=1,...,K} v_i$ denote the highest index among the skills, given $(\mathbf{h}_0, \mathbf{P}_0)$. From (17),

$$\begin{aligned} V(0) &= \overline{M} + \int_0^{\overline{M}} V_{-k}^*(M) \frac{\partial}{\partial M} V_k^*(M) \mathrm{d}M \\ &= \overline{M} - \left(V_k^*(\overline{M}) V_{-k}^*(\overline{M}) - V_k^*(0) V_{-k}^*(0) \right) + \int_0^{\overline{M}} V_k^*(M) \mathrm{d}V_{-k}^*(M) \end{aligned}$$

$$= \overline{M} - \left(\overline{M} - V_k^*(0)V_{-k}^*(0)\right) + \int_0^{\overline{M}} V_k^*(M)\mathrm{d}V_{-k}^*(M)$$

$$= V_k^*(0)V_{-k}^*(0) + \int_0^{v^*} V_k^*(M)\mathrm{d}V_{-k}^*(M), \tag{18}$$

where the second equality follows from integration by parts, the third from the fact that $V_k^*(\overline{M}) = \overline{M}$ and $V_{-k}^*(\overline{M}) = 1$, and the fourth from the fact that $\frac{\mathrm{d}V_{-k}^*(M)}{\mathrm{d}M} = 0$ for all $M \geq v^*$.

Without loss of generality suppose the skill with highest index given $(\mathbf{h}_0, \mathbf{P}_0)$ is skill 1. From Lemma 1, for all values of M below v_1, $V_1^*(M) = \Lambda_1 V_1^*(M)$ (note that this need not be true in general, and is a consequence of Lemma 1). Hence, from (18), $V(0) = \Lambda_1 V(0)$. Furthermore, from (18), since $V_i^*(M) \geq \Lambda_i V_i^*(M)$ for all i and M, and since $\frac{\mathrm{d}V_{-i}^*(M)}{\mathrm{d}M} \geq 0$ for all i, it must be that $V(0) \geq \Lambda_i V(0)$ for all i.[17] Therefore, since we have shown that $V(0) = \Lambda_1 V(0)$ and that $V(0) \geq \Lambda_i V(0)$ for all skills $i = 1, ..., K$, this implies $V(0) = \max_{i=1,...,K} \Lambda_i V(0)$. Thus, V solves the Bellman equation. ■

Proof of Proposition 2. The result follows from Proposition 1, once we establish that the indices (9) are a special case of the indices v and that (10) characterizes the graduation region $\mathcal{G}_k$, as defined in Section 2.

Skill k's graduation region $\mathcal{G}_k$ (as defined in (5)) can be derived as follows. For a given state $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$, consider first the following "one-step-look-ahead" inequality, which guarantees the value of entering the labor market immediately as a skill-k specialist is weakly greater than the expected value of studying k for a *single* additional period and then entering the labor market:

$$h_{k,t} \geq \delta \mathbb{E}\left[h_{k,t+1} | (h_{k,t}, \alpha_{k,t}, \beta_{k,t})\right]. \tag{19}$$

Rewritten as

$$h_{k,t} \geq \frac{\delta}{1-\delta} \mathbb{E}\left[a_{k,t} | (h_{k,t}, \alpha_{k,t}, \beta_{k,t})\right], \tag{20}$$

the condition states that one period of foregone earnings is greater than the expected marginal benefit of continuing to study for an additional period. Using (8), we rewrite

[17] That $\frac{\mathrm{d}V_{-k}^*(M)}{\mathrm{d}M} \geq 0$ follows from the convexity of each V_i^*, $i \neq k$.

(19) as

$$h_{k,t}(1-\delta) \geq \delta\nu_k \left(\frac{\alpha_{k,t}}{\alpha_{k,t}+\beta_{k,t}}\right). \tag{21}$$

It is easy to verify that this inequality guarantees that the value of entering the labor market immediately as a k-specialist exceeds the value of studying k for *at least* one additional period. This follows from that fact that the skill-k stopping problem – deciding when to stop studying k and enter the market as a k-specialist, ignoring all other skills – is *monotone* (Chow and Robbins 1961).[18]

Note that the number of successes the individual has experienced studying skill k in the periods prior to period t is equal to $(h_{k,t} - h_{k,0})/\nu_k$. As a result, $\alpha_{k,t} = h_{k,t}/\nu_k$. It therefore follows that the work region $\mathcal{G}_k$ is the set of skill-k states that satisfy (10). Next, denote by $m_{k,t}$ the number of periods the individual has studied skill k prior to period t, and note that $\alpha_{k,t} + \beta_{k,t} = \alpha_{k,0} + \beta_{k,0} + m_{k,t}$. Given (10), the number of periods m_k^* the individual spends studying skill k before reaching a state in the work region, is equal to the smallest integer greater than $\frac{\delta}{1-\delta} - (\alpha_{k,0} + \beta_{k,0})$. That is, $m_k^* = 0$ if $\frac{\delta}{1-\delta} < \alpha_{k,0} + \beta_{k,0}$, and otherwise

$$m_k^* = \lceil \frac{\delta}{1-\delta} \rceil - (\alpha_{k,0} + \beta_{k,0}). \tag{22}$$

Clearly, if $m_{k,t} \geq m_k^*$, the state $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$ of skill k is in the work region, and hence the index of the skill is $v_k(h_{k,t}, \alpha_{k,t}, \beta_{k,t}) = h_k/(1-\delta)$. Now suppose $m_{k,t} < m_k^*$. From (4), the index of the skill is equal to the expected discounted payoff from studying the skill for exactly an additional $m_k^* - m_{k,t}$ periods, and then entering the labor market as a skill-k specialist. Given the current state $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$ and $m_{k,t}$, the distribution of the number of "successes" during these $m_k^* - m_{k,t}$ periods of studying skill k is Binomial with parameters $(m_k^* - m_{k,t}, \alpha_{k,t}/(\alpha_{k,t} + \beta_{k,t}))$, and the expected number of successes is $(m_k^* - m_{k,t})\left(\alpha_{k,t}/(\alpha_{k,t} + \beta_{k,t})\right)$. The index of skill k is therefore:

$$\begin{aligned} v_k(h_{k,t}, \alpha_{k,t}, \beta_{k,t}) &= \frac{\delta^{m_k^*-m_{k,t}}}{1-\delta}\left(h_{k,t} + \nu_k(m_k^* - m_{k,t})\left(\frac{\alpha_{k,t}}{\alpha_{k,t}+\beta_{k,t}}\right)\right) \\ &= \frac{h_{k,t}}{1-\delta}\delta^{m_k^*-m_{k,t}}\left(\frac{m_k^* - m_{k,t} + \alpha_{k,t} + \beta_{k,t}}{\alpha_{k,t}+\beta_{k,t}}\right) \end{aligned}$$

[18]We say the skill-k stopping problem is monotone if, for any period t and $(h_{k,t}, \alpha_{k,t}, \beta_{k,t})$, $h_{k,t} \geq \delta\mathbb{E}\left[h_{k,t+1}|(h_{k,t}, \alpha_{k,t}, \beta_{k,t})\right]$ implies that for any realization of $a_{k,t}$, and resulting $(h_{k,t+1}, \alpha_{k,t+1}, \beta_{k,t+1})$, $h_{k,t+1} \geq \delta\mathbb{E}\left[h_{k,t+2}|(h_{k,t+1}, \alpha_{k,t+1}, \beta_{k,t+1})\right]$.

$$= \frac{h_{k,t}}{1-\delta}\left(\frac{\lceil\frac{\delta}{1-\delta}\rceil\delta^{\lceil\frac{\delta}{1-\delta}\rceil-(\alpha_{k,t}+\beta_{k,t})}}{\alpha_{k,t}+\beta_{k,t}}\right),$$

where the second equality follows from the fact that $\nu_k\alpha_{k,t} = h_{k,t}$, and the fourth equality follows from (22). ■

Proof of Proposition 3. *Part 1.* The educational markup of a given skill k is equal to $\Omega_k(x) = \lceil\frac{\delta}{1-\delta}\rceil\delta^{\lceil\frac{\delta}{1-\delta}\rceil}\cdot\frac{\delta^{-x}}{x}$. The derivative of $\frac{\delta^{-x}}{x}$ is $\frac{\delta^{-x}(-1-xln(\delta))}{x^2}$, which is negative if and only if $-1-xln(\delta)<0$. The latter condition is equivalent to $-ln(\delta)<\frac{1}{x}$. Since the skill has not yet reached its graduation region, $\frac{\delta}{1-\delta}>x$, which can be rearranged as $\frac{1}{x}>\frac{1}{\delta}-1$. The derivative of Ω_k is therefore negative, as $\frac{1}{\delta}-1>-ln(\delta)$ for all $\delta\in(0,1)$. Next, note that $\Omega_k(x)>1$ is equivalent to $\lceil\frac{\delta}{1-\delta}\rceil\delta^{\lceil\frac{\delta}{1-\delta}\rceil-x}>x$, which indeed holds whenever $\lceil\frac{\delta}{1-\delta}\rceil>x$.

Part 2. We now show that the marginal reduction in Ω_k is diminishing; that is, that $\Omega_k(x+1)-\Omega_k(x)$ is decreasing in x. To prove the claim, we must show that the derivative of $\delta^{-x}\left(\frac{1}{\delta(x+1)}-\frac{1}{x}\right)$ is negative. Indeed, the derivative of the latter is given by $\delta^{-x}\left[\frac{xln(\delta)+1}{x^2}-\frac{(x+1)ln(\delta)+1}{\delta(x+1)^2}\right]$, which is negative since $\frac{xln(\delta)+1}{x^2}<\frac{(x+1)ln(\delta)+1}{\delta(x+1)^2}$. ■

Proof of Proposition 4. The proof follows from arguments similar to those in Gossner, Steiner, and Stewart (2021) and is therefore omitted. ■

Proof of Proposition 5. The proof follows from arguments similar to those in the proof of Proposition 3. ■

Proof of Proposition 6. Denote the policy described in the statement of the proposition by π^*. Recall that the index of any skill k in a general value state $(h_k, P_k)\notin\mathcal{E}_k$ is defined as in (14). Let the index of any skill in a state $(h_k, P_k)\in\mathcal{E}_k$ be equal to ∞. For any skill k in a state $(h_k, P_k)\notin\mathcal{E}_k$, let $\lambda_k^*(h_k, P_k)$, denote the labor-market-entry rule that solves the maximization problem in (14). Analogously to the proof of Proposition 1, this rule can be shown to be invariant in the size of a fictitious retirement option, and hence it will be optimal to follow λ_k^* whenever skill j is selected. The new complication complementarities introduce is in the choice *between* skills.

Denote the the policy described in the statement of Proposition 6 by π^*. Under this policy, skills in a state with general value are all treated the same, and have priority over skills in state $(h_k, P_k) \notin \mathcal{E}_k$. Denote the index of the former by ∞. Among skills in states $(h_k, P_k) \notin \mathcal{E}_k$, those with a higher index are preferred. Among skills in states $(h_k, P_k) \in \mathcal{E}_k$, all skills have the same indices. Indices are therefore ordered in increasing order, with ∞ preferred to any finite index. Denote this ordering by $\succeq$.

The optimal stopping times in (14) satisfy the following crucial property. Suppose that a skill k is in state $(h_{kt}, P_{kt}) \notin \mathcal{E}_k$ in period t. Then $\tau_k^*(h_{k,t}, P_{k,t})$ solving the maximization problem in (14) is the first time $\hat{t} > t$ in which $v_k^C(h_{k,t}, P_{k,t}) \succeq v_k^C(h_{k,\hat{t}}, P_{k,\hat{t}})$. That is, it is the first time at which the skill's index drops to a value that is less preferred under the order $\succeq$.

We now use an interchange argument to show that this policy is optimal. Let π^0 be a policy that chooses skill i in period 0 and then proceeds according to the policy π^* from period 1 onward. Following standard results in the literature on Markov decision processes, in order to establish the optimality of π^*, it is sufficient to show that the expected discounted payoff under π^0 is no greater than that under π^*, given any initial state $(\mathbf{h}_0, \mathbf{P}_0)$.

Let $(\mathbf{h}_0, \mathbf{P}_0)$ be the initial state of the agent's problem. Consider the policy π^0. If π^0 selects in period 0 the same skills as π^* would have, the two policies coincide in all periods. Therefore, suppose that π^0 selects skill i in period 0, while π^* would have selected skill $j \neq i$ in period 0. This means $v_j^C(h_{j,0}, P_{j,0}) \succeq v_i^C(h_{i,0}, P_{i,0})$. Furthermore, despite the fact that π^0 proceeds according to π^* from period 1 onward, π^0 need not select skill j in period 1, since the skill i's state may change after it is studied in period 0.

Define $\tau_k^*(h_k, P_k)$ to be the first time such that, starting from the state (h_k, P_k) and repeatedly choosing k, the index of skill k is no longer better than $v_k^C(h_k, P_k)$ according to $\succeq$. Note that this time is stochastic. Denote by σ_1 the (stochastic) time at which a skill other than i is selected under π^0. Without loss of optimality, assume this will be skill j. As j has not been selected yet, its state in period σ_1 is equal to that of period 0. Let $\tau_j^*(h_{j,0}, P_{j,0})$ be the optimal stopping time in the definition of the index of j given state $(h_{j,0}, P_{j,0})$. Setting $\sigma_2 = \tau_j^*(h_{j,0}, P_{j,0})$, π^0 will therefore choose skill j from period σ_1 until (at least) period $\sigma_1 + \sigma_2 - 1$. At time $\sigma_1 + \sigma_2$, the skill i's index will be $v_i^C(h_{i,\sigma_1}, P_{i,\sigma_1})$, skill j's index will be $v_j^C(h_{j,\sigma_2}, P_{j,\sigma_2})$, and the index of all other skills remains $v_k^C(h_{k,0}, P_{k,0})$.

Define the policy π^1 that initially invests in skill j during periods $0, ..., \sigma_2 - 1$, then invests in skill i, during periods $\sigma_2, ..., \sigma_2 + \sigma_1 - 1$, and then coincides with π^* thereafter.

We now show that the expected payoff under π^1 is weakly greater than under π^0. Denote by $u_k(h_k, P_k)$ the payoff from skill k given state (h_k, P_k). If the state (h_k, P_k) is in skill k's graduation region, then $u_k(h_k, P_k) = h_k^{\eta_k}$. Otherwise, $u_k(h_k, P_k) = 0$. The expected payoff during periods $0, ..., \sigma_1 + \sigma_2 - 1$ under π^1 is equal to

$$\prod_{k \neq i,j} h_{k,0}^{\eta_{k,0}} \left\{ h_{i,0}^{\eta_{i,0}} \mathbb{E} \left(\sum_{t=0}^{\sigma_2 - 1} \delta^t u_j(h_{j,t}, P_{j,t}) \right) + \mathbb{E} \left(\delta^{\sigma_2} h_{j,\sigma_2}^{\eta_{j,\sigma_2}} \right) \mathbb{E} \left(\sum_{t=0}^{\sigma_1 - 1} \delta^t u_i(h_{i,t}, P_{i,t}) \right) \right\}. \tag{23}$$

Similarly, under π^0, the expected payoff during these periods is equal to

$$\prod_{k \neq i,j} h_{k,0}^{\eta_{k,0}} \left\{ h_{j,0}^{\eta_{j,0}} \mathbb{E} \left(\sum_{t=0}^{\sigma_1 - 1} \delta^t u_i(h_{i,t}, P_{i,t}) \right) + \mathbb{E} \left(\delta^{\sigma_1} h_{i,\sigma_1}^{\eta_{i,\sigma_1}} \right) \mathbb{E} \left(\sum_{t=0}^{\sigma_2 - 1} \delta^t u_j(h_{j,t}, P_{j,t}) \right) \right\}. \tag{24}$$

Denote by $\Delta(\pi^1, \pi^0)$ the difference between the expected discounted payoff under π^1 and its counterpart under π^0. Subtracting (24) from (23) and rearranging, we have that $\Delta(\pi^1, \pi^0)$ is equal to

$$\begin{aligned} \prod_{k \neq i,j} h_{k,0}^{\eta_{k,0}} \Bigg\{ & \mathbb{E} \left(\sum_{t=0}^{\sigma_2 - 1} \delta^t u_j(h_{j,t}, P_{j,t}) \right) \left(h_{i,0}^{\eta_{i,0}} - \mathbb{E} \left(\delta^{\sigma_1} h_{i,\sigma_1}^{\eta_{i,\sigma_1}} \right) \right) \\ & - \mathbb{E} \left(\sum_{t=0}^{\sigma_1 - 1} \delta^t u_i(h_{i,t}, P_{i,t}) \right) \left(h_{j,0}^{\eta_{j,0}} - \mathbb{E} \left(\delta^{\sigma_2} h_{j,\sigma_2}^{\eta_{j,\sigma_2}} \right) \right) \Bigg\}. \end{aligned} \tag{25}$$

We now verify that $\Delta(\pi^1, \pi^0) \geq 0$. Recall that $v_j^C(h_{j,0}, P_{j,0}) \succsim v_i^C(h_{i,0}, P_{i,0})$. We must consider the following cases.

Case 1. Suppose that $(h_{j,0}, P_{j,0})$ is a state of general value and $(h_{i,0}, P_{i,0})$ is not. Then, by the definition of $\sigma_2 = \tau_j^*(h_{j,0}, P_{j,0})$, $h_{j,0}^{\eta_{j,0}} - \mathbb{E}\left(\delta^{\sigma_2} h_{j,\sigma_2}^{\eta_{j,\sigma_2}}\right) < 0$, and because i is not of general value, $h_{i,0}^{\eta_{i,0}} - \mathbb{E}\left(\delta^{\sigma_1} h_{i,\sigma_1}^{\eta_{i,\sigma_1}}\right) \geq 0$. This guarantees that $\Delta(\pi^1, \pi^0) \geq 0$.

Case 2. Suppose that $(h_{j,0}, P_{j,0})$ and $(h_{i,0}, P_{i,0})$ are both states of general value. Note that this implies $u_j(h_{j,t}, P_{j,t}) = u_i(h_{i,t}, P_{i,t}) = 0$ for all $t = 0, ..., \sigma_2 - 1$, as these periods are necessarily spent studying. So we have $\Delta(\pi^1, \pi^0) = 0$ in this case.

Case 3. Suppose that $(h_{j,0}, P_{j,0})$ and $(h_{i,0}, P_{i,0})$ are both not states of general interest. Then $v_j^C(h_j, P_j) \geq v_i^C(h_i, P_i)$ and both indices are finite. Furthermore, $h_{i,0}^{\eta_{i,0}} - \mathbb{E}\left(\delta^{\sigma_1} h_{i,\sigma_1}^{\eta_{i,\sigma_1}}\right) \geq 0$ and $h_{j,0}^{\eta_{j,0}} - \mathbb{E}\left(\delta^{\sigma_2} h_{j,\sigma_2}^{\eta_{j,\sigma_2}}\right) \geq 0$, and by (14),

$$\frac{\mathbb{E}\left(\sum_{t=0}^{\sigma_1 - 1} \delta^t u_i(h_{i,t}, P_{i,t})\right)}{h_{i,0}^{\eta_{i,0}} - \mathbb{E}\left(\delta^{\sigma_1} h_{i,\sigma_1}^{\eta_{i,\sigma_1}}\right)} \leq v_i^C(h_{i,0}, P_{i,0}) \leq v_j^C(h_{j,0}, P_{j,0}) = \frac{\mathbb{E}\left(\sum_{t=0}^{\sigma_2 - 1} \delta^t u_j(h_{j,t}, P_{j,t})\right)}{h_{j,0}^{\eta_{j,0}} - \mathbb{E}\left(\delta^{\sigma_2} h_{j,\sigma_2}^{\eta_{j,\sigma_2}}\right)}.$$

Rearranging and multiplying by $\prod_{k\neq i,j} h_{k,0}^{\eta_{k,0}}$, we have that $\Delta(\pi^1,\pi^0)\geq 0$,

We have therefore shown that the expected payoff under π^1 is weakly greater than under π^0. Note that if π^1 coincides with π^* during the periods $\sigma_2,...,\sigma_2+\sigma_1-1$, then π^1 and π^* are identical and the proof is complete. Otherwise, modify π^1 to a new policy π^2, repeating the argument in the preceding paragraphs. We can proceed inductively to construct a sequence of policies $(\pi^0,\pi^1,\pi^2,...)$, such that (i) given the initial state $(\mathbf{h}_0,\mathbf{P}_0)$, π^{s+1} yields an expected discounted payoff weakly greater than π^s, and (ii) the expected discounted payoff under π^s converges to the expected discounted payoff under π^* as $s\to\infty$. It follows that the expected discounted payoff under π^0 is no greater than under π^*, which completes the proof. ■

Proof of Proposition 7. Recall that the condition determining whether a skill is in a state for which it has general value is given by (13). At any period t, for any skill k, and any $h_{k,t}$ and $\eta_{k,t}$, a first-order stochastic improvement in the distribution Ψ_k relaxes the condition in (13). ■

Proof of Proposition 8. Recall that the condition determining whether a skill is in a state for which it has general value is given by (13). At any period t, for any skill k, and any $h_{k,t}$ and $\eta_{k,t}$, a mean preserving spread of the distribution Ψ_k relaxes the condition in (13). This is because because $h_{k,t+\tau}^{\eta_{k,t+\tau}}$ is a convex function of $\eta_{k,t+\tau}$, such a mean preserving spread must increase $\mathbb{E}(h_{k,t+\tau}^{\eta_{k,t+\tau}})$. Therefore, at any state and for any realization path, $\mathbb{E}\left[h_{k,t+\tau}^{\eta_{k,t+\tau}}\right]$ increases, making the inequality easier to satisfy. General education will therefore last longer. ■

B Computational Appendix

This appendix contains additional computational details on the simulations in Figures 1 and 2. The simulations illustrate the model's dynamics under the optimal policy by simulating the academic history of 100,000 college students. Each model period corresponds to one college course credit. Allowing eight course credits per year, as per the standard college schedule, we initialize individuals at age 18 and calibrate the model so that college (endogenously) lasts an average of four years (32 periods). For the illustration of educational interventions in Figure 2, the parameterization additionally targets an 18% share of STEM college graduates, as in the data. Education interventions at the start of Freshman, Sophomore, Junior, and Senior year correspond to model periods 1, 9, 17, and 25, respectively.

Despite its parsimony, the parametric model has many more degrees of freedom than are necessary for the illustrative exercise. To simplify it further, we set the parameters of the human capital technology, skill wages, and initial conditions to unity for all skills, so that $w_k = 1$, $h_{0,k} = 1$ and $v_k = 1$ for all k in all simulations. To calibrate each simulation, we vary the discount factor (δ) and parameters of the population ability distributions (α_k, β_k). Together, these parameters determine the optimal years of schooling, as δ determines the opportunity cost of time, and the ability distributions determine the average returns to schooling and, via rational expectations, the degree of uncertainty individuals face in knowing their own abilities. Clearly, many combinations of these parameters can generate the same targeted years of schooling. The illustrations select among these calibrations to produce visually pronounced dynamics and avoid degenerate cases.

For Figure 1, the discount factor is set to $\delta = 0.975$ and $(\alpha_k, \beta_k) = (1, 22)$ for all k. For Figure 2, the discount factor is set to $\delta = 0.970$, the STEM (S) ability distribution is $(\alpha_S, \beta_S) = (2, 10)$, and the non-STEM (NS) ability distribution is $(\alpha_{NS}, \beta_{NS}) = (3, 10)$.

For the education intervention simulations in Figure 2, the indices in the optimal policy do not admit closed-form solutions. As a result, we calculate the indices numerically using the recursive methods in Sonin (2008). In general, while closed-form expressions produce the most efficient computation, a variety of methods exist to calculate the indices of the optimal policy, even when their precise analytical form is not known (for example, see also Varaiya, Walrand, and Buyukkoc 1985; Katehakis and Veinott 1987).