The Parental Investment Rat Race: College Admissions Tournament and Household Welfare*

Hyunjae Kang

September 2026
(First version: November 2022)

Abstract

How much of a household's educational investment is driven by competition rather than direct human-capital returns? I develop and estimate a dynamic tournament model of admissions competition embedding skill formation, where households choose tutoring expenditure, tutoring hours, and self-study hours. Competition drives 77% of tutoring expenditure during secondary school, and 89% in the final year, holding direct returns to human capital fixed. Less than half of self-study time is competition-driven. Although both investments erode non-cognitive skills, tournament and cognitive returns outweigh these costs. Policies targeting educational investment directly improve mobility at a substantial welfare cost, while admissions reforms can achieve mobility gains at a smaller welfare cost.

JEL Classification Codes: D15, D64, I21, I22, I24, I26, J62
Keywords: Parental Investment, Child's Efforts, Student Competition, Skill Formation, Non-cognitive skills

*Institute of Economic Research, Kyoto University. Email: jay.kang@kier.kyoto-u.ac.jp. I am immensely grateful to Steven Stern and Juan Pantano for their invaluable guidance and support. I would also like to thank Meta Brown and David Wiczer for their valuable advice. This paper benefited from helpful comments from Francesco Agostinelli, Peter Arcidiacono, Minji Bang, Elena Capatina, Heepyung Cho, Sunghun Cho, Juan Carlos Conesa, Olivier De Groote, Marcos Fernandes, Chao Fu, Yujung Hwang, Shoya Ishimaru, Fedor Iskhakov, Changhui Kang, Shinhyuk Kang, Michael Keane, Minsub Kim, Sunham Kim, Lance Lochner, Nirav Mehta, Gabriel Mihalache, Mark Montgomery, Joseph Mullins, John Rust, Warren Sanderson, Todd Stinebrickner, Michèle Tertilt, Michela Tincani, Takashi Unayama, Greg Veramendi, Makoto Watanabe, Ken Yamada, Minchul Yum, Basit Zafar, Weilong Zhang, Yiyi Zhou, and many seminar/conference participants. This paper supersedes a previous version circulated as "Dynamic Competition in Parental Investment and Child's Efforts". All errors are my own.

1 Introduction

Competition for limited seats at prestigious colleges generates a "rat race" of household investments. How much parental investment is driven by competition rather than by direct human-capital returns? Parental investment is known to play an important role in forming early childhood skills (Cunha, Heckman and Schennach, 2010; Del Boca, Flinn and Wiswall, 2014) and intergenerational persistence of earnings (Bolt, French, Hentall-MacCuish and O'Dea, 2026; Gayle, Golan and Soytas, 2022). By adolescence, however, a substantial share of parental investment can be driven by competitive motives. In the United States, rising competition for elite universities has been linked to increases in parental time investment (Ramey and Ramey, 2010; Mahler, Tertilt and Yum, 2025). In East Asia, parents frequently dedicate a significant portion of their income to private tutoring to improve their children's relative standing for admission to prestigious universities even though evidence suggests that private tutoring has only modest effects on absolute test score improvements (Ryu and Kang, 2013; Kang and Park, 2021).[1] In such settings, competitive externalities may push parental investment beyond the level justified by its direct human-capital returns.[2]

This paper asks two questions. First, how much of parental investment and child self-effort is driven by the competitive externalities of the college admissions tournament rather than by direct human-capital returns? Second, how do policies that alter admission competition affect household welfare and intergenerational mobility? Parental investment can affect outcomes through two channels: it can improve a student's relative position in the admission tournament, and it can raise human capital directly. Disentangling these channels matters for welfare. If a substantial share of investment is driven by competition rather than direct human-capital returns, policies that reduce the demand for such investment may improve welfare. Self-study is less strongly tied to parental background than tutoring expenditure, so omitting it would misstate the role of admission competition in intergenerational mobility. Moreover, self-effort itself may respond to competitive incentives.

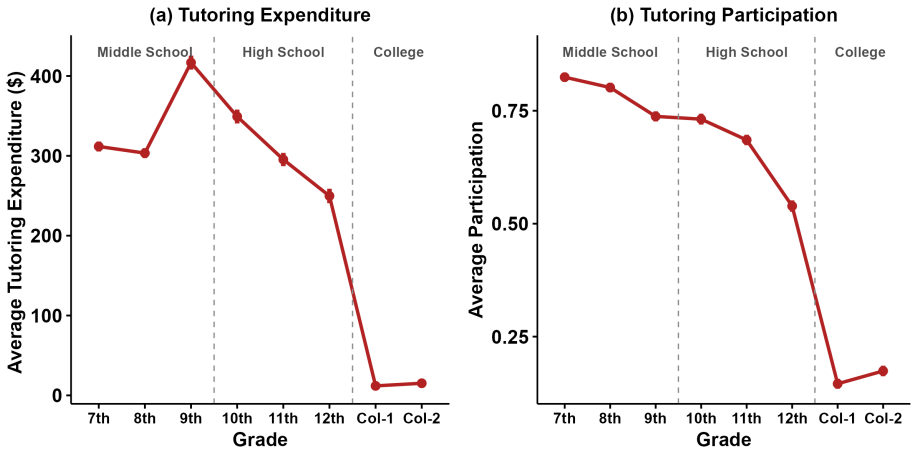
To answer these questions, I build and estimate a dynamic tournament model of college admission with endogenous skill formation. I estimate the model using longitudinal data from South Korea on household income, tutoring expenditure and hours, children's time allocation, and administrative test scores, linked to local tutoring-market measures and auxiliary earnings data. In the model, households choose tutoring quality, tutoring hours, and child self-study during the secondary school years. These choices produce final test scores, which determine college-tier assignment through

[1] Notably, for China, South Korea, and Japan, the reported participation rates for private tutoring in secondary school are 48.7%, 61.8%, and 51.9%, respectively (Bray, 2022).

[2] Suggestive of the competitive motive, private tutoring collapses immediately once the admission outcome is realized (Figure 1).

endogenous admission cutoffs. College tier, together with the child's skill, generates future earnings. The estimated model allows me to quantify the competition-driven component of investment and to evaluate how changes in admission competition affect tutoring, self-study, household welfare, and intergenerational mobility.

Figure 1: Private Tutoring Expenditure and Participation in Tutoring



Source: Korea Educational Longitudinal Study 2005, Korean Educational Development Institute.

The institutional features of South Korea provide a useful setting for studying how household income shapes children's educational outcomes. With relatively homogeneous secondary-school quality, private tutoring expenditure is a major channel through which parents influence these outcomes. Substantial earnings differences across college tiers create incentives for such investment.[3]

I document descriptive evidence supporting the tournament framework. Using longitudinal data from South Korea, I find that college tier, a vertical ordering of colleges, strongly predicts both initial earnings and earnings growth of alumni, and the estimates remain economically and statistically significant after controlling for college entrance exam scores. This evidence is consistent with recent empirical work on elite college premiums (Zimmerman, 2019; Sekhri, 2020; Jia and Li, 2021; Lee and Koh, 2023).[4]

A second empirical fact suggests that parental investment and self-study have different implications for intergenerational mobility. While parental investment varies substantially with household income, hours of self-study show limited correlation with it. At the same time, both inputs positively affect child outcomes. If the two are substitutable, income-constrained households could compensate for lower parental investment through increased self-study. Omitting self-study from the analysis could therefore overstate the role of parental investment in child outcomes and distort conclusions

[3]Several countries share these institutional features, as discussed in Section 2.

[4]Using data on students majoring in science at University of California campuses, Arcidiacono, Aucejo and Hotz (2016) show that mismatch between minority students' preparedness and higher-ranked campuses decreases the likelihood of graduation.

about welfare and intergenerational mobility.

Motivated by this evidence, I develop and estimate an equilibrium dynamic tournament model of college admission competition. The model embeds the rank-order tournament of Lazear and Rosen (1981) into a dynamic model of altruistic households who value their child's lifetime income. Each period, given its state variables, a household chooses parental investment and the child's self-study, which produce the child's test scores and non-cognitive skills. The ranking of final test scores and the number of available seats determine college tier assignment through endogenous admission cutoffs. Lifetime income depends on the assigned tier and on the child's skills, so investment is productive as well as positional. Unlike the classic tournament, the prize structure here is not designed by a principal to elicit effort: tier premia are set by the labor market and the number of seats.

Two features of the model are central to answering the research questions. First, the tournament structure with a fixed number of seats incorporates both channels through which household investment affects lifetime earnings: enhancing the child's relative ranking in the admission competition and improving earnings directly through human capital. This allows me to quantify the competition-driven share of parental investment and self-study, and because the model is dynamic, to quantify how this share evolves across grades. Second, because the number of seats is an explicit element of the model, I can simulate how exogenous changes in college seats affect welfare and intergenerational mobility, and compare these counterfactuals against a tutoring ban of the kind implemented in China in 2021.

I estimate the model using Maximum Simulated Likelihood augmented with moment conditions from auxiliary data sources. The estimates suggest that the marginal effects of both investments decline substantially over age, with minimal effects in the final period. Despite these low marginal effects, households maintain high investment. The tournament structure sustains investment incentives even when investment barely raises final test scores. Also, investments harm children's non-cognitive skills. However, households invest nonetheless because the cognitive channel dominates in earnings.

Using the estimated model, I quantify the share of parental investment and self-study driven by competition. I remove the tournament channel by equalizing the tier component of lifetime earnings across colleges, so that earnings depend on the child's absolute skills but not on admission rank. The model shows that 77% of tutoring expenditure and 47% of self-study are driven by college admission competition. The competition-driven share of tutoring expenditure rises sharply with age, reaching 89% in 12th grade. This large competition-driven share motivates the evaluation of whether policies that reduce the demand for tutoring can actually improve welfare.

Finally, I compare policies that directly regulate tutoring with reforms to college

admissions that alter the incentives for educational investment. Tutoring bans and an hours curfew restrict investment directly, while admission lotteries and seat expansion change how college opportunities are allocated or supplied. A tutoring ban substantially increases intergenerational mobility at a welfare cost of 5.5 percent of lifetime earnings. A curfew reduces tutoring hours, but households partially offset the restriction through more expensive tutoring per hour and more self-study. An admission lottery modestly reduces investment while largely preserving human capital, at a substantially smaller welfare cost. Expanding top-tier seats reduces tutoring expenditure across the calibrated range of premium declines. A 50 percent expansion of seats at the most selective colleges, under the baseline calibration of the earnings-premium response, raises elite access from the bottom parental-income quintile by roughly a third at a small welfare cost.

Related Literature and Contributions. First, this paper relates to the literature on parental investment (Becker and Tomes, 1979; Del Boca, Flinn and Wiswall, 2014; Doepke and Zilibotti, 2017; Lee and Seshadri, 2019; Bolt, French, Hentall-MacCuish and O'Dea, 2026), and particularly to work in which social interactions or market competition shape household investments in children. One strand studies peer and neighborhood effects. Agostinelli (2018) and Agostinelli, Doepke, Sorrenti and Zilibotti (2026) model equilibrium parental investment when children's skill accumulation depends partly on their peer or neighborhood environment, which parents may themselves influence. This paper studies a different externality: competition for a fixed number of college seats, in which one household's investment reduces other households' admission chances. Ramey and Ramey (2010) provide an early theoretical account of rising parental time investment as a response to competition for elite college admissions, while Bodoh-Creed and Hickman (2019) develop and estimate a structural model of an admissions contest to study pre-college human-capital investment.[5] I contribute to this literature along three dimensions. First, the tournament is dynamic, allowing competitive incentives to evolve over six years of skill formation. Second, separately observed tutoring and self-study at each grade allow me to estimate the competition-driven component of each investment and its age profile. Third, academic effort also shapes non-cognitive skills, whose earnings consequences enter the household's investment decisions.

Second, this paper relates to the literature on child development, skill formation, and their intergenerational implications (Cunha, Heckman and Schennach, 2010; Agostinelli and Wiswall, 2025; Daruich, 2026; Caucutt, Lochner, Mullins and Park, 2026). In particular, Del Boca, Flinn and Wiswall (2014) build and estimate a dynamic model of

[5] Grau (2018) estimates a static tournament model of college competition in Chile, abstracting from parental investment.

parental investment and cognitive development that separately identifies parental time and monetary investments. Emphasizing children's own time investment, Del Boca, Flinn, Verriest and Wiswall (2026) develop a Stackelberg model of parent-child interaction, while De Groote (2025) quantifies the role of student effort in an academic tracking system.[6] Fu and Mehta (2018) jointly model school and parental effort to study how ability tracking affects student achievement. Caucutt, Lochner, Mullins and Park (2026) use input-price variation to distinguish parental time from market-based investments. Here, purchased tutoring and children's own study time play analogous roles, with local tutoring supply providing variation in the relative cost of the two inputs. These papers treat the returns to investment as operating through skill production. I depart from this framework in two ways. First, households in my model invest not only to raise their children's skills but also to improve their rank in the admissions tournament. Second, I allow tutoring and self-study to affect cognitive and non-cognitive skills differently. A longitudinal education dataset containing tutoring expenditure, tutoring hours, self-study, non-cognitive measures, and test scores, combined with later earnings data, allows me to estimate these age-specific production channels within one framework.

Third, this paper relates to recent work on education competition in East Asia and its policy responses. Kim, Tertilt and Yum (2024) study low fertility in South Korea through a model of a "status externality," in which parents care about their children's human capital relative to that of other children. Gu and Zhang (2026) study college-admission competition in China using a heterogeneous-agent life-cycle model. More recently, Kim (2026) studies how education competition affects fertility, representing college admissions and access to preferred careers through a single assignment contest. I contribute to this policy discussion by jointly estimating household preferences and skill-production technologies from longitudinal observations of investment choices and subsequent skills, treating self-study as a separate alternative to purchased tutoring. This empirical framework allows me to quantify how substitution between tutoring and self-study and their effects on cognitive and non-cognitive skills shape the welfare and mobility consequences of tutoring regulation and admissions reform.

The rest of the paper is organized as follows. I describe the institutional features in Section 2. In Section 3, I document empirical facts that motivate the dynamic tournament model. Section 4 introduces the tournament model. Section 5 describes the estimation strategy. Section 6 presents the estimation results. Section 7 presents the decomposition of household investments. I present the counterfactual policy analyses

[6]Several studies estimate the relationship between self-study hours and academic achievement (Cooper, Robinson and Patall, 2006; Stinebrickner and Stinebrickner, 2008; Todd and Wolpin, 2018). This paper also complements the literature on private tutoring, which largely uses regression-based methods to study tutoring and academic outcomes (Stevenson and Baker, 1992; Cheo and Quah, 2005; Tansel and Bircan, 2005; Dang, 2007; Ono, 2007; Ryu and Kang, 2013; Kang and Park, 2021).

in Section 8 and conclude in Section 9.

2 Key Institutional Features

The institutional features of South Korea's education system offer a useful setting for understanding how parents' resources shape child's outcome through college admission competition. In this section, I explain the key institutional features of the country: the high-stakes college entrance exam, hierarchical college structure represented by college tiers, homogeneous secondary schools, and an established private tutoring market. While these features are particularly pronounced in Korea, similar characteristics exist in many other countries, which makes the insights of the paper's results broadly relevant.

2.1 High-Stakes College Entrance Exam

The College Scholastic Ability Test (CSAT), taken at the end of 12th grade, serves as the primary determinant of college admissions in Korea.[7] The exam's centrality to academic outcomes makes it a defining feature of the Korean education system. Its importance is reflected in nationwide accommodations on exam day - aircraft takeoffs and landings are suspended during English listening tests, while firms and government offices delay their workday to help students avoid traffic congestion.

Students receive both raw and standardized scores for subjects including Korean, Mathematics, English, and electives. Based on these scores and published college cutoffs, students apply to up to three colleges. Educational consulting firms release detailed predictions of admission cutoffs, which closely match actual outcomes, allowing students to target colleges matching their performance level.

This high-stakes examination system parallels college entrance mechanisms in many countries. China's Gaokao and France's Baccalauréat similarly serve as crucial determinants for elite college admission.[8] While the U.S. Scholastic Aptitude Test (SAT) also plays an important role, it differs in that other factors like grade-point-average and extra-curricular activities carry substantial weight in admissions decisions.

2.2 Hierarchical College Structure and College-Tier

College quality in Korea, as in many countries, exhibits a clear hierarchical structure with substantial implications for labor market outcomes. Empirical studies consistently

[7] In South Korea, the share of seats allocated through holistic review, in which test scores are not the only determinant of admission, has increased in recent years. In 2019, 24.9% of students were admitted through this route (Bastedo, 2021). The working sample of this paper predates this change.

[8] Other examples include Turkey's Yükseköğretim Kurumları Sınavı, Brazil's Exame Nacional do Ensino Médio, and Malaysia's Sijil Pelajaran Malaysia.

show that graduating from an elite college significantly affects future earnings and career trajectories across various national contexts.[9]

Korea's college hierarchy has remained remarkably stable over time (Kim and Lee, 2006; Kim, 2014), enabling clear categorization of institutions into distinct tiers based on admission selectivity and outcomes.[10] Based on the "cutoff sheet" published by Uway (2011), one of the major education consulting firms, I classify colleges into four ordered tiers. Tier 1 comprises the most selective universities, the top 1% of the CSAT distribution; the Tier 2 and Tier 3 cutoffs lie at approximately the top 5% and 15%. Tier 4 consists of the remaining college graduates and high school graduates who do not attend college.[11] The specific member universities of each tier are detailed in Appendix B. The stability of these status distinctions, coupled with their significant impact on labor market outcomes (which I document in Section 3), creates strong incentives for competitive investment in college admission preparation.

2.3 Homogeneous Secondary School

Secondary schools in Korea are highly homogeneous in both curriculum and quality, while the private tutoring market is well developed.[12] This combination creates a unique environment for studying how household resources translate into educational outcomes.

The homogeneity of schools stems from multiple policy features. First, the curriculum is uniform and under strict government control, applying to both public and private schools. Even private schools have limited autonomy in curriculum and tuition decisions,[13] which makes them very different from private secondary schools in the United States. Second, successive school-equalization policies have led to similar education quality across schools.[14] Student assignment to middle schools occurs through a random lottery within residential districts, mitigating selection effects.[15] The only ex-

[9] See, for example, Hoekstra (2009) for the United States, MacLeod *et al.* (2017) for Colombia, Zimmerman (2019) for Chile, Anelli (2020) for Italy, Sekhri (2020) for India, and Jia and Li (2021) for China.

[10] Notably, the government has strictly regulated enrollment quotas for universities in the greater Seoul area, where most prestigious universities are concentrated, since the 1980s.

[11] Within Tier 4, the early-career earnings difference between college attendees and non-attendees is not statistically different from zero, unconditionally ($+0.058$ log points, $t = 1.6$) or conditional on CSAT (-0.044, $t = -1.5$); non-attendance also concentrates well below the Tier 3 cutoff.

[12] On the other hand, competition between private and public schools is extensively studied in the US setting. See Epple and Romano (1998); Epple, Figlio and Romano (2004); Epple and Romano (2008); Epple, Romano and Urquiola (2017, 2021).

[13] One of the few decisions private secondary schools in Korea can make independently is teacher hiring. Park, Behrman and Choi (2013) provide evidence that the difference in teacher quality between private and public secondary schools in Korea is not significant.

[14] See Section II of Kim and Lee (2010) for a description of the history of the school equalization policy. As of 2010, the high school equalization policy has been adopted in all major cities in South Korea.

[15] Papers in the literature exploit this random assignment feature to estimate the effects of various independent variables of interest on educational outcomes. See, for example, Kang (2007), Park, Behrman and Choi (2013), and Park, Behrman and Choi (2018).

ception to this central assignment is specialized high schools, which account for just 3% of total enrollment.

2.4 Private Tutoring Market

Alongside the relatively homogeneous school system exists an established private tutoring market, responsible for 2.8% of GDP (Nam, 2007). Korean parents spend on average 9% of their income on private tutoring, primarily through hagwon (cram schools), one-on-one tutoring, group tutoring, and online classes.[16] Hagwons account for the largest share of enrollment and spending, especially at the middle- and high-school levels. Providers are densely located around residential and school neighborhoods, giving students broad access to supplementary instruction, although quality varies with price. With the centralized school curriculum, tutors can specialize in preparing students for school and entrance exams. Tutoring therefore acts as an effective substitute for parental time in academic support.

This institutional setting, combining relatively homogeneous schools with an extensive and readily accessible private tutoring market, provides a transparent environment for studying how household resources translate into educational outcomes. The limited variation in school quality helps isolate the effects of household investments, while tutoring offers a clear measure of parental monetary investment.

3 Descriptive Evidence

3.1 Data: Korean Educational Longitudinal Study 2005

The Korean Educational Longitudinal Study 2005 (KELS) provides comprehensive data for this paper's analysis. It contains measures of the determinants and outcomes of parental investment and student effort: parental income and education, private tutoring expenditure, students' time allocation, standardized test scores, and non-cognitive traits. This is an unusually rich combination within a single longitudinal dataset.

In particular, KELS tracks a nationally representative sample of 6,908 students (1% of the country's 703,914 7th graders) from 2005 through 2020. The survey consists of two stages, annual surveys during secondary school (2005-2010) and biennial surveys during college and early career (2012-2020). Household income and private tutoring expenditure are collected each year. Monetary values are in units of 10,000 KRW, approximately USD 9 at the average exchange rate over the 2005–2010 survey waves. The hours spent in tutoring activities and the hours spent on self-study are collected as a weekly average. Academic performance in primary school is measured as an ordered

[16] See Bray (1999, 2021) for a comprehensive cross-country comparison of private tutoring.

Table 1: Sample Moments

School grade	7th		8th		9th	
	Mean	Stdev	Mean	Stdev	Mean	Stdev
Tutoring Expenditure	27.41	24.25	25.95	24.13	33.34	36.94
Hours of Self-Study	5.43	5.33	5.90	5.20	6.27	5.37
Hours of Tutoring	10.34	8.76	8.74	7.50	9.02	9.58
HH Income	425.58	222.07	413.69	202.45	443.98	246.48
Tutoring Participation	0.795	0.404	0.774	0.418	0.708	0.455
Number of Tutoring Centers	493	429	530	458	560	478
Test Scores	422.00	61.53	421.41	63.38	422.21	63.13
Observations			3,524			

School grade	10th		11th		12th	
	Mean	Stdev	Mean	Stdev	Mean	Stdev
Tutoring Expenditure	38.80	39.71	32.92	38.36	27.78	42.16
Hours of Self-Study	7.18	5.68	7.98	6.14	13.31	9.39
Hours of Tutoring	7.04	6.60	5.83	6.32	5.02	7.00
HH Income	426.17	240.37	397.43	244.85	374.26	216.96
Tutoring Participation	0.742	0.437	0.657	0.475	0.509	0.500
Number of Tutoring Centers	632	497	654	503	756	587
Test Scores	–	–	–	–	408.67	66.38
Observations			3,524			

	Mean	Stdev
Parental education match		
Both parents college	0.173	0.378
Father college only	0.172	0.377
Mother college only	0.013	0.115
Neither college	0.642	0.479
6th grade Test Score	417.44	59.85
Early Labor Market Income	217.11	75.12
Observations	3,524	

Note: Means and standard deviations over the estimation sample. Tutoring expenditure and household income are in 10,000 KRW per month, approximately USD 9 per unit at the 2005–2010 average exchange rate (income top-coded at a common bracket across waves); hours of tutoring and self-study are weekly; tutoring centers are counted within 5 km of the student's school. Tutoring expenditure, tutoring hours, and self-study hours are unconditional: non-participants enter as zeros, the convention used in the model tables. In survey waves 1–3, expenditure is recorded only for participants, and zeros are imputed for observed non-participants. Conditional-on-participation means equal the unconditional means divided by the participation rate. Early labor-market income is monthly, in 10,000 KRW.

discrete measure answered by the household. For 7th to 9th grades, the administrative test scores are of achievement tests standardized at the national level. For 12th grade, the administrative College Scholastic Ability Test (CSAT) score is available. Table C.1 details the sample selection criteria and their effects.

KELS also measures non-cognitive skill with an adolescent self-concept battery: identical items on social acceptance, family relations, and physical appearance at grades 7, 8, 9, and 12, supplemented by two peer-relations measures at grades 10 and 11. These self-reports are the age-appropriate counterpart to the mother-reported behavior-problem scales of CNLSY-based work (Cunha, Heckman and Schennach, 2010).

Table 1 presents sample moments from KELS. A notable pattern is that, while average self-study hours increase over age, tutoring hours show a decreasing trend.[17] Household income moments remain stable throughout the sample period. Parental education, collected in the first year, is assumed constant given the relatively short time frame.

I supplement KELS with two additional data sources: (i) local variation in tutoring supply and (ii) labor-market outcomes later in life. First, I combine Statistics Korea's Census on Establishments with administrative records on school locations. The establishment census reports the number of hagwons in each eup-myeon-dong (neighborhood-level administrative area) in each KELS wave. Using KELS school identifiers, I geocode schools and nearby administrative-area centroids and construct the number of hagwons located within a 5-kilometer radius of each student's school. I then link each student to the year-specific local hagwon count and classify observations into quintiles of local tutoring supply. This measure reflects the availability of tutoring providers near each school and serves as a proxy for households' costs of finding suitable tutoring.

Second, because KELS observes individuals only at the beginning of their labor-market careers, I use the Korean Labor and Income Panel Study (KLIPS) to estimate age–earnings profiles that vary by college tiers. These profiles help identify the model's lifetime return to college placement rather than treating income observed in KELS at young ages as a measure of lifetime earnings. Details about KLIPS are provided in Appendix C.

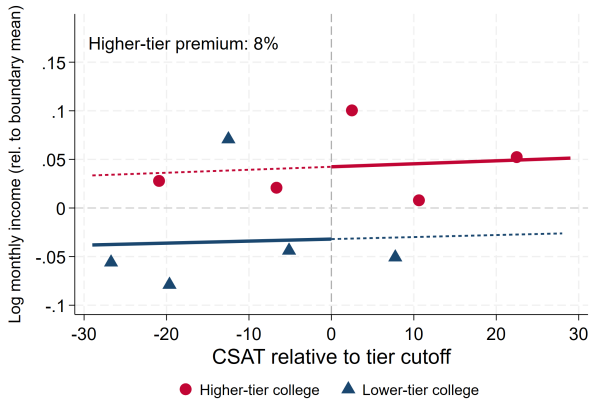
3.2 Elite College Premium

College tier is strongly associated with early-career earnings, even after controlling for CSAT scores. Panel (A) of Figure 2 summarizes three adjacent-tier comparisons: Tiers 1–2, Tiers 2–3, and Tiers 3–4. I center CSAT scores at each pair's admission boundary and stack the three comparisons. The two series therefore represent the higher- and lower-ranked colleges within each pair, rather than two fixed college tiers. Across

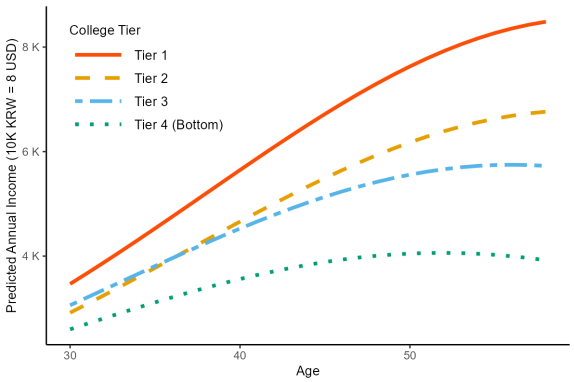
[17] I discuss this empirical pattern and its implications on structural estimation further in Section 3.3.

Figure 2: College Tiers and Earnings

(A) Early earnings



(B) Lifetime earnings



Note: Panel (A) stacks three adjacent-tier comparisons: Tier 1 versus Tier 2, Tier 2 versus Tier 3, and Tier 3 versus Tier 4. Within each comparison, CSAT scores are centered at the corresponding tier boundary, so zero denotes that boundary. The centered observations are then pooled to plot binned mean log monthly income at labor-market entry (KELS, age $\approx$ 28). Circles denote the higher-ranked tier within each pair, and triangles denote the lower-ranked tier. Thus, each plotted series combines different college tiers across the three comparisons. Solid segments show higher-tier attendance above the boundary and lower-tier attendance below it. Dotted segments show the reverse. Income is winsorized at the 1st and 99th percentiles. The figure describes earnings differences by attended tier at similar relative CSAT scores. Panel (B): predicted annual income (10,000 KRW, real 2015 base) by college tier over the life-cycle (ages 30–58), from the KLIPS pooled-OLS estimates in Panel B, Column (1) of Table C.3 (Tier 4 = baseline; birth year 1992).

these comparisons, higher-tier attendance is associated with about 8% higher earnings at similar scores relative to the corresponding boundary. Panel A of Table C.3 reports the corresponding regression analysis, with Tier 4 as the reference category. Without controls, Tier 1 graduates earn about 0.34 log points (roughly 40%) more than Tier 4 graduates (Column 1); controlling for CSAT and field of study, the premium remains a significant 0.17 log points (about 18%, Column 3).[18] These tier differentials are already visible at labor-market entry, when the KELS cohort is about 28 years old.

College tier is also associated with differences in alumni's earnings growth. Using KLIPS, I estimate a pooled-OLS regression of log annual income on tier dummies interacted with age. Columns (1)–(3) of Panel B of Table C.3 provide the OLS estimates for the regression equation,

$$\ln y_{it} = \sum_{j=1}^{J-1} (\beta_j + \delta_j age_{it}) D_{i,j}^{Tier} + Z_{it}\gamma + \varepsilon_{it}^{y} \tag{1}$$

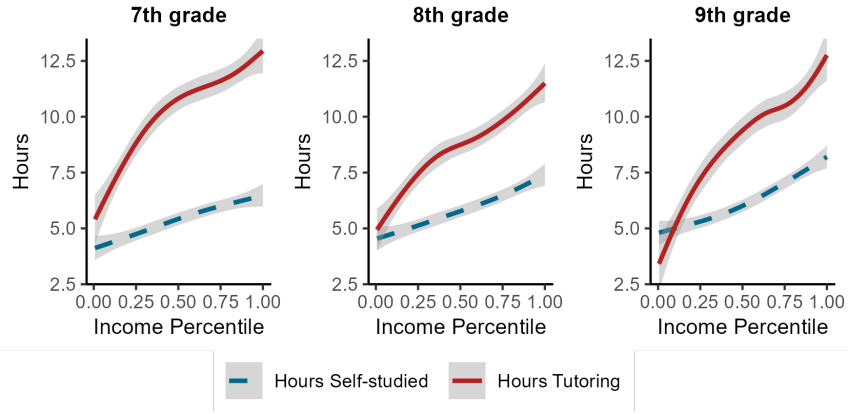
where $D_{i,j}^{Tier}$ is a dummy variable indicating that person i graduated from a tier j college,

[18] To the best of my knowledge, among studies using Korean datasets, these are the first estimates of the early-career earnings premium associated with college tier that condition on CSAT scores. The CSAT- and field-adjusted premium is smaller than the unconditional tier gradient reported in earlier work, including the 50.5% Tier 1–bottom-tier difference reported by Lee and Koh (2023).

and Z_{it} is the set of explanatory variables including age, squared age, birth year, and gender of person i.[19] Panel (B) of Figure 2 plots predicted annual income by college tier using Column (1) of Panel B in Table C.3. The earnings gaps widen substantially with age. The Tier 1–age interaction remains positive and statistically significant when entrance scores and college major are added in Columns (2) and (3), although these specifications use a much smaller subsample with observed entrance scores. Taken together, the descriptive evidence indicates that earnings differences by college tier are visible at labor-market entry and widen over the life cycle.

3.3 Income Gradient of Student's Self-Efforts

Figure 3: Income Gradient in Effort Decision



Source: Korea Educational Longitudinal Study 2005
Note: The gray regions are confidence bands with a significance level of 0.05. I include only households who do not have missing information on the following variables: tutoring expenditure, CSAT scores, and household income.

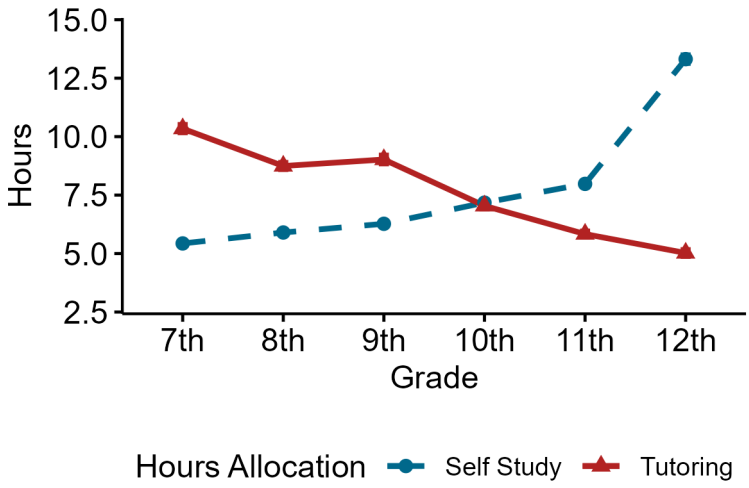
Hours of self-study are less responsive to parental income than tutoring hours, suggesting that children's own study effort is an important channel to account for when quantifying the consequences of parental investment. Figure 3 plots tutoring and self-study hours against parental income for students in grades 7–9 using local linear regression. The tutoring-hours gradient is substantially steeper than the self-study gradient. Additional regression evidence in Appendix C shows that the covariation between parental income and self-study falls substantially once parental education is controlled for, and is small within households: the income coefficient drops to 0.15 with parental-education controls and to 0.036 with household fixed effects. The corresponding income gradient in tutoring hours, on the other hand, remains similar. Together, these patterns indicate that purchased tutoring is more responsive to household resources than is self-study.

Differences in observable household characteristics can therefore generate different allocations between parental investment and child effort. Omitting either margin may

[19]The purpose of the birth year dummy variable is to capture the cohort difference in workers' income.

bias the model's quantification of the effects of parental investment on child outcomes, household welfare, and intergenerational mobility.

Figure 4: Student's Time Allocation by Grades



Source: Korea Educational Longitudinal Study 2005.
Note: I include only households who do not have missing information on the following variables: tutoring expenditure, CSAT scores, and household income.

Students' time allocation and tutoring participation change substantially as they progress through secondary school, which motivates a framework that allows both the returns to investment and selection into tutoring participation to evolve over time. Figure 4 shows how average self-study and tutoring hours vary by grade. Tutoring hours decline with grade, but self-study hours increase. By 12th grade, students spend almost three times as many hours on self-study as on tutoring.

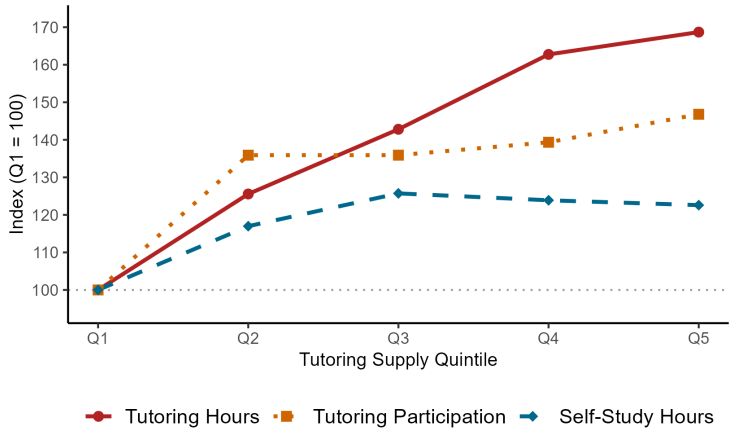
Tutoring participation also falls from 80% in 7th grade to 51% in 12th grade. Thus, later-grade tutoring hours are observed among an increasingly selected group of students. Ignoring this extensive-margin selection may bias estimates of the returns to tutoring, particularly during late adolescence. The changing allocation between self-study and tutoring also suggests that their marginal effects on academic outcomes may vary by grade.[20] Together, these patterns motivate a dynamic framework for estimating the production function and household behavior.

3.4 Local Tutoring Supply and Substitution

Local variation in private-tutoring supply affects households' educational investment choices. The number of nearby *hagwons* (private tutoring centers) varies substantially across students' school neighborhoods, from rural areas with few options to

[20]Several studies report that the effects of parental investment decline with children's age (Cunha, Heckman and Schennach, 2010; Del Boca, Flinn and Wiswall, 2014; Del Boca, Monfardini and Nicoletti, 2017).

Figure 5: Hours Investment by Local Tutoring Supply



Note: Supply quintiles are based on the log of the average number of hagwons within 5 km of the student's school in grades 7–12. The corresponding mean numbers of hagwons are 42, 279, 645, 949, and 1,470 for supply quintiles $r = 1, \ldots, 5$. Both series are normalized to 100 at the lowest supply quintile. Source: KELS 2005 merged with Statistics Korea's Census on Establishments.

dense urban markets with hundreds of providers within 5 km. This variation changes the local choice set and the search costs that households face when purchasing tutoring.[21] Figure 5 presents how tutoring and self-study vary across local-supply quintiles. Students in high-supply areas shift their time heavily toward tutoring relative to self-study, a pattern consistent with substitution between the two inputs.

I incorporate this geographic variation by allowing tutoring-participation and search costs to differ across local hagwon-supply quintiles. The supply measure therefore provides an observed cost shifter that helps distinguish access to tutoring from household resources and preferences.[22]

4 Dynamic Model of College Admissions Tournament

Motivated by the competition-driven nature of parental investment and the descriptive evidence that college tiers generate different lifetime earnings profiles, I develop a dynamic tournament model of college admissions competition. The framework builds on the rank-order tournament literature pioneered by Lazear and Rosen (1981) and its applications to college admissions (Grau, 2018; Tincani, Kosse and Miglino, 2022).

The model inherits the competition framework of rank-order tournaments. Households compete for a fixed set of ordered prizes, and effort responds to the prize spread and to proximity to the thresholds separating them. Applying the tournament model

[21] While local tutoring density likely affects both the physical choice set, such as provider variety and travel distance, and the search costs of finding a suitable match, the data do not allow for the separate identification of these mechanisms. The regional supply variable should therefore be interpreted as capturing the net effect of market thickness on the effective cost of tutoring.

[22] This approach has a similar spirit with the identification logic of Caucutt, Lochner, Mullins and Park (2026), which I discuss further in the estimation section.

to the education setting departs from its standard version in three ways. Effort is productive as well as positional, since tutoring and self-study also build human capital rewarded independently of rank. The prize structure is not designed by a principal. Tier premia are set by the labor market and seat counts by regulation.

The economy consists of N households, indexed by i, that make educational investments over decision periods $t = 1, \ldots, M$. In each period, household i chooses tutoring participation (d_{it}), tutoring quality (p_{it}) and tutoring hours (x_{it}) conditional on participating, and the child's self-study time (s_{it}). These choices affect skill accumulation and ultimately the child's rank in the college-admission tournament. College tiers are assigned on the basis of final rank. The payoff from a given tier is nevertheless household-specific. Conditional on tier assignment, the child's lifetime earnings also depend on accumulated cognitive and non-cognitive skills. Households therefore compete for college-tier opportunities through their educational investment choices while valuing those opportunities differently according to their children's states.

I model the household as a unitary decision maker. This abstraction allows the model to focus on competition across households rather than bargaining or contracting within households. It is particularly useful in this setting, where both parents and students have strong incentives to perform well on a high-stakes college entrance examination. This assumption does not imply that adolescents are passive or that within-family conflicts are absent. Rather, the allocation $\{d_{it}, p_{it}, x_{it}, s_{it}\}_{t=1}^{t=M}$ can be viewed as the reduced-form outcome of the intra-household process that determines tutoring and self-study time.[23] The model therefore characterizes the household's resulting educational investment choices.

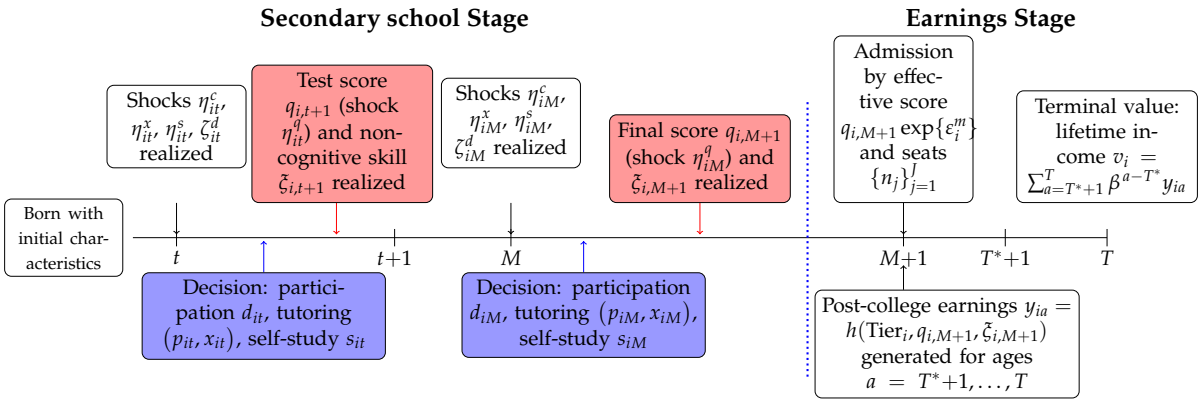
4.1 Timeline

Figure 6 illustrates the timeline of the model. The model has two stages: a secondary school stage during which households choose educational investments, and a post-school earnings stage during which labor market outcome is realized. The schooling stage begins when the student enters 7th grade. Each household is endowed with an income stream $\{w_{it}\}_{t=1}^{M}$, parental education m_i, initial test score q_{i1}, and non-cognitive skill ξ_{i1}. Households also face a local tutoring-supply environment $\{r_{it}\}_{t=1}^{M}$, where r_{it} denotes the supply quintile determined by the number of tutoring centers near the student's school. Households also differ in permanent unobserved type k, summarized by $(\lambda_k^c, \lambda_k^x, \lambda_k^s)$. These type-specific components affect the marginal utility of consumption, the taste costs of tutoring and self-study, respectively.

At the beginning of each schooling period t, the household observes shocks to the

[23] Del Boca, Flinn, Verriest and Wiswall (2026) study parent–child decision making by modeling parental conditional cash transfers that incentivize children's self-investment.

Figure 6: Model Timeline



Note: "Initial characteristics" comprise the parental income stream $\{w_{it}\}_{t=1}^M$, parental education m_i, the initial test score q_{i1}, the initial non-cognitive skill ξ_{i1}, the local tutoring-supply environment $\{r_{it}\}_{t=1}^M$, and the permanent unobserved type k, summarized by $(\lambda^c_k, \lambda^x_k, \lambda^s_k)$. In each schooling period $t = 1, \ldots, M$ the household observes its preference and participation-cost shocks, decides whether to participate in tutoring and, if so, the price and hours (p_{it}, x_{it}), chooses self-study s_{it}, and the next-period test score and non-cognitive skill are produced. After period M, college tiers are assigned by the ranking of the effective score $q_{i,M+1}e^{\varepsilon^m_i}$ given the fixed seat vector $\{n_j\}_{j=1}^J$; earnings begin after college graduation at age T^*+1 and depend on the assigned tier, the final score, and terminal non-cognitive skill $\xi_{i,M+1}$.

marginal utility of consumption, tutoring, and self-study, $(\eta^c_{it}, \eta^x_{it}, \eta^s_{it})$. These shocks capture time-varying unobserved components that are not explained by observed states or permanent unobserved heterogeneity. Conditional on these shocks and the observed state variables, the household chooses tutoring quality p_{it}, tutoring hours x_{it}, and self-study hours s_{it} to maximize its value function, subject to the budget and time constraints. The next-period test score $q_{i,t+1}$ is then produced after the realization of the test-score shock η^q_{it}. This process repeats until the final test score $q_{i,M+1}$ is generated at the end of high school, $t = M$.

Students are assigned to college tiers according to the fixed number of seats in each tier and the ranking of their effective admission score, defined as $\ln q_{i,M+1} + \varepsilon^m_i$: the log of the final test score plus an idiosyncratic matching shock. Let n_j denote the number of seats in Tier j. The students with the top n_1 effective scores are assigned to Tier 1, the next n_2 students to Tier 2, and so on until Tier $J-1$. The bottom tier is the residual tier, consisting of students below the cutoff for Tier $J-1$ and students who do not attend college.[24] The matching shock captures idiosyncratic admission factors, including preferences over college location, that may lead students with similar raw test scores to attend different tiers. The assigned college tier affects subsequent earnings, which also depend on the final test score as a measure of cognitive skill, non-cognitive skill, and an idiosyncratic income shock, as described below.

4.2 The Preliminaries of the Tournament

[24]The implicit assumption for the bottom tier is that all students graduate from high school. The high school dropout rate in South Korea is less than 2 percent.

4.2.1 Lifetime Income: The Prize of the Tournament

The prize for going to a more prestigious college tier is a higher expected lifetime income awarded to the student, which motivates the household to make investments. For a household i, college tier j gives expected lifetime income v_{ij}. The tier-specific lifetime income v_{ij} is the discounted sum of the predicted income of the graduates. In particular,

$$v_{ij} = \sum_{a=T^*+1}^{T} \beta^{a-T^*} \hat{y}_{ija} \tag{2}$$

where $\hat{y}_{ija}$ is the predicted income of the student in household i at age a if she graduates from tier j, T^* is the age when the student graduates from college, T is the retirement age, and β is the discount factor fixed at 0.95.[25]

Post-college annual log earnings are given by

$$\begin{aligned} \ln y_i = {} & a_{0y} + b_{z0} \ln q_{i,M+1} + b_h \xi_{i,M+1} + b_{a1} \widetilde{\text{age}} + b_{a2} \widetilde{\text{age}}^2 \\ & + \sum_{j=1}^{J-1} D_{ij} \left(\tau_j + b_{a1,j} \widetilde{\text{age}} \right) + \sigma_y \varepsilon_i^y(\text{age}) \end{aligned} \tag{3}$$

where D_{ij} indicates attendance at Tier j, $\widetilde{\text{age}} = \text{age} - \text{age}^*$ is years since labor-market entry, and Tier 4 is the reference category. The parameter τ_j is the tier premium at entry, while $b_{a1,j}$ captures the difference in earnings growth relative to Tier 4. The age profile has a common linear term b_{a1} and quadratic term b_{a2}. The parameters b_{z0} and b_h are the earnings returns to the final test score and non-cognitive skill, respectively, and $\varepsilon_i^y(\text{age})$ is a standard normal income shock.[26] The income equation (3) allows the tier prize v_{ij} to vary across individuals through the final CSAT score $q_{i,M+1}$ and non-cognitive skill $\xi_{i,M+1}$.[27]

4.2.2 College Seats, Tiers, and Cutoffs

Lifetime income is highest in Tier 1 and decreases across lower-ranked tiers:

$$v_{i1} > v_{i2} > \cdots > v_{i,J-1} > v_{iJ}.$$

[25] The average interest rate was around 5% in South Korea in 2010.

[26] The age-profile parameters are identified by moment conditions from KLIPS described in Appendix C.

[27] I restrict tournament prizes to pecuniary returns. Gong, Lochner, Stinebrickner and Stinebrickner (2019) estimate college consumption benefits using an Euler-equation approach and detailed data on consumption, desired borrowing, earnings expectations, and risk and time preferences. Applying this approach to admissions competition would additionally require identifying how these benefits differ across college tiers.

Student i is assigned to Tier j if and only if

$$\ln \widetilde{Q}_{j-1} > \ln q_{i,M+1} + \varepsilon_i^m \geq \ln \widetilde{Q}_j$$

where ε_i^m is the matching shock and $\ln \widetilde{Q}_j$ is the log-score cutoff between Tiers j and $j+1$. Let n_j denote the number of seats in Tier j and $N_j = \sum_{\ell=1}^{j} n_\ell$ the cumulative number through Tier j. The log cutoff $\ln \widetilde{Q}_j$ is the N_jth-highest effective log score in the cohort, so exactly N_j students attain Tier j or higher.

The cutoff vector $\{\widetilde{Q}_j\}_{j=1}^{J-1}$ is the channel through which competition enters the model. To obtain Tier j or a higher tier, a student must outperform enough competitors to rank among the N_j highest final scores. Because final scores are determined by households' sequential investment choices through period M, the cutoff vector is endogenously determined by dynamic competition across households.

The value of obtaining a higher tier should be interpreted jointly with the earnings equation (3). If tier effects are small relative to the direct returns to cognitive and noncognitive skills, incentives to compete for rank are weak and direct human-capital returns dominate households' investment choices. The maintained ordering of tier prizes could fail if mismatch costs were sufficiently large (Arcidiacono, Aucejo, Coate and Hotz, 2014). This possibility appears less relevant in the Korean setting where college completion rates are high. The earnings estimates below provide an empirical assessment of the ordering of tier returns.

4.2.3 Investment Choices of the Households

Parental Investment: One of the two modes of household investment is parental monetary investment, namely private tutoring. To model the expenditure choice jointly with the hours decision, I assume that each household chooses the unit price (quality) of tutoring p_{it} and tutoring hours (quantity) x_{it} to increase the child's test score.[28] The total amount of tutoring expenditure e_{it} is

$$e_{it} = p_{it} x_{it}.$$

Tutoring expenditure is subject to two constraints. A household cannot spend more on tutoring than its income (i.e., $e_{it} \leq w_{it}$).[29] In addition, tutoring hours are bounded by the child's maximum available time, denoted by h. While the income constraint varies across households, the child's available time is constant across households.

[28] In the data, expenditure is measured per month and tutoring hours per week, so $p_{it} = e_{it}/x_{it}$ is monthly spending per weekly hour of tutoring, which is proportional to, but about 4.3 times, an hourly price. Since p_{it} enters the model as a choice variable with its own scale, the constant of proportionality is absorbed by the cost and taste parameters; all tables report p_{it} in these units.

[29] I assume no borrowing or saving.

Note that the time choice is solely about the time use of the child, which means I do not model the time allocation of parents. The data suggest that the majority of parents do not teach their children themselves in middle school periods, and very few parents use their time to teach their child in the high school periods. A few potential explanations can be given for this empirical fact. As students grow, the test materials become more and more difficult to be taught by parents. Also, if there exists an established tutoring market, it would be a safer option for parents in terms of increasing student's test score. Note that the model concerns a regime with a high-stakes standardized test where full-time tutors would have a comparative advantage in preparing students for exams over parents.

Child's hours of self-study: Hours of self-study is the other household's mode of investment in the tournament. Each household chooses the child's self-study hours s_{it} jointly with tutoring hours x_{it}, subject to the time constraint $x_{it} + s_{it} \leq h$. Unlike parental monetary investment, the time endowment available for self-study does not vary across households because each child is allocated the same amount of time.

The cost of self-study, however, can vary substantially with observed student and household characteristics. I allow the utility component from child self-study (α_s) to vary across four parental-education configurations, depending on whether neither parent, only the father, only the mother, or both parents hold a college degree.[30]

I also allow for unobserved heterogeneity in preferences for self-study. Some students may prefer studying independently to reviewing the same material with tutors, while others may prefer working with tutors to studying alone. I capture this variation through unobserved types in the taste for self-study, namely λ_s.

4.2.4 Test Score and Non-Cognitive Skills Production Function

The final test score is the result of accumulated dynamic choices of the household along with its given initial conditions. The initial academic performance q_{i1} is exogenously given and proxied by academic performance in primary school. Although an earlier measure of the initial child's ability would be more desirable, this is the earliest time period that the academic performance data of KELS are available. The three choices affecting test scores are quality of tutoring p_{it}, hours of tutoring x_{it}, and hours of self-study s_{it}. I allow that the quantity (hours) and quality (unit price) of the tutoring activity have different intensities in contributing to the test score production. Denoting κ as intensity of quality of tutoring, the transformed tutoring input is specified as

$$\tilde{e}_{it} = p_{it}^{\kappa} x_{it}^{1-\kappa}. \tag{4}$$

[30]Related to this, Del Boca, Flinn, Verriest and Wiswall (2026) document a positive gradient between parental education and children's self-investment time.

Allowing for the separate contribution of quality and hours captures their distinct roles in producing subsequent test scores. The split parameter κ is common across grades.

For each time $t = 1, 2, ..., M$, the test score $q_{i,t+1}$ is produced following

$$q_{i,t+1} = g(\theta_t^q, q_{it}, p_{it}, x_{it}, s_{it}, \eta_{it}^q)$$

where η_{it}^q is the test score shock, and θ^q is the set of relevant parameters for the test score production. The inclusion of the test score produced in the previous period, q_{it}, allows that the previous test score has its own effects in generating subsequent test score (Cunha and Heckman, 2007). Furthermore, I allow a subset of production parameters to vary across grades to capture changes in the effectiveness of tutoring and self-study. As coursework becomes more advanced, a given amount of educational investment may generate smaller gains in test scores, making tutoring and self-study less productive in later grades. In addition, the relative importance of each investment may vary by student's grade level, as suggested in Figure 4. To capture these changing effects, I let the marginal effects parameters δ_{qt}, ν_t, δ_{et}, and δ_{st} be different for each period $t = 1, 2, \ldots, M$.

For estimation, I parameterize the production function g as a Constant Elasticity of Substitution (CES) production function, which is specified as

$$q_{i,t+1} = A_t q_{it}^{\delta_{qt}} \left[\delta_{et}(1 + \tilde{e}_{it})^{\phi} + \delta_{st}(1 + s_{it})^{\phi} \right]^{\frac{\nu_t}{\phi}} \exp\{\eta_{it}^q\} \tag{5}$$

where A_t is total factor productivity, ν_t is the parameter of marginal effect of the combined effort choices, and ϕ is the parameter governing substitution between tutoring and self-study. The marginal effect of the total effort decision is captured by ν_t, while the relative importance of the tutoring expenditure and hours of self-study are captured by δ_{et} and δ_{st}, respectively. The test-score shock enters in levels as $\exp\{\eta_{it}^q\}$, so η_{it}^q is the shock to the log score.

Household choices also affect the student's non-cognitive skill, ξ_{it}. Time devoted to tutoring and self-study may influence non-cognitive development through academic pressure or changes in opportunities for social interaction. I allow the effects of these two inputs to differ across grades:

$$\xi_{i,t+1} = \gamma_{\xi,0} + \gamma_{\xi,1,t}\xi_{it} + \gamma_{\xi,x,t}\log(1 + x_{it}) + \gamma_{\xi,s,t}\log(1 + s_{it}) + \varepsilon_{\xi,it}, \tag{6}$$

where $\varepsilon_{\xi,it}$ is an idiosyncratic shock to non-cognitive skill formation at time t. Because non-cognitive skill enters the earnings equation (3), households account for these effects alongside the gains in test scores and admission prospects when choosing tutoring and self-study.

4.3 Preferences

Flow Utility: The flow utility of the unitary household has four components: household consumption c_{it}, the child's tutoring hours x_{it}, the child's self-study hours s_{it}, and the child's residual leisure time $\ell_{it} = h - x_{it} - s_{it}$. I denote the corresponding baseline taste parameters by α_c, α_x, α_s, and α_ℓ, respectively.

I do not restrict the signs of α_x and α_s. Thus, tutoring and self-study may provide either positive or negative flow utility to the household. Because the associated taste shocks are positive, they change the magnitude but not the sign of these utility components. These terms represent the gross utility of the unitary household, and I remain agnostic about how the utility costs and benefits of the child's time allocation are distributed among individual household members.

I allow the utility coefficients associated with the child's time allocation to depend on parental education. Let $m_i \in \{0, F, M, B\}$ indicate that neither parent, only the father, only the mother, or both parents hold a college degree, and let $g_t \in \{ms, hs\}$ denote the middle-school ($t \leq 3$) and high-school ($t \geq 4$) stages. For $z \in \{x, s\}$, define

$$\alpha_{it}^z = \alpha_z \exp\{\mu^z(m_i, g_t)\}, \qquad \mu^z(m_i, g_t) = \begin{cases} 0 & \text{if } m_i = 0, \\ \mu^z_{F,g_t} & \text{if } m_i = F, \\ \mu^z_{M,g_t} & \text{if } m_i = M, \\ \mu^z_{B,g_t} & \text{if } m_i = B. \end{cases} \tag{7}$$

The normalization $\mu^z(0, g_t) = 0$ makes households in which neither parent holds a college degree the reference group; the shifters act as log multipliers on the taste weights and are allowed to differ between the middle-school and high-school stages.

Households also differ in their access to tutoring providers. As discussed in Section 3.4, I divide local tutoring markets into five supply quintiles, $r \in \{1, \ldots, 5\}$, based on the number of tutoring centers within 5 kilometers of the student's school. Quintile 1 has the lowest local supply and quintile 5 the highest. I allow the cost of obtaining higher-quality tutoring to vary across these local markets. In particular, define

$$\alpha_r^c(p_{it}) = \alpha_c \exp\{-\Gamma_r(p_{it})\}, \tag{8}$$

$$\Gamma_r(p_{it}) = \gamma_1 p_{it} + \gamma_{2r} p_{it}^2 \tag{9}$$

where γ_1 is common across local markets and γ_{2r} varies across tutoring-supply quintiles. The search-cost term applies only when the household purchases tutoring. In estimation, the linear coefficient γ_1 is not separately identified from the participation hurdle introduced below and is fixed at zero, so the estimated search cost operates through the quintile-specific curvature γ_{2r}.

The household's flow utility $u(p_{it}, x_{it}, s_{it}, \varepsilon^u_{it})$ is

$$u(\cdot) = \begin{cases} \alpha^c_r \varepsilon^c_{it} \log(c_{it}) + \alpha^x_{it} \varepsilon^x_{it} \log(x_{it}) + \alpha^s_{it} \varepsilon^s_{it} \log(1 + s_{it}) + \alpha_\ell \log(\ell_{it}) & \text{if } x_{it} > 0, \\ \alpha_c \varepsilon^c_{it} \log(c_{it}) + \alpha^s_{it} \varepsilon^s_{it} \log(1 + s_{it}) + \alpha_\ell \log(\ell_{it}) & \text{if } x_{it} = 0. \end{cases} \tag{10}$$

Here, ε^c_{it}, ε^x_{it}, and ε^s_{it} are time-varying shocks to the utility from consumption, tutoring, and self-study, respectively, and $\varepsilon^u_{it} = (\varepsilon^c_{it}, \varepsilon^x_{it}, \varepsilon^s_{it})'$. I decompose these shocks as

$$\log \varepsilon^z_{it} = \lambda^z_k + \eta^z_{it}, \qquad z \in \{c, x, s\}, \tag{11}$$

where λ^z_k is the type-specific component and η^z_{it} is the time-varying component. The vector of time-varying components satisfies

$$(\eta^c_{it}, \eta^x_{it}, \eta^s_{it})' \sim N(\mu_{\eta t}, \Omega^\eta), \qquad \mu_{\eta t} = (0,\, 0,\, \mu_s \mathbf{1}[t = M])', \tag{12}$$

so the mean vector is zero except for the grade-12 shift μ_s in the study-taste component reported in Table 3. Thus, conditional on permanent household type k, which captures persistent differences in preferences for consumption, tutoring, and self-study, the utility shocks are jointly log-normal. The test-score innovation is not a utility shock and is introduced separately in the test-score production process. Note that I do not specify the utility flow from the current test score. The role of the current test score is limited to the stepping stone for the final test score. The specification of future value is introduced with the recursive formulation at the end of the subsection.

Terminal Value: Expected lifetime income is the terminal value of the model, which drives the dynamic investments of the tournament model. Note that the tier-specific lifetime income v_{ij} defined in equation (2) is a function of the college tier, the final test score, and non-cognitive skills. The expected lifetime income is a weighted sum,

$$\int_{\varepsilon_y} \sum_{j=1}^{J} \left\{ \ln(v_{ij}) * \text{Prob}(\ln \tilde{Q}_{j-1} \geq \ln q_{i,M+1} + \varepsilon^m_i \geq \ln \tilde{Q}_j \Big| \Omega_{iM}) \right\} g(\varepsilon_y) \tag{13}$$

where $\text{Prob}(\ln \tilde{Q}_{j-1} \geq \ln q_{i,M+1} + \varepsilon^m_i \geq \ln \tilde{Q}_j \Big| \Omega_{iM})$ is the probability of getting into college tier j, and Ω_{iM} denotes the household's information set in the final schooling period. The randomness of the admission probability comes from the test score shock η^q_{it} and the idiosyncratic matching shock ε^m_i, with composite dispersion $\sigma_{\text{total}} = (\sigma^2_q + \sigma^2_a)^{1/2}$. Each student has a different probability of going to a college tier j as they have different characteristics affecting the evolution of the test scores. The disparity among students in terms of going to each college tier leads to the discrepancies in ex-

pected lifetime income, which generates heterogeneous incentives among households. A higher expected lifetime income leads to a larger terminal value for the household, which makes it more appealing for the parents to invest in the child.[31]

Budget and Time Constraints: The choices of the household are restricted by the budget and the time constraints. The budget constraint is given by

$$c_{it} + p_{it}x_{it} \leq w_{it} \tag{14}$$

where w_{it} is household income, and the time constraint is

$$x_{it} + s_{it} \leq h \tag{15}$$

where h is student's disposable time. I define h as the maximum time each student can use every week.[32]

State Variables: There are observed and unobserved state variables in the dynamic model. The two skill stocks q_{it} and ξ_{it} are the endogenous state variables that are actively changing based on the household decisions over time. Both skill stocks are known to the household. To the econometrician, q_{it} is observed through test scores, while the non-cognitive stock ξ_{it} remains latent and is measured with error. The set Z_{it} collects other observed state variables including parental education m_i, parental income w_{it}, and the quintile r_{it} of the number of tutoring centers within neighborhood (i.e., $Z_{it} = \{m_i, w_{it}, r_{it}\}$). The set of unobserved state variables Ψ_{it} includes the set of unobserved shocks and the type-specific heterogeneity: the taste shocks $(\eta^c_{it}, \eta^x_{it}, \eta^s_{it})$, the tutoring participation-cost shock ζ^d_{it}, and the type-k loadings on the consumption, tutoring, and study tastes. Based on the timeline, the time-varying innovations to the two skills are not unobserved state variables for the current period: the test-score shock and the non-cognitive skill innovation $\varepsilon_{\xi,it}$ are realized after the period's decisions. (i.e., $\Psi_{it} = \{\eta^c_{it}, \eta^x_{it}, \eta^s_{it}, \zeta^d_{it}, \lambda^c_k, \lambda^x_k, \lambda^s_k\}$).

As depicted in Figure 6, each household learns about the realization of the consumption shock η^c_{it}, the disutility shock to hours of tutoring η^x_{it}, the disutility shock to hours of self-study η^s_{it}, and the participation-cost shock ζ^d_{it} at the beginning of each period. However, it does not know the test score shock η^q_{it} or the non-cognitive skill innovation $\varepsilon_{\xi,it}$ before it makes a decision. Therefore, it makes a set of choices based on the expectation over η^q_{it}, $\varepsilon_{\xi,it}$, and the next period's shocks $\eta^c_{i,t+1}$, $\eta^x_{i,t+1}$, $\eta^s_{i,t+1}$, and $\zeta^d_{i,t+1}$, conditional on observed state variables and type-specific unobserved heterogeneity.

Each household chooses whether to participate in tutoring by comparing the max-

[31] Under the assumed normal distribution of the composite test-score and matching shock, each tier's admission probability is the difference between normal cumulative probabilities evaluated at its standardized upper and lower cutoffs. Appendix D presents the resulting terminal-value expression.

[32] I assume each student can use 9 hours every day for activities other than hours spent in regular school.

imized values with and without tutoring, given its current state and realized shocks. Participation incurs the per-period utility cost $\chi_{r_{it}} + \sigma_d \zeta_{it}^d$, where $\chi_{r_{it}}$ is a fixed participation hurdle that varies across local tutoring-supply quintiles, ζ_{it}^d is a standardized transitory participation-cost shock, and σ_d governs its dispersion. A lower realization of ζ_{it}^d makes participation more attractive. This hurdle is distinct from the search-cost function $\Gamma_r(p_{it})$, which varies with tutoring quality conditional on participation.

Household Value Function: Building upon the preceding model components, the household's problem for $t = 1, \ldots, M$ can be written recursively as

$$V_{it}(q_{it}, \xi_{it}, \mathbf{Z}_{it}, \mathbf{\Psi}_{it}) = \max_{\substack{d_{it} \in \{0,1\}; \\ a_{it} = \{c_{it}, p_{it}, x_{it}, s_{it}\}}} \left\{ u^{d_{it}}\left(c_{it}, x_{it}, s_{it}, \boldsymbol{\varepsilon}_{it}^u\right) - d_{it}\left(\chi_{r_{it}} + \sigma_d \zeta_{it}^d\right) + \beta E_t\left[V_{i,t+1}\left(q_{i,t+1}, \xi_{i,t+1}, \mathbf{Z}_{i,t+1}, \mathbf{\Psi}_{i,t+1}\right)\right] \right\} \tag{16}$$

where the choice vector $a_{it} = (c_{it}, p_{it}, x_{it}, s_{it})$ satisfies the budget and time constraints (14) and (15), with $p_{it} = x_{it} = 0$ when $d_{it} = 0$, and $u^{d_{it}}$ is the corresponding branch of the flow utility in equation (10). The expectation integrates over the shocks described in the timeline above: the current period's skill innovations and the next period's taste and participation-cost shocks.

The participation rule implied by equation (16) is

$$d_{it} = \mathbf{1}\left\{\widetilde{V}_{it}^1 - V_{it}^0 \geq \chi_{r_{it}} + \sigma_d \zeta_{it}^d\right\} \tag{17}$$

where $\widetilde{V}_{it}^1$ is the maximized gross value of the tutoring branch before paying the participation hurdle and V_{it}^0 is the maximized value of the no-tutoring branch.

At the final schooling period, the continuation term in equation (16) is the expected utility from the student's lifetime income:

$$\beta E_M\left[V_{i,M+1}\right] = \alpha_v \sum_{j=1}^{J} \ln\left(v_{ij}\right) \times \text{Prob}\left(\ln \tilde{Q}_{j-1} \geq \ln q_{i,M+1} + \varepsilon_i^m \geq \ln \tilde{Q}_j \,\middle|\, \Omega_{iM}\right) \tag{18}$$

where $\text{Prob}(\ln \tilde{Q}_{j-1} \geq \ln q_{i,M+1} + \varepsilon_i^m \geq \ln \tilde{Q}_j \mid \Omega_{iM})$ is the tournament object described in expression (13), and v_{ij} is the tier-specific lifetime income defined in equation (2), which depends on the assigned tier and on both terminal skills. Because the log-earnings specification is additively separable, the common returns to terminal skills factor out. This allows the tier components of the expected log prize to be multiplied directly by the marginal admission probabilities, which is detailed in Appendix D. The parameter α_v measures the weight that the household places on the child's future lifetime income relative to its current flow utility. The household thus trades off current

consumption and the time costs of tutoring and self-study against future skills and, ultimately, the child's expected lifetime income.

4.4 Tournament Equilibrium

Let μ_t denote the cross-sectional distribution of household states at the beginning of period t. I assume a continuum of households.

Assumption 1. *The initial distribution μ_1, the model primitives, and the distributions of idiosyncratic shocks are common knowledge. Conditional on household states, the shocks are independent across households.*

The initial distribution includes the joint distribution of initial cognitive and noncognitive skills, parental education, household wage paths, permanent household types, and regional tutoring-supply types. As the population is a continuum, the exact law of large numbers implies that individual uncertainty washes out in the aggregate. Households can therefore anticipate the terminal distribution of admission outcomes and the associated college-tier cutoffs.[33]

Let $\tilde{Q} = (\tilde{Q}_1, \ldots, \tilde{Q}_{J-1})$ denote a guessed vector of admission cutoffs. Given $\tilde{Q}$, households solve the dynamic programming problem in equations (16) and (18), yielding value functions $V_t(\cdot; \tilde{Q})$ and Markov policy rules $\pi_t(\cdot; \tilde{Q})$. Starting from μ_1, these policies generate the subsequent distribution of household states and, ultimately, the distribution of the terminal test score $q_{i,M+1}$.

With the matching shock ε_i^m, the admission index $\ln q_{i,M+1} + \varepsilon_i^m$ determines college-tier assignment. Given the fixed tier capacity shares, the induced distribution of this admission index generates a new vector of cutoffs, denoted $\mathcal{T}(\tilde{Q})$.

Definition 1. *A recursive tournament equilibrium consists of a cutoff vector $\tilde{Q}^*$, household value functions and policy rules, and a distribution path such that: (i) households optimize given $\tilde{Q}^*$; (ii) the distribution path is generated by the resulting policy rules; and (iii) the guessed cutoffs clear the admissions market:*

$$\mathcal{T}(\tilde{Q}^*) = \tilde{Q}^*.$$

Although the distribution of households determines the equilibrium cutoffs, an individual household responds to this distribution only through the cutoff vector. Households therefore do not need to track the entire distribution of other households when making their investment choices. I prove existence of the equilibrium in Appendix A.

[33] See Judd (1985) and Sun (2006) for the exact law of large numbers in continuum economies.

5 Estimation Strategy

I estimate the model parameters by maximum simulated likelihood (MSL), exploiting the joint household-level distribution of tutoring participation, prices, and hours, study hours, test scores, and non-cognitive measurements across grades, together with early labor-market outcomes. In addition to the likelihood, I include moment conditions constructed from KLIPS that target lifecycle earnings-growth differentials by college tier and age. These moments identify the tier-specific lifetime income paths that constitute the tournament prizes.

5.1 The Likelihood Function

The individual specific likelihood contribution is defined by

$$\begin{aligned}&\mathcal{L}_i(\theta|q_{i1}, \{w_{it}\}_{t=1}^M, m_i, \{r_{it}\}_{t=1}^M) \\ &= \sum_{k=1}^K \pi_k \left[\left(\prod_{t=1}^M \mathcal{L}_{it}(\theta|Z_{it}, \lambda_k) \right) \mathcal{L}_i^{\text{tier}}(\theta|q_{i,M+1}, \lambda_k)\, \mathcal{L}_i^y\left(\theta|\text{Tier}_i, q_{i,M+1}, \{\tilde{\xi}_{it}\}_{t=1}^M, \lambda_k\right) \right] \end{aligned} \tag{19}$$

conditional on the initial test score q_{i1}, the household income stream $\{w_{it}\}_{t=1}^M$, parental education m_i, and the regional tutoring supply $\{r_{it}\}_{t=1}^M$. Here π_k is the probability of being type k; the period contributions $\mathcal{L}_{it}$ are the likelihood contributions of the household's choices, test scores, and non-cognitive measurements. $\mathcal{L}_i^{\text{tier}}$ is the likelihood of the observed college tier given the terminal test score and the number of seats for each college tier, which identifies the admission-noise dispersion σ_a, and $\mathcal{L}_i^y$ is the likelihood contribution of early labor-market earnings given the assigned tier, the terminal score, and the non-cognitive measurement history.

The period-specific choice contribution $\mathcal{L}_{it}(\theta|Z_{it}, \lambda_k)$ takes one of four forms depending on tutoring and self-study participation:

$$\begin{aligned}\mathcal{L}_{it}(\theta|Z_{it}, \lambda_k) = &\Big[\Pr(x_{it} > 0 \mid \tilde{\boldsymbol{\eta}}_{it})\, f(p_{it}, x_{it}, s_{it}, q_{i,t+1}, \xi_{i,t+1}) \Big]^{d_{it}^x d_{it}^s} \\ &\times \Big[\Pr(x_{it} > 0,\, s_{it} = 0 \mid \tilde{\boldsymbol{\eta}}_{it})\, f(p_{it}, x_{it}, q_{i,t+1}, \xi_{i,t+1} \mid s_{it} = 0) \Big]^{d_{it}^x (1-d_{it}^s)} \\ &\times \Big[\Pr(x_{it} = 0 \mid \tilde{\boldsymbol{\eta}}_{it})\, f(s_{it}, q_{i,t+1}, \xi_{i,t+1} \mid x_{it} = 0) \Big]^{(1-d_{it}^x) d_{it}^s} \\ &\times \Big[\Pr(x_{it} = 0,\, s_{it} = 0 \mid \tilde{\boldsymbol{\eta}}_{it})\, f(q_{i,t+1}, \xi_{i,t+1} \mid x_{it} = 0,\, s_{it} = 0) \Big]^{(1-d_{it}^x)(1-d_{it}^s)},\end{aligned}$$

where $d_{it}^x = 1$ indicates tutoring participation, $d_{it}^s = 1$ indicates non-zero self-study hours. The joint choice densities are evaluated by inverting the first-order conditions for the preference shocks, with the associated Jacobian term. The selection probabilities

are conditional on $\tilde{\boldsymbol{\eta}}_{it}$, the shocks recovered from the choices observed in each branch. The participation shock loads on the taste shocks, and the joint corner probabilities are the implied conditional normal probabilities.

As test scores are unobserved in grades 10 and 11, the latent score states in these years are integrated out by forward simulation. Missing choice variables are handled the same way: for household-years in which tutoring or study choices are not recorded, the corresponding choices are simulated using the structural model, so that every household in the estimation sample contributes its full history to the likelihood.

The structural parameters are estimated by maximizing the joint objective

$$\max_{\theta} \;\; \log \mathcal{L}(\theta) \;-\; (\bar{\mathbf{y}}^{\,\mathrm{sim}}(\theta) - \bar{\mathbf{y}}^{\,\mathrm{KLIPS}})' \mathbf{W} (\bar{\mathbf{y}}^{\,\mathrm{sim}}(\theta) - \bar{\mathbf{y}}^{\,\mathrm{KLIPS}}),$$

where $\log \mathcal{L}(\theta) = \sum_{i=1}^{N} \log \mathcal{L}_i(\theta | q_{i1}, \{w_{it}\}_{t=1}^{M}, m_i, \{r_{it}\}_{t=1}^{M})$. The second term incorporates the KLIPS lifecycle earnings data, where $\bar{\mathbf{y}}^{\,\mathrm{sim}}(\theta)$ and $\bar{\mathbf{y}}^{\,\mathrm{KLIPS}}$ are 21×1 vectors of simulated and empirical mean log incomes for the tier-by-age-bin cells defined in Appendix C.5. $\mathbf{W}$ is a diagonal weighting matrix with entries $n_c / \sigma^2_{j(c)}$, where n_c is the cell size and $\sigma^2_{j(c)}$ is the tier-specific earnings variance.

5.2 Identification

Parameters of the model are classified into the productivity parameters of the test-score and non-cognitive skill technologies, the earnings parameters, the preference parameters, and the parameters governing unobserved heterogeneity and shocks.

For production parameters, spatial variation in regional tutoring supply helps identify the substitution parameter ϕ. As supply shifts the effective cost of tutoring but not that of self-study, it induces variation in the investment allocations that is not from contemporaneous household tastes or skill shocks. Specifically, this identification relies on the assumption that, conditional on a household's initial states and permanent unobserved type, local tutoring supply is orthogonal to time-varying skill shocks and does not directly enter the test-score production function. Residential sorting based on unobserved academic demand could confound this restriction. To the extent that such sorting reflects persistent household traits, the permanent unobserved types capture it, and the orthogonality above is maintained as an assumption for any remaining dependence.[34] This supply-induced reallocation between tutoring and self-study affects the subsequent test score, which helps separate the substitutability of the two inputs from the overall scale of effort. This identification is similar in spirit to Caucutt, Lochner,

[34]School transfers, the channel through which residential re-sorting would operate during the schooling period, are rare: about 5 percent of students in the sample ever change schools across the six survey waves. Education-motivated moves in Korea concentrate before secondary-school entry, and such pre-period heterogeneity is absorbed by the initial states and permanent types.

Mullins and Park (2026), who use relative-demand conditions and variation in child-care prices relative to parental wages and goods prices to estimate substitution among parental-time, home-goods, and market-child-care inputs. They emphasize that identifying curvature parameters requires substantial relative-price variation. In my setting, local tutoring supply provides substantial spatial variation in effective tutoring costs: the number of hagwons within 5 kilometers ranges from about 40 in the lowest supply quintile to about 1,470 in the highest.

The remaining parameters of the test-score technology are identified by the covariation between the subsequent score $q_{i,t+1}$, the current score q_{it}, and the current inputs (p_{it}, x_{it}, s_{it}). To capture time-varying returns, I identify the productivity parameters separately for each grade. The substitution parameter ϕ and the relative contribution of tutoring quality and tutoring hours, κ, do not vary by grade.

The same regional supply variation also helps identify the tutoring-cost parameters. Participation differences across the five tutoring-supply quintiles discipline the participation costs, so that a household in a neighborhood with more tutors faces a lower hurdle to participate. Among participating households, variation in chosen tutoring quality across supply quintiles disciplines the search cost described in equation (9).

The non-cognitive skill parameters are identified by the measurement panel, following the dynamic latent factor framework of Agostinelli and Wiswall (2025). Each measure loads on the latent skill with an intercept, a loading, and independent measurement error. The location and scale of the latent skill are normalized once, in the initial period (grade 7), by fixing the loading and intercept of the social self-concept battery to one and zero. The location and scale of the skill in later waves are identified by an age-invariance restriction. Since this battery is administered with identical items and scales in every wave in which it appears (grades 7, 8, 9, and 12), its loading and intercept are equated to their grade-7 values; in grades 10 and 11, where only the peer-relations parcels are collected, the latent scale is carried by the estimated law of motion and the parcels' freely estimated loading. The remaining measurement parameters—the family and body loadings and intercepts, the shared loading of the grade-10–11 peer-relations parcels, and all error variances—are estimated, with contemporaneous cross-measure covariances within each wave separating measurement error from latent dispersion (Appendix Table E.2). Given this measurement system, the investment loadings $\gamma_{\xi,s,t}$ and $\gamma_{\xi,x,t}$ are identified by the covariation between a household's study and tutoring choices and subsequent non-cognitive-skill measurements. The loading of earnings on terminal non-cognitive skill is identified by the covariance between earnings residuals and the measurement-implied terminal skill.

The tier effects on labor-market outcomes are disciplined by earnings differences among students with similar CSAT scores who attend different college tiers, as described in Section 3. The tier-specific intercepts, the return to the CSAT score, and the

residual dispersion in the early labor-market equation are identified jointly by the earnings observations, CSAT scores, and assigned tiers. The admission-noise dispersion σ_a is identified by the tier-score mismatches in $\mathcal{L}_i^{\text{tier}}$, as described above. The lifetime component consists of the tier-specific intercepts, growth rates, and common curvature of the income paths that constitute the tournament prizes and is identified by the tier-by-age moments described in Appendix C.5.

The taste parameters α_c, α_x, and α_s and the prize weight α_v directly affect the value function. The flow utility weights are normalized to sum to one, with the leisure weight defined residually as $\alpha_\ell = 1 - \alpha_c - \alpha_x - \alpha_s$.[35] I do not let the taste parameters vary by period, so each parameter captures the overall value the household attaches to the corresponding choice across periods. The levels of the choice variables identify the taste weights, while the prize weight α_v is identified from how combined investments respond to the household's lifetime earnings.

The longitudinal structure of the data identifies time-invariant unobserved heterogeneity through persistent within-household covariation in choices and test-score residuals. The time-varying shocks η_{it}^z $(z = c, x, s)$ capture what remains unexplained by observables and permanent types. Their covariance matrix, including correlations between the participation shock and the taste shocks, is identified by the joint distribution of the inverted continuous-choice residuals and tutoring participation outcomes.

Finally, conditional on permanent type, I assume that the time-varying preference, participation, and skill shocks are independent of the initial test score q_{i1}, parental education m_i, the parental income path $\{w_{it}\}_{t=1}^M$, and regional tutoring supply $\{r_{it}\}_{t=1}^M$. I also assume that these shocks are serially independent over time, while allowing for the contemporaneous correlations described above.

6 Estimation Results

6.1 Estimates

Estimates suggest that both skills are highly persistent. Table 2 presents the estimates of the two production functions. For test scores, the persistence estimates range from 0.440 to 0.923: previous achievement is the dominant input into current achievement. This is the "skills beget skills" mechanism (Cunha and Heckman, 2007; Cunha, Heckman and Schennach, 2010). Non-cognitive self-productivity ranges from 0.49 to 1.30 and exceeds one in the first and last transitions. Neither persistence profile is monotone in age. Self-productivity above one is admissible over the finite six-year

[35] The estimation also permits negative taste weights, for example tutoring entering as a disutility. In that case, the utility scale is instead fixed by setting the leisure weight to one, $\alpha_\ell = 1$. At the estimated values, all taste weights are positive, so the sum-to-one normalization applies.

Table 2: Structural Estimates

Panel A. Test-score production						
	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$
Previous Test Score (δ_{qt})	0.440	0.708	0.683	0.923	0.585	0.554
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
Effort Productivity (ν_t)	0.035	0.020	0.023	0.034	0.022	0.013
	(0.001)	(<0.001)	(0.001)	(<0.001)	(0.001)	(<0.001)
Share of Tutoring Effort (δ_{et})	0.579	0.630	0.472	0.424	0.363	0.322
	(0.007)	(0.007)	(0.005)	(0.005)	(0.005)	(0.004)
Log TFP ($\ln A_t$)	3.334	1.730	1.878	0.365	2.453	2.636
	(0.016)	(0.017)	(0.016)	(0.019)	(0.014)	(0.015)
Substitution Parameter (ϕ)	0.141	Intensity of Quality (κ)			0.495	
	(0.003)				(0.004)	
Panel B. Non-cognitive skill production						
	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$
Hours of Self-Study ($\gamma_{\xi,s,t}$)	-0.038	-0.028	-0.039	-0.026	-0.026	-0.037
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Hours of Tutoring ($\gamma_{\xi,x,t}$)	-0.031	-0.027	-0.029	-0.018	-0.015	-0.017
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Self-Productivity ($\gamma_{\xi,1,t}$)	1.300	0.972	0.887	0.494	0.831	1.117
	(0.016)	(0.011)	(0.011)	(0.006)	(0.010)	(0.052)
Intercept ($\gamma_{\xi,0}$)	0.092	Shock SD (σ_{ε_ξ})			0.235	
	(0.003)				(0.003)	
Initial SD ($\sigma_{\xi 0}$)	0.291	Measurement SD (σ_m)			0.478	
	(0.003)				(0.004)	
Panel C. Early labor-market earnings and admission						
Constant (a_{0y})	2.813	Return to CSAT Score (b_{z0})			0.415	
	(0.018)				(0.002)	
Return to Non-cognitive Skill (b_h)	0.350	Earnings Shock SD (σ_y)			0.467	
	(0.004)				(0.014)	
Tier Premiums (τ_1, τ_2, τ_3)	0.241	0.069	0.034	Admission Noise SD (σ_a)		0.119
	(0.007)	(0.012)	(0.011)			(0.002)

Note: Standard errors in parentheses: household-block bootstrap with re-estimation in each replicate. Periods $t = 1, \ldots, 6$ correspond to grades 7–12. Panel A: the test-score technology combines the tutoring-effort composite $\tilde{e}_t = p_t^{\kappa} x_t^{1-\kappa}$ and self-study through a CES aggregate with substitution parameter ϕ; the share of self-study is implied ($\delta_{st} = 1 - \delta_{et}$); ϕ and κ are common across grades. Panel B: law of motion for non-cognitive skill in equation (6), measured by the six-wave non-cognitive survey battery with post-investment timing; σ_m is the measurement-error dispersion of the anchor (social self-concept) measure, and the grade-10–11 peer-relations parcels carry a freely estimated loading, intercept, and dispersions. All non-cognitive parameters are estimated; the self-productivity $\gamma_{\xi,1,t}$ varies by grade. Panel C: earnings equation (3), with tier premiums expressed relative to Tier 4 (the baseline tier); the admission-noise dispersion σ_a is identified by observed tier-score mismatches; the age-profile parameters of equation (3), identified by the KLIPS moments of Appendix C.5, and the implied tier gaps in log lifetime income are reported in Table E.3. Preference, parental-education, and participation-cost parameters are reported in Table 3.

Table 3: Structural Estimates (continued)

Panel D. Preferences					
Prize Weight (α_v)	14.351		Taste for Consumption (α_c)		0.0726
	(0.368)				(0.0014)
Taste for Tutoring (α_x)	0.0951		Taste for Self-Study (α_s)		0.0630
	(0.0007)				(0.0010)
Leisure Weight (α_ℓ, implied)	0.769		Gr. 12 Study Shift (μ_s)		0.679
	(0.002)				(0.008)
Panel E. Parental education taste shifters					
	Study taste		Tutoring taste		
	Middle sch.	High sch.	Middle sch.	High sch.	
Father only	0.134	0.137	0.061	0.045	
	(0.018)	(0.016)	(0.012)	(0.013)	
Mother only	0.047	0.008	0.025	0.013	
	(0.034)	(0.022)	(0.022)	(0.020)	
Both college	0.259	0.229	0.087	0.085	
	(0.021)	(0.014)	(0.010)	(0.009)	
Panel F. Tutoring participation and search costs (by supply quintile)					
	$r=1$	$r=2$	$r=3$	$r=4$	$r=5$
Participation Cost (χ_r)	0.120	0.103	0.103	0.101	0.090
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Search Cost (γ_{2r})	0.0009	0.0007	0.0007	0.0006	0.0005
	(0.0001)	(<0.0001)	(<0.0001)	(<0.0001)	(<0.0001)
Hurdle Shock SD (σ_d)	0.0554		Loading on tutoring taste (ρ)		0.034
	(0.0016)				(0.020)
			Loading on study taste (b_{ns})		-0.288
					(0.013)

Note: Continuation of Table 2. Standard errors in parentheses: household-block bootstrap. Panel D: the flow-utility weights are normalized to sum to one, so the leisure weight is implied (its standard error is computed from the same bootstrap replicates); the prize weight α is outside the normalization; the grade-12 study-shock mean μ_s captures cramming before the final examination (the grade-12 tutoring-shock mean is fixed at zero). Panel E: father×mother college categories; taste shifts relative to the baseline of neither parent college-educated; middle-school and high-school shifts are estimated separately. Panel F: a household participates in tutoring when its participation value net of the quintile-specific cost χ_r exceeds the hurdle shock $\sigma_d \zeta^d$; ρ and b_{ns} are the loadings of the hurdle shock on the tutoring- and study-taste shocks; the search cost is quadratic in tutoring quality with quintile-specific curvature γ_{2r} (the linear term is fixed at zero).

horizon. Together with the estimated innovation and initial variances, these coefficients imply bounded dispersion over the sample period.

The productivity of contemporaneous effort, by contrast, declines steeply as students approach the final examination. The estimated effort parameters fall over time, with the final period effect plummeting to 0.013. Figure 7 presents the computed average marginal effects for hours of tutoring and self-study. In the final period, the average marginal effect on the log test score is about 0.002 for self-study and 0.0004 for hours of tutoring. These magnitudes are small relative to the dispersion of achievement. With a standard deviation of log test scores of about 0.17, a one-standard-deviation increase in weekly self-study hours raises the final test score by about 0.11 standard deviations, and the same increase in tutoring hours by only 0.02. In grade 7, the corresponding effects are 0.27 and 0.08 standard deviations, so the productivity of tutoring falls by about three quarters between middle school and the final year, and that of self-study by about 60 percent. These time-varying effects are consistent with the child development literature, which finds that the productivity of investments in children declines with age (Cunha, Heckman and Schennach, 2010; Del Boca, Flinn and Wiswall, 2014), and the estimates here extend that pattern into late adolescence. The small marginal effects of tutoring are also in line with the modest average achievement effects reported in the Korean private tutoring literature (Kang and Park, 2021). Despite such minimal impact in the final period, households maintain substantial investment in both tutoring and self-study.[36] The tournament rationalizes this pattern: when admission to an elite tier turns on rank, even small score advantages are worth defending, so investment persists despite low marginal returns. Together, the two estimates concentrate the value of investment early. Declining productivity makes early effort more effective, while strong persistence carries its gains into later grades. These findings speak to the importance of investment timing emphasized by Caucutt and Lochner (2020).

Test-score-directed effort erodes non-cognitive skill. Panel B of Table 2 shows that hours of both inputs reduce non-cognitive skill contemporaneously: all twelve estimated coefficients are negative, with the self-study coefficients stable across grades and the tutoring coefficients somewhat smaller in absolute magnitude in high school. With broadly stable coefficients on log hours, the final year's high effort levels generate a larger non-cognitive cost. The effect is economically meaningful. Because hours enter the law of motion (6) in logs, a one-standard-deviation increase in final-year study hours from the mean raises $\log(1+s)$ by approximately 0.5. This lowers non-cognitive skill by 0.02 on impact, about 6 percent of a standard deviation of the initial skill distribution ($\sigma_{\xi 0} = 0.291$). Given the estimated persistence, these effects carry forward into subsequent periods. Because non-cognitive skill carries an independent earnings

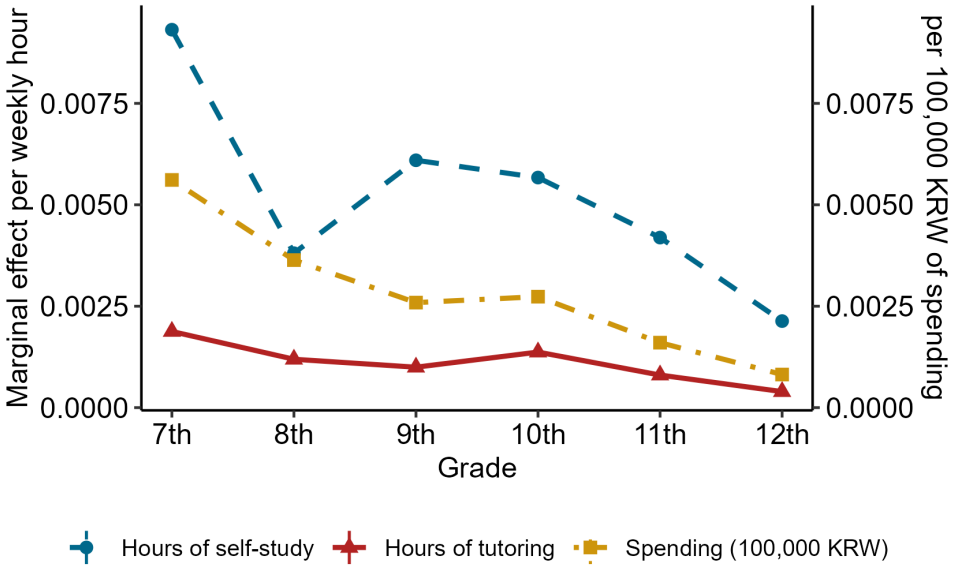
[36] As I show in Section 6.2, the high investments in the data are captured by the estimated model.

return, this reduction constitutes a cost in the household's investment problem.[37]

Households nevertheless maintain investment because the cognitive channel dominates in earnings: the test score earns both the continuous score gradient and the tournament tier premia, while the non-cognitive channel has no tournament component.

Finally, tutoring and self-study are gross substitutes in producing test scores. The estimated substitution parameter of 0.141 implies an elasticity of substitution of 1.16. Previous work emphasizes the importance of substitution patterns between parental time and purchased inputs in children's skill production (Caucutt, Lochner, Mullins and Park, 2026). Here, substitution occurs between purchased tutoring and students' own study time. This substitutability has important implications for understanding intergenerational mobility. Without accounting for self-study as a substitute for tutoring, models would likely overstate the role of parental investment in intergenerational persistence of earnings. Self-study provides income-constrained households another channel to compete in the college admissions tournament, potentially mitigating the advantages of families who can afford extensive tutoring.

Figure 7: Average Marginal Effects of Hours Allocation



Note: This figure presents the average of marginal effects of hours of self-study and hours of tutoring over time. Due to the functional form of the test score function, the marginal effects differ by each individual. The marginal effects are computed using the first order derivative with respect to hours of self-study (s_{it}) or hours of tutoring (x_{it}) and the estimated parameters.

Panel C reports the earnings equation that generates the tournament stakes. Conditional on the final test score and non-cognitive skill, admission to a Tier 1 college raises early-career earnings by 27 percent relative to Tier 4, with premia of 7 and 3 percent for Tiers 2 and 3. The continuous returns operate alongside the tier ladder: a 10 percent

[37]The signal share of the anchor measure is roughly one-half of its variance in each wave in which it is collected, and the peer-relations parcels reach 64–70 percent (Appendix Table E.2).

higher CSAT score raises earnings by about 4 percent within a tier, and non-cognitive skill carries an independent return.

Since tier premia compound over the life cycle through tier-specific earnings growth, the entry premia alone do not explain the whole stake. The implied gap in the present value of lifetime earnings between a Tier 1 graduate and a Tier 4 graduate is 0.70 log points, roughly a doubling. These lifetime values constitute the tournament prizes, and their magnitude is what makes small score advantages near the admission cutoffs worth defending.

Panel D of Table 3 presents the estimates of the preference parameters. The flow-utility weights on consumption, tutoring hours, and self-study hours are normalized to sum to one together with the weight on leisure, so the leisure weight is an implied estimate. The estimated weights are small relative to leisure (0.073 for consumption, 0.095 for tutoring, and 0.063 for self-study), and all are positive: tutoring carries direct consumption value beyond its effect on scores, which matters for the welfare evaluation of policies in Section 8. The prize weight α_v is estimated at 14.35 under the model's flow-utility normalization.[38] Its interpretation depends on the terminal payoff specification and discounting, so its magnitude is not directly comparable to the flow-utility weights. The grade-12 study shifter μ_s of 0.68 captures the additional taste for study in the examination year.

Parental education shifts preferences mainly through the taste for self-study. Relative to households where neither parent attended college, both-college households have a study-taste weight about 30 percent higher in middle school and 26 percent higher in high school, while differences in tutoring-taste weights across parental-education groups are substantially smaller. Households where only the father attended college fall in between, with a study-taste weight about 14 percent higher than the baseline in both school stages. For households where only the mother attended college, the differences from the baseline are smaller and the standard errors are large, which reflects the small number of such households in the sample.

Permanent unobserved heterogeneity enters through household types on three margins: the taste for self-study, the marginal utility of consumption, and the taste for tutoring. The mixture assigns probabilities of 0.22, 0.34, and 0.44 to the three types (Appendix Table E.1), so no single type dominates. Type differences operate mainly through consumption utility and the taste for tutoring. The middle type also has a higher taste for self-study (0.107 against the baseline 0.063).

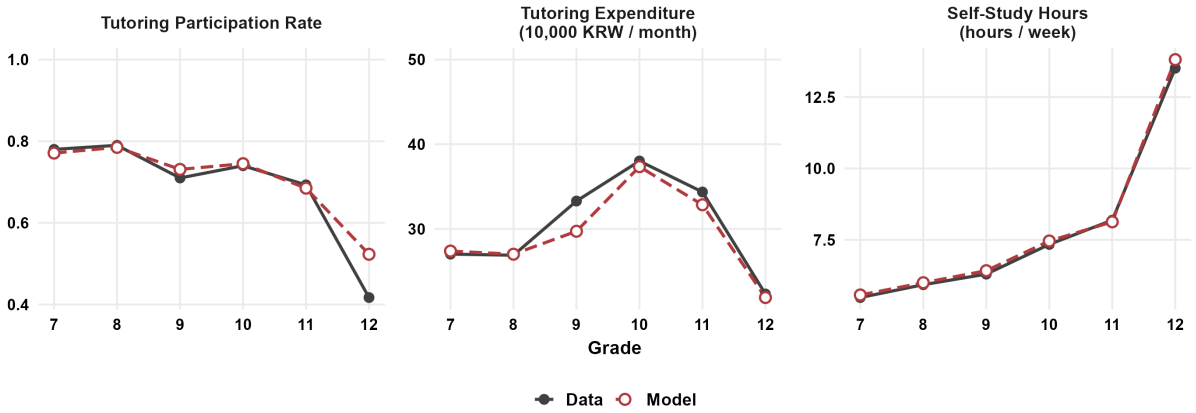
Students with a stronger taste for self-study find studying alone more attractive, reducing their incentive to participate in tutoring. The estimated relationship between the study-taste shock and the participation shock works in the opposite direction: the

[38] Unlike Del Boca, Flinn and Wiswall (2014), who specify terminal utility over the child's skill level, I specify terminal utility over lifetime earnings generated by skills and college tier.

negative coefficient ($b_{ns} = -0.29$) lowers the participation hurdle. The first effect dominates, so these students participate less overall. The corresponding coefficient on the tutoring-taste shock (ρ) is small and imprecisely estimated.

6.2 Model Fit

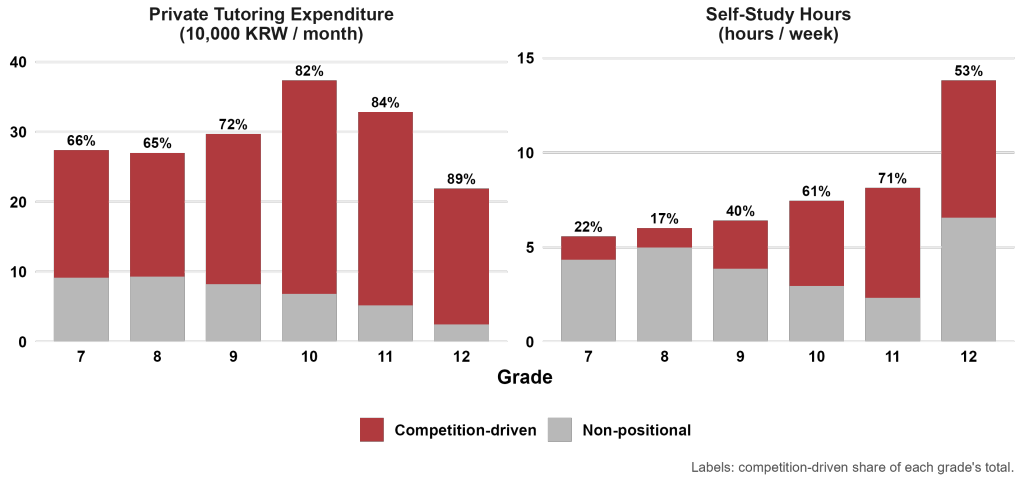
Figure 8: Model Fit: Investment by Grade



Note: Grade-specific means in the data and in the model simulation at the estimated parameters. Tutoring expenditure is per household including non-participants, in 10,000 KRW per month ($\approx$ USD 9 at the 2005–2010 average exchange rate). Self-study is weekly hours.

The model fits the key empirical moments and the overall distributions of the data well, with three visible misses. Table F.1 shows the model reproduces the means and standard deviations of tutoring participation, tutoring prices and hours, self-study time, and test scores; the notable gaps are grade-12 participation (0.523 in the model against 0.417 in the data), grade-9 tutoring hours (11.2 against 12.9), and the grade-12 score dispersion (0.146 against 0.170). The grade-12 participation moment also differs from the descriptive rate of about 51 percent reported in Section 3: the fit table conditions on the estimation sample's observations used in estimating household choices, while the descriptive figure uses all grade-12 responses. Figure 8 compares model and data trajectories of tutoring participation, tutoring expenditure, and self-study hours over time. Figure F.1 evaluates distributional fit beyond these moments with quantile–quantile plots of simulated against empirical distributions for self-study, tutoring hours and prices, test scores, the non-cognitive measures, and log earnings, which track the 45-degree line closely outside the upper tails of tutoring hours, late-grade study hours, and final scores.

Figure 9: Competition-Driven versus Non-Positional Investment by Grade



Note: Each bar is the baseline mean private tutoring expenditure per household including non-participants, in 10,000 KRW per month (≈ USD 9 at the 2005–2010 average exchange rate; left panel), or mean weekly hours of self-study (right panel), simulated at the estimated parameters. The non-positional component is the investment that would remain even if the college admissions tournament were removed. Specifically, I simulate the case in which earnings depend only on the absolute CSAT score and non-cognitive skills, not on the college tier. The competition-driven component is the investment that disappears once the tournament is removed. Percentages report its share of each grade's total.

7 Decomposition of Household Investments

Using the estimated tournament model, I quantify the household investments that are driven by the college admission competition. The competition-driven component is the investment that disappears when the tournament is removed. Specifically, I equalize the tier component of lifetime earnings across colleges at its seat-weighted baseline average. This removes the earnings incentive to improve admission rank while preserving total prize mass and holding direct returns to skills fixed. The continuous returns to the CSAT score and to non-cognitive skills in equation (3) are unchanged.[39]

A substantial share of private tutoring expenditure is driven by competition. On average, across the six grades, 76.6 percent of expenditure disappears when the tournament is removed. Figure 9 reports the decomposition by grade. Baseline expenditure is roughly flat through middle school, peaks at 37 (10,000 KRW per month) in grade 10, and falls to 22 in the CSAT year as participation drops, matching the pattern in the data. The competition-driven share rises from 66 percent in grade 7 to 89 percent in grade 12, while the non-positional component falls almost monotonically, from 9.2 to 2.5. Final-year tutoring is thus almost entirely driven by the tournament. The high final-year share is consistent with the estimated marginal effects in Figure 7, which show that ad-

[39]The choice of the common level is innocuous: with log utility, a common shift in all tier payoffs does not affect investment choices, and the decomposition is numerically identical when tiers are instead equalized at the lowest-tier level.

Table 4: Aggregate Decomposition of Investment by Incentive Channel

	Spending (₩10,000/mo)	Participation	Spending | Participation	Self-study (h/wk)
No score-based incentives	1.2	0.28	4.3	3.0
+ Cognitive return only	6.9	0.38	18.1	4.2
+ Tournament only	16.8	0.54	31.1	5.8
+ Both (full model)	29.4	0.71	41.6	7.9

Note: Six-grade averages, simulated at the estimated parameters with admission cutoffs re-clearing so that baseline tier enrollment equals the observed seat counts. Spending is the unconditional mean monthly tutoring expenditure (non-participants included); spending given participation is the ratio of means. Spending is in units of 10,000 KRW per month ($\approx$ USD 9 at the 2005–2010 average exchange rate). "No score-based incentives" flattens all tier premia and sets the score's continuous earnings return (b_{z0}) to zero; the second row restores b_{z0} while tier premia stay flat; the third instead restores the tier prizes; the final row is the full model. The non-cognitive earnings return (b_h) and all preference parameters are held at their estimates in every scenario.

ditional investment moves test scores little by grade 12. With limited score effects, the direct earnings return alone would justify only minimal spending, and the expenditure that remains is sustained by the tournament.

The competition-driven share of self-study is uniformly lower. On average, 47.2 percent of study hours disappear when the tournament is removed. The right panel of Figure 9 presents the decomposition for weekly hours of self-study. The share rises from about 20 percent in middle school (22 percent in grade 7, dipping to 17 percent in grade 8) to a peak of 71 percent in grade 11, then falls back to 53 percent in grade 12. Most of the CSAT-year jump in study hours is non-positional. Students study more in the final year even net of the tournament motive. Across all grades, the competition-driven share of self-study is well below the corresponding share for tutoring expenditure.

The non-positional component itself combines two distinct motives. To understand these mechanisms, I simulate the model under different combinations of these three motives: competition, the direct human-capital return, and preferences. The first two operate through the earnings equation and the tournament structure. The competition motive reflects the discrete tier prizes allocated by rank in the admission tournament, and the direct motive reflects the continuous earnings return to cognitive skill (b_{z0}), which remains active even when rank is irrelevant. The preference motive operates through the flow utility, the value households attach to the activity itself. To quantify each motive, I switch the channels on one at a time and compare simulated tutoring and self-study with the baseline.

Table 4 shows that, for tutoring, no single channel alone adequately explains the observed demand. The cognitive return alone supports spending of 6.9 and the tour-

nament alone 16.8, against 29.4 in the full model, so the interaction between the two accounts for roughly a quarter of baseline spending.

Unlike private tutoring expenditure, preference plays an important role in shaping the demand for self-study. With all score-based incentives removed, students still study 3.0 hours per week, nearly 40 percent of the baseline 7.9 hours, whereas tutoring spending collapses to 1.2, barely 4 percent of its baseline of 29.4. Restoring the cognitive return alone raises study to 4.2 hours, and the full set of score-based incentives raises it to 7.9: incentives scale study time up by a factor of 2.6. This pattern is consistent with the utility estimates. The flow-utility weight on self-study is positive ($\alpha_s = 0.063$). Thus, study time enters the household problem partly as a consumption good rather than purely as an investment input. A potential explanation is parent-child interaction. Study time at home is partly the outcome of an intra-household mechanism, such as parental supervision or rewards for studying, which this model captures in reduced form through the study taste.[40]

The decomposition also speaks to the human-capital-versus-signaling literature, with a caveat on interpretation. The competition-driven share measures the source of investment incentives, namely prizes allocated by rank, not the social content of the prize itself. The tier premium that drives the tournament could include several channels: employer screening on college identity, alumni networks, peer effects, and any value added by elite colleges. These are indirect returns that accrue only through admission, so the investment they induce is competition-driven in the decomposition even when the underlying channel raises productivity. The share of the premium that is pure screening is not separately identified here. Evidence from Korean hiring data, where college tier strongly predicts job offers while college performance adds little (Cho, Kang, Kang and Kim, 2024), suggests the screening component is substantial.[41] The welfare analysis below does not require taking a stand: prizes are estimated from observed earnings and held to the same rules in every counterfactual. But the interpretation of competition-driven investment as purely wasteful depends on the share of the screening component.

The decomposition implies that policy leverage lies in the competition channel. The next section examines reforms that change the supply side of the tournament directly.

[40] See Del Boca, Flinn, Verriest and Wiswall (2026) for a model in which parents incentivize child study time through monetary transfers.

[41] Appendix Table C.2 reports the regression: the probability of a final interview request at a major conglomerate by college tier and college GPA.

8 Policy Analysis

The previous section shows that competition accounts for a substantial share of household investment. On the one hand, households might be better off under an alternative regime or policies that suppress the current demand for tutoring. On the other hand, tutoring also increases household welfare, through higher lifetime earnings from the human capital it builds and through its consumption value. This creates a policy tension. Direct regulation may reduce inefficient positional investment, but it may also sacrifice the benefits of investment, including higher lifetime earnings, cognitive skills, and the consumption value households derive from tutoring through the preference channel. I compare policies that restrict tutoring with policies that alter the competitive returns generated by college admissions.

I first consider direct regulations on private tutoring. These include (i) a complete tutoring ban and (ii) a curfew that limits tutoring hours. The ban is motivated by policies implemented in Korea and China, while the hours cap is motivated by tutoring curfews previously imposed in Korea. Then, I consider reforms to the resource side of college admissions, following recent proposals in the literature (Mahler, Tertilt and Yum, 2025; Fang and Liu, 2026). These include (iii) a threshold-based admission lottery and (iv) an expansion of college seats. For each counterfactual, I re-solve the household decision problem and the endogenous admission cutoffs that clear college capacity. I evaluate each policy by its effects on tutoring expenditure, self-study hours, test scores, non-cognitive skills, intergenerational persistence of earnings, upward mobility, and household welfare.

(a) **Tutoring Bans** ($p = x = 0$): Under this policy, private tutoring is prohibited while self-study remains unrestricted. Korea introduced a nationwide private-tutoring ban in 1980, and China's 2021 Double Reduction policy restricted academic-subject tutoring during compulsory education through grade 9. Motivated by these policies and by the rising competition-driven share of investment across grades, I simulate three versions: a complete ban in grades 7–12, a middle-school ban in grades 7–9 reflecting the scope of the Chinese policy, and a high-school ban in grades 10–12 targeting the years in which the competition-driven share is largest. Households anticipate the ban when choosing investment in the unrestricted grades, so a high-school ban can shift tutoring forward into middle school rather than eliminate it. Because tutoring is unavailable at any quality level during the banned grades, households respond within those grades only by reallocating toward self-study.

(b) **Tutoring-Hours Cap**: Korean local governments have imposed curfews on the operating hours of private tutoring institutions. Since the model measures weekly

tutoring hours rather than the time of day at which tutoring occurs, I model a curfew as a stylized cap of six weekly tutoring hours.[42] Unlike the ban, the cap leaves the quality margin open: households may respond by purchasing more expensive tutoring per hour, by studying more, or both.

(c) **Admission Lottery**: Fang and Liu (2026) propose a threshold-based admission lottery to mitigate excessive competition: students scoring above a specified threshold enter a lottery and are randomly allocated across the colleges above it. In the tournament model this is equivalent to merging the tiers above the threshold into one, because college tiers differ only in the lifetime income they offer: for household investment decisions, the lottery is equivalent to a merged tier offering the seat-weighted average log prize.[43] I consider two versions: a lottery above the Tier 2 cutoff, covering the elite tiers, and a lottery above the Tier 3 cutoff. The lottery removes the marginal return to rank within the eligible pool while preserving the incentive to reach the threshold.

(d) **Tier-1 Seat Expansion**: I expand the number of seats at Tier 1 colleges by 50 percent, holding Tier 2 and Tier 3 seat counts fixed.[44] Admission cutoffs adjust endogenously, and households reoptimize their investments. Two opposing forces shape the tutoring response. Additional seats make Tier 1 admission attainable for more students, potentially strengthening investment incentives, while a decline in the Tier 1 earnings premium reduces the stakes of the tournament. I assume that the expansion lowers the Tier 1 log lifetime-income premium by $\ln(1.5)/\sigma$. The benchmark $\sigma = 2.1$, motivated by the estimates in Kim (2005), implies a reduction of 0.19 log points.[45] Appendix Table G.1 reports the tutoring response over the full grid of expansion sizes (2.5–100 percent) and elasticities.

Table 5 reports the effects of each policy on household investment, skills, intergenerational mobility, and welfare. Welfare is reported as the equivalent proportional change in lifetime earnings, the unit used in the discussion below.

Policies that directly restrict private tutoring reduce it at the expense of household welfare. A complete tutoring ban improves intergenerational mobility relative to the baseline. The intergenerational elasticity of earnings (IGE) falls from 0.109 to 0.078, and elite admission from the bottom parental-income quintile nearly doubles,

[42] This follows the simulation strategy of Ho and Wang (2024).

[43] Because utility is logarithmic in the prize, the expected utility of random assignment equals the utility of the seat-weighted average of the merged tiers' log premia, so the lottery and the merged tier induce identical behavior.

[44] Relatedly, Kang, Koh and Park (2025) study a mid-1990s relaxation of university admission quotas in South Korea and find that deregulation reduced private tutoring participation among high school students by 9–10 percentage points, approximately one-third of the baseline rate.

[45] Kim (2005) studies the earnings response to changes in the relative supply of college- and high-school-educated workers in South Korea.

Table 5: Policy Comparisons: Investment, Skills, Mobility, and Welfare

	Tutoring expenditure	Self-study hours	Mean log final test score ($\bar{q}_{12}$)	Non-cognitive skill ($\bar{\xi}_T$)	Welfare (% lifetime inc.)	IGE	$P(\text{elite} \mid Q_1)$
Baseline	29.4	7.90	5.985	−0.005	—	0.109	0.027
Tutoring ban (gr. 7–12)	0.0	9.47	5.964	0.106	−5.5	0.078	0.051
Tutoring ban (gr. 7–9)	15.3	8.56	5.974	0.051	−3.0	0.101	0.031
Tutoring ban (gr. 10–12)	14.1	8.81	5.974	0.049	−2.6	0.094	0.040
Hours curfew (6h)	18.3	8.51	5.982	0.022	−2.3	0.099	0.027
Lottery above Tier 2	27.3	7.55	5.981	0.008	−0.5	0.110	0.028
Lottery above Tier 3	26.2	7.37	5.979	0.016	−0.7	0.108	0.042
Tier-1 expansion (+50%, $\sigma = 2.1$)	26.2	7.37	5.979	0.016	−0.3	0.107	0.037

Note: Simulated at the estimated parameters. Admission allocations are recomputed under each policy's seat capacities and assignment rules. Tutoring expenditure and self-study hours are unconditional means over grades 7–12 (non-participants enter as zero); expenditure is monthly, in units of 10,000 KRW ($\approx$ USD 9 at the 2005–2010 average exchange rate). Welfare is the equivalent proportional change in lifetime earnings: $\beta^{M-1}\alpha_v \ln(1+\lambda) = \Delta W$, where ΔW is the change in ex-ante expected lifetime utility relative to the baseline. IGE is the OLS slope of the child's log lifetime income on log parental income. $P(\text{elite} \mid Q_1)$ is the probability of admission to Tier 1–2 for a child in the bottom parental-income quintile. In the lottery scenarios, tiers within the eligible pool are assigned randomly. For Tier-1 seat expansion, σ governs the earnings-premium response, with baseline calibration $\sigma = 2.1$. Appendix Table G.1 reports tutoring responses across expansion sizes and premium elasticities.

but the welfare cost is 5.5 percent of lifetime earnings and the mean final score declines. Households partially substitute toward self-study, which increases from 7.9 to 9.5 hours per week. Nevertheless, total effort hours fall, because the increase in self-study replaces only part of the banned tutoring hours. Non-cognitive skill rises as total effort decreases, reflecting the estimated technology in which intensive academic effort erodes non-cognitive development. The grade-ranged bans each cost roughly half the complete ban (3.0 percent of lifetime earnings for the middle-school ban, 2.6 percent for the high-school ban), with households shifting part of the banned spending into the unrestricted grades. The high-school ban delivers more of the mobility gain. The hours curfew operates on three margins at once. Tutoring hours are forced down for the households the cap binds, and some of them exit tutoring altogether: participation falls from 0.71 to 0.65. Households that remain shift toward quality, purchasing more expensive tutoring per hour, and toward self-study, which rises to 8.5 hours per week. The hours restriction dominates the quality response, so total expenditure falls by roughly 38 percent, and all effects are intermediate between the baseline and the full ban.

A lottery that randomizes assignment among the top tiers reduces both tutoring expenditure and self-study. Randomizing assignment within Tiers 1–2 lowers average tutoring expenditure by 7 percent, while extending the lottery through Tier 3 lowers it by 11 percent. Investment persists because households still compete to enter the eligible pool. Crossing its lower boundary now provides access to the lottery over higher-tier prizes, preserving a strong incentive to invest. Welfare losses range from 0.5 to 0.7

percent of lifetime earnings, substantially below those under tutoring bans.

Expanding Tier 1 seats reduces tutoring expenditure across the range of premium declines. At the calibrated $\sigma = 2.1$, expanding Tier 1 seats by 50 percent reduces tutoring expenditure by 11 percent and self-study by 7 percent. The tutoring reduction is comparable to that under the Tier 3 lottery.[46] The expansion attains this reduction with a smaller welfare loss of 0.3 percent of lifetime earnings.

The sensitivity analysis clarifies the role of premium adjustment. Holding the Tier 1 premium fixed ($\sigma = \infty$) provides a useful benchmark that isolates the effect of additional admission opportunities, but is less plausible as a long-run scenario, because the premium reflects the scarcity of top-tier graduates. In this benchmark, tutoring rises by 5 percent. Allowing the premium to decline reverses this response. Tutoring falls in every reported finite-elasticity scenario, with reductions ranging from 0.5 percent at $\sigma = 6$ to 11 percent at $\sigma = 2.1$. Thus, the effect of seat expansion depends on the balance between greater admission opportunities and smaller earnings differences across tiers.

Overall, the simulations distinguish direct regulation of educational investment from reforms to admissions. Tutoring restrictions induce substitution toward self-study and trade welfare for mobility. Bans generate the largest reduction in intergenerational income persistence, but primarily through leveling down, at welfare costs of up to 5.5 percent of lifetime earnings. Admission lotteries and the calibrated Tier 1 expansion reduce competitive investment at substantially smaller welfare costs. The Tier 3 lottery delivers a larger improvement in elite access for children from low-income households, while the expansion entails a smaller welfare loss.[47]

Tier 1 expansion arguably offers a more institutionally feasible route than an admission lottery because it relaxes longstanding restrictions on the number of seats colleges can offer while preserving rank-based admissions.[48] To the extent that these restrictions bind, relaxing them would allow colleges to accommodate additional demand.

9 Conclusion

I develop and estimate a dynamic tournament model of college admissions in which households compete using private tutoring and self-study. The model rationalizes high levels of investment despite their modest impact on test scores: what sustains invest-

[46] The two policies yield the same aggregate investment and skill outcomes at the reported precision, but differ in their allocation of college seats and their welfare and mobility effects.

[47] At the benchmark calibration, Tier 1 expansion delivers roughly twice the tutoring expenditure reduction per unit of household welfare loss as the complete tutoring ban, and more than twice that of the Tier 3 lottery. These comparisons concern the simulated policy changes and exclude the costs of supplying additional college seats.

[48] See Kang, Koh and Park (2025) for a discussion of these regulatory limits in South Korea.

ment is the need to defend rank near tier boundaries, not the direct human-capital return. Quantitatively, about three quarters of tutoring expenditure is competition-driven and this share rises to 89 percent in the final year of high school, precisely when the marginal effects are smallest. Self-study is different: more than half of it, including most of the final-year surge, survives the removal of the tournament and is non-positional. The policy experiments show that the instrument matters. Tutoring restrictions induce substitution toward self-study and improve mobility at substantial welfare cost. Admission lotteries and the calibrated Tier 1 expansion reduce competitive investment at considerably smaller welfare costs. The expansion experiment also highlights two opposing forces: additional admission opportunities can stimulate investment, while the accompanying decline in the Tier 1 premium weakens competitive incentives.

These findings suggest two directions for future research. First, incorporating wealth transfers into the admissions tournament would clarify how parents allocate resources between securing college access and transferring wealth directly, and whether policies that reduce educational competition shift intergenerational advantage toward wealth transmission. Second, allowing colleges to adjust capacity, tuition, and educational quality in response to the tutoring regulations and admissions reforms considered in this paper would clarify how higher education supply responses reshape parental investment and admission competition.

References

Agostinelli, F. (2018) Investing in Children's Skills: An Equilibrium Analysis of Social Interactions and Parental Investments, *Unpublished Manuscript*.

Agostinelli, F., Doepke, M., Sorrenti, G. and Zilibotti, F. (2026) It Takes a Village: The Economics of Parenting with Neighborhood and Peer Effects, *Journal of Political Economy*, **134**, 313–365.

Agostinelli, F. and Wiswall, M. (2025) Estimating the technology of children's skill formation, *Journal of Political Economy*, **133**, 846–887.

Anelli, M. (2020) The Returns to Elite University Education: a Quasi-Experimental Analysis, *Journal of the European Economic Association*, **18**, 2824–2868.

Arcidiacono, P., Aucejo, E., Coate, P. and Hotz, V. J. (2014) Affirmative action and university fit: Evidence from Proposition 209, *IZA Journal of Labor Economics*, **3**, 1–29.

Arcidiacono, P., Aucejo, E. M. and Hotz, V. J. (2016) University differences in the graduation of minorities in STEM fields: Evidence from California, *American Economic Review*, **106**, 525–62.

Bastedo, M. (2021) Holistic admissions as a global phenomenon, in *Higher Education in the Next Decade*, Brill, pp. 91–114.

Becker, G. S. and Tomes, N. (1979) An equilibrium theory of the distribution of income and intergenerational mobility, *Journal of Political Economy*, **87**, 1153–1189.

Bodoh-Creed, A. and Hickman, B. R. (2019) Pre-College Human Capital Investments and Affirmative Action: A Structural Policy Analysis of US College Admissions, *Working Paper*.

Bolt, U., French, E., Hentall-MacCuish, J. and O'Dea, C. (2026) Intergenerational altruism and transfers of time and money: A life-cycle perspective, *Working Paper*.

Borsch-Supan, A., Hajivassiliou, V. and Kotlikoff, L. J. (1992) Health, Children, and Elderly Living Arrangements: A Multiperiod-Multinomial Probit Model with Unobserved Heterogeneity and Autocorrelated Errors, in *Topics in the Economics of Aging*, University of Chicago Press, pp. 79–108.

Bray, M. (1999) *The shadow education system: Private tutoring and its implications for planners*, UNESCO International Institute for Educational Planning.

Bray, M. (2021) Shadow Education in Europe: Growing Prevalence, Underlying Forces, and Policy Implications, *ECNU Review of Education*, **4**, 442–475.

Bray, M. (2022) Shadow Education in Asia and the Pacific: Features and Implications of Private Supplementary Tutoring, in *International Handbook on Education Development in Asia-Pacific* (Eds.) W. O. Lee, P. Brown, A. L. Goodwin and A. Green, Springer Nature Singapore, Singapore, pp. 1–23.

Caucutt, E., Lochner, L., Mullins, J. and Park, Y. (2026) Child skill production: Accounting for parental and market-based time and goods investments, *Journal of Political Economy*, **134**, 150–209.

Caucutt, E. M. and Lochner, L. (2020) Early and Late Human Capital Investments, Borrowing Constraints, and the Family, *Journal of Political Economy*, **128**, 1065–1147.

Cheo, R. and Quah, E. (2005) Mothers, Maids and Tutors: An Empirical Evaluation of their Effect on Children's Academic Grades in Singapore, *Education Economics*, **13**, 269–285.

Cho, D., Kang, T., Kang, H. and Kim, S. (2024) College quality vs. post college efforts in elite firm hiring, *Unpublished Manuscript*.

Cooper, H., Robinson, J. C. and Patall, E. A. (2006) Does Homework Improve Academic Achievement? A Synthesis of Research, 1987–2003, *Review of Educational Research*, **76**, 1–62.

Cunha, F. and Heckman, J. (2007) The Technology of Skill Formation, *American Economic Review*, **97**, 31–47.

Cunha, F., Heckman, J. J. and Schennach, S. M. (2010) Estimating the technology of cognitive and noncognitive skill formation, *Econometrica*, **78**, 883–931.

Dang, H.-A. (2007) The determinants and impact of private tutoring classes in Vietnam, *Economics of Education Review*, **26**, 683–698.

Daruich, D. (2026) The macroeconomic consequences of early childhood development policies, *Journal of Political Economy*, **134**, 2023–2073.

De Groote, O. (2025) Dynamic effort choice in high school: Costs and benefits of an academic track, *Journal of Labor Economics*, **43**, 467–502.

Del Boca, D., Flinn, C., Verriest, E. and Wiswall, M. (2026) Parenting with patience: Parental incentives and child development, *Journal of Political Economy*, **134**, 210–284.

Del Boca, D., Flinn, C. and Wiswall, M. (2014) Household choices and child development, *Review of Economic Studies*, **81**, 137–185.

Del Boca, D., Monfardini, C. and Nicoletti, C. (2017) Parental and child time investments and the cognitive development of adolescents, *Journal of Labor Economics*, **35**, 565–608.

Doepke, M. and Zilibotti, F. (2017) Parenting With Style: Altruism and Paternalism in Intergenerational Preference Transmission, *Econometrica*, **85**, 1331–1371.

Epple, D., Figlio, D. and Romano, R. (2004) Competition between private and public schools: testing stratification and pricing predictions, *Journal of Public Economics*, **88**, 1215–1245.

Epple, D. and Romano, R. (2008) Educational vouchers and cream skimming, *International Economic Review*, **49**, 1395–1435.

Epple, D. and Romano, R. E. (1998) Competition between private and public schools, vouchers, and peer-group effects, *American Economic Review*, pp. 33–62.

Epple, D., Romano, R. E. and Urquiola, M. (2017) School vouchers: A survey of the economics literature, *Journal of Economic Literature*, **55**, 441–92.

Epple, D., Romano, R. E. and Urquiola, M. (2021) Is Education Different? A Review of the Voucher Literature and Lessons for Implementation, *The Routledge Handbook of the Economics of Education*, pp. 150–168.

Fang, H. and Liu, C. (2026) The confucian fertility paradox: Education competition and ultra-low fertility in east asia, Tech. rep., National Bureau of Economic Research.

Fu, C. and Mehta, N. (2018) Ability tracking, school and parental effort, and student achievement: A structural model and estimation, *Journal of Labor Economics*, **36**, 923–979.

Gayle, G.-L., Golan, L. and Soytas, M. A. (2022) What is the source of the intergenerational correlation in earnings?, *Journal of Monetary Economics*.

Geweke, J. (1988) Antithetic acceleration of Monte Carlo integration in Bayesian inference, *Journal of Econometrics*, **38**, 73–89.

Gong, Y., Lochner, L., Stinebrickner, R. and Stinebrickner, T. (2019) The Consumption Value of College, *NBER Working Paper Series*.

Grau, N. (2018) The impact of college admissions policies on the academic effort of high school students, *Economics of Education Review*, **65**, 58–92.

Gu, S. and Zhang, L. (2026) The macroeconomic consequences of competition for college admissions, *International Economic Review*.

Han, C., Kang, C. and Lee, S. H. (2016) Measuring Effort Incentives in a Tournament with many Participants: Theory and Application, *Economic Inquiry*, **54**, 1240–1250.

Ho, C. and Wang, Y. (2024) Social institutions and low birth rates, Economics and Statistics Working Paper 11-2024, Singapore Management University, School of Economics.

Hoekstra, M. (2009) The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach, *The Review of Economics and Statistics*, **91**, 717–724.

Jia, R. and Li, H. (2021) Just above the exam cutoff score: Elite college admission and wages in China, *Journal of Public Economics*, **196**, 104371.

Judd, K. L. (1985) The law of large numbers with a continuum of IID random variables, *Journal of Economic Theory*, **35**, 19–25.

Kang, C. (2007) Classroom peer effects and academic achievement: Quasi-randomization evidence from South Korea, *Journal of Urban Economics*, **61**, 458–495.

Kang, C., Koh, Y. and Park, Y. (2025) University regulations and private tutoring, *Working Paper*.

Kang, C. and Park, Y. (2021) Private Tutoring and Distribution of Student Academic Outcomes: An Implication of the Presence of Private Tutoring for Educational Inequality, *The Korean Economic Review*, **37**, 287–326.

Kim, D. (2026) The race for elite destinations: Education competition and low fertility in Korea, arXiv Working Paper 2608.27980, Simon Fraser University.

Kim, D. I. (2005) Growth in college education and wage differentials in Korea, *Seoul Journal of Economics*, **18**, 87–123.

Kim, S. (2014) Effects of academic cliques on the first job offers of college graduates, *Journal of Educational Studies*, **45**.

Kim, S. and Lee, J. (2010) Private Tutoring and Demand for Education in South Korea, *Economic Development and Cultural Change*, **58**, 259–296.

Kim, S. and Lee, J.-H. (2006) Changing Facets of Korean Higher Education: Market Competition and the Role of the State, *Higher Education*, **52**, 557–587.

Kim, S., Tertilt, M. and Yum, M. (2024) Status externalities in education and low birth rates in Korea, *American Economic Review*, **114**, 1576–1611.

Lavy, V., Palumbo, M. and Stern, S. (1998) Simulation of multinomial probit probabilities and imputation of missing data, *Advances in Econometrics*, **13**, 145–180.

Lazear, E. P. and Rosen, S. (1981) Rank-Order Tournaments as Optimum Labor Contracts, *Journal of Political Economy*, **89**, 841–864.

Lee, J. and Koh, Y. (2023) University ranking and the lifetime wage gap, *The Korean Journal of Economic Studies*, **71.2**.

Lee, S. Y. T. and Seshadri, A. (2019) On the Intergenerational Transmission of Economic Status, *Journal of Political Economy*, **127**, 855–921.

MacLeod, W. B., Riehl, E., Saavedra, J. E. and Urquiola, M. (2017) The Big Sort: College Reputation and Labor Market Outcomes, *American Economic Journal: Applied Economics*, **9**, 223–261.

Mahler, L., Tertilt, M. and Yum, M. (2025) *Policy concerns in an era of low fertility: The role of social comparisons and intensive parenting*, Centre for Economic Policy Research.

Nam, K. (2007) Time Series Trend of the Scale of Private Supplementary Tutoring, *The Journal of Economics and Finance of Education*, **16**, 57–79.

Ono, H. (2007) Does examination hell pay off ? A cost–benefit analysis of "ronin" and college education in Japan, *Economics of Education Review*, **26**, 271–284.

Park, H., Behrman, J. R. and Choi, J. (2013) Causal Effects of Single-Sex Schools on College Entrance Exams and College Attendance: Random Assignment in Seoul High Schools, *Demography*, **50**, 447–469.

Park, H., Behrman, J. R. and Choi, J. (2018) Do single-sex schools enhance students' STEM (science, technology, engineering, and mathematics) outcomes?, *Economics of Education Review*, **62**, 35–47.

Ramey, G. and Ramey, V. A. (2010) The Rug Rat Race, *Brookings Papers on Economic Activity*, **41**, 129–199.

Ryu, D. and Kang, C. (2013) Do Private Tutoring Expenditures Raise Academic Performance? Evidence from Middle School Students in South Korea: Do Tutoring Expenditures Raise Performance?, *Asian Economic Journal*, **27**, 59–83.

Sekhri, S. (2020) Prestige Matters: Wage Premium and Value Addition in Elite Colleges, *American Economic Journal: Applied Economics*, **12**, 207–225.

Stern, S. (1997) Simulation-Based Estimation, *Journal of Economic Literature*, **35**, 2006–2039.

Stevenson, D. L. and Baker, D. P. (1992) Shadow Education and Allocation in Formal Schooling: Transition to University in Japan, *American Journal of Sociology*, **97**, 1639–1657.

Stinebrickner, R. and Stinebrickner, T. R. (2008) The causal effect of studying on academic performance, *The BE Journal of Economic Analysis & Policy*, **8**.

Stinebrickner, T. R. (1999) Estimation of a duration model in the presence of missing data, *Review of Economics and Statistics*, **81**, 529–542.

Sullivan, P. (2009) Estimation of an Occupational Choice Model when Occupations are Misclassified, *Journal of Human Resources*, **44**, 495–535.

Sun, Y. (2006) The exact law of large numbers via Fubini extension and characterization of insurable risks, *Journal of Economic Theory*, **126**, 31–69.

Tansel, A. and Bircan, F. (2005) Effect of private tutoring on university entrance examination performance in Turkey, IZA Discussion Paper 1609, Institute for the Study of Labor.

Tincani, M., Kosse, F. and Miglino, E. (2022) Subjective beliefs and inclusion policies: Evidence from college admissions, CEPR Discussion Paper 17177, Centre for Economic Policy Research.

Todd, P. and Wolpin, K. I. (2018) Accounting for mathematics performance of high school students in Mexico: Estimating a coordination game in the classroom, *Journal of Political Economy*, **126**, 2608–2650.

Uway (2011) College admission cutoff scores by institution, Korean college-admissions information service, https://www.uway.com/, admission cutoff data for the 2011 entering cohort.

Zimmerman, S. D. (2019) Elite Colleges and Upward Mobility to Top Jobs and Top Incomes, *American Economic Review*, **109**, 1–47.

Appendices

Appendix A: Equilibrium Existence **50**

Appendix B: Member colleges for each college tier **54**

Appendix C: Descriptive Evidence **55**
Appendix C.1: Korean Labor and Income Panel Study 55
Appendix C.2: Selection Rules and Effects 55
Appendix C.3: Firm Interview Regression and Income Equation Estimates . . 56
Appendix C.4: Hours Allocation Regression 58
Appendix C.5: Moment Conditions from KLIPS 59

Appendix D: Likelihood Function Details **60**

Appendix E: Additional Estimation Results **63**

Appendix F: Model Fit **66**

Appendix G: Additional Counterfactual Analyses **68**

Appendix A: Equilibrium Existence

This appendix proves that there exists a recursive equilibrium of the model in Section 4. I define a map $\mathcal{T}$ that takes a guessed cutoff vector $\tilde{Q}$ to the cutoff vector that clears all tier seats when households best respond to $\tilde{Q}$. An equilibrium is exactly a fixed point $\mathcal{T}(\tilde{Q}^*) = \tilde{Q}^*$, and I show that such a point exists by Brouwer's theorem.

Actions and feasible sets. Let h denote the per-period time endowment (maximum available hours). In each period the household chooses an action

$$a = (d, p, x, s),$$

where $d \in \{0, 1\}$ indicates participation in private tutoring, p is the hourly tutoring price (quality), x is tutoring hours, and s is self-study hours. The household consumes $c = w - px$ and enjoys leisure $\ell = h - x - s$. A household that does not tutor ($d = 0$) sets $p = x = 0$. Fix constants $\underline{c} > 0$, $\underline{\ell} > 0$, $\underline{x} > 0$, and $0 < \underline{p} < \overline{p} < \infty$, whose role is stated in Assumption A.1. The no-tutoring and tutoring choice sets are

$$\begin{aligned}
\mathcal{A}^0 &= \{(d, p, x, s) = (0, 0, 0, s) : \ 0 \le s \le h - \underline{\ell}\}, \\
\mathcal{A}^1(w) &= \left\{(d, p, x, s) : \ d = 1, \ p \in [\underline{p}, \overline{p}], \ x \ge \underline{x}, \ s \ge 0, \ x + s \le h - \underline{\ell}, \ w - px \ge \underline{c}\right\},
\end{aligned}$$

and the household chooses from $\mathcal{A}(w) = \mathcal{A}^0 \cup \mathcal{A}^1(w)$. The lower bound $\underline{x} > 0$ is the smallest admissible level of tutoring hours and ensures $\log x$ is well defined when tutoring is chosen. Note that $\mathcal{A}^0$ is nonempty whenever $h > \underline{\ell}$ and delivers consumption $c = w$, so $\mathcal{A}(w)$ is nonempty for every $w \ge \underline{c}$; $\mathcal{A}^1(w)$ may be empty for low w, in which case the household does not tutor. Each $\mathcal{A}(w)$ is compact.

Let

$$\mathcal{Q} = \left\{Q \in [\underline{Q}, \overline{Q}]^{J-1} : Q_1 \ge Q_2 \ge \cdots \ge Q_{J-1}\right\}$$

denote the set of feasible cutoff vectors, a nonempty compact convex set.

Assumption A.1 (Equilibrium existence conditions)**.**

(a) *The time endowment satisfies $h > \underline{\ell}$, and household resources are finite and satisfy $w_{it} \ge \underline{c}$ for every household and period.*

(b) *Flow utilities and the laws of motion for the household states are continuous in states and choices.*

(c) *For each tier j, the lifetime-income payoff $v_j(q, \xi)$ is finite, strictly positive, and continuous. Expected current and continuation payoffs are finite and admit an integrable bound*

that is uniform over $\tilde{Q} \in \mathcal{Q}$.

(d) For every $\tilde{Q} \in \mathcal{Q}$, the household's optimal choice is unique for almost every realization of the household states and preference shocks.[49]

(e) Admission is determined by the log effective admission score $z_i = \ln \widehat{q_{i,M+1}} + \omega_i$, where $\ln \widehat{q_{i,M+1}}$ is the deterministic component of the terminal log test score given the household's initial conditions, shock history, and optimal choices, and $\omega_i \equiv \eta^q_{iM} + \varepsilon^m_i$ is the composite of the terminal test-score innovation and the admission noise. Equivalently, $z_i = \ln q_{i,M+1} + \varepsilon^m_i$, where $\ln q_{i,M+1} = \ln \widehat{q_{i,M+1}} + \eta^q_{iM}$ is the realized terminal score, as in the main text. ω_i has a strictly positive, continuous density on $\mathbb{R}$, and the cumulative seat shares satisfy $0 < \overline{N}_j < 1$, $j = 1, \ldots, J-1$.

(f) The bounds $\underline{Q}$ and $\overline{Q}$ are sufficiently wide that every seat-clearing cutoff vector generated from a guess in $\mathcal{Q}$ lies in $\mathcal{Q}$.

Lemma A.1. *Fix a guessed cutoff vector $\tilde{Q} \in \mathcal{Q}$.*

(a) The household dynamic programming problem admits a solution: in every period and at every state the maximum is attained on $\mathcal{A}(w_{it})$, and the value functions are finite.

(b) The value functions are continuous in $\tilde{Q}$ at every state, and jointly continuous in the household states and $\tilde{Q}$ at every state at which the tutoring branch is either infeasible or strictly feasible, i.e., $w_{it} \neq w_0 \equiv \underline{c} + \underline{p}\,\underline{x}$.

(c) The optimal policy correspondence is upper hemicontinuous in the states and $\tilde{Q}$ away from the threshold w_0, and at every realization of the states and shocks at which the optimizer is unique, the optimal policy is a continuous function of $\tilde{Q}$.

Proof. Each of the two choice sets, $\mathcal{A}^0$ and $\mathcal{A}^1(w_{it})$, is closed and bounded, so their union $\mathcal{A}(w_{it})$ is compact. Consumption and leisure are bounded away from zero, and tutoring hours are bounded away from zero whenever tutoring is chosen, so all logarithmic components of flow utility are continuous and bounded on the relevant choice sets.

In the final schooling period, the expected tournament payoff is continuous in household states, choices, and $\tilde{Q}$, and finite by Assumption A.1(c). I apply the theorem of the maximum separately to the maximization over $\mathcal{A}^0$ and over $\mathcal{A}^1(w_{iM})$: each objective is continuous, and each of the two constraint correspondences is nonempty, compact-valued, and continuous on its domain. A maximizer therefore exists in each nonempty

[49] I do not assume concavity of the household objective: the prize component is convex below and concave above each admission cutoff, and this non-concavity is the economic content of the tournament: investment responds most strongly near thresholds. Multiple local optima can therefore arise in principle, but exact indifference between them is a measure-zero event under continuously distributed shocks, so the optimal choice is unique for almost every realization.

subproblem, the two restricted value functions are continuous, and the final-period value function, as their pointwise maximum, is finite, continuous in $\tilde{Q}$ everywhere, and continuous in the states away from the income threshold w_0.

Suppose the period-$t+1$ value function is finite and continuous. Continuity of the laws of motion (Assumption A.1(b)), together with the integrable bound that is uniform over $\mathcal{Q}$ (Assumption A.1(c)), implies that the expected continuation value is finite and continuous by the dominated convergence theorem. The period-t objective is therefore continuous on a compact choice set, and the same argument applies unchanged. Backward induction establishes parts (a) and (b) for all $t = 1, \ldots, M$.

For part (c), the theorem of the maximum also yields upper hemicontinuity of the optimal-policy correspondence in the states and $\tilde{Q}$. At any realization of the states and shocks at which the maximizer is unique, an upper hemicontinuous correspondence that is singleton-valued is continuous; the optimal policy is therefore a continuous function of $\tilde{Q}$ at such realizations.

The overall value function is the pointwise maximum of the two restricted value functions. For $w_{it} < w_0$ only the no-tutoring branch is feasible, for $w_{it} > w_0$ both branches are and the maximum of two continuous functions is continuous; at $w_{it} = w_0$ the tutoring branch becomes feasible at the forced minimum bundle, so the value function may jump upward in the income direction. Continuity in $\tilde{Q}$ at every fixed state is unaffected, because the branch domains do not vary with $\tilde{Q}$.

□

Lemma A.2. *The cutoff update map $\mathcal{T} : \mathcal{Q} \to \mathcal{Q}$ is continuous.*

Proof. For a guessed cutoff vector $\tilde{Q}$, write the log effective admission score of Assumption A.1(e) as

$$z_i(\tilde{Q}) = \ln \widehat{q_{i,M+1}}(\tilde{Q}) + \omega_i,$$

where $\ln \widehat{q_{i,M+1}}(\tilde{Q})$ is the deterministic part given household i's initial conditions, shock history, and optimal choices under $\tilde{Q}$, and ω_i is the remaining noise, independent of it. Let F_ν and f_ν denote the distribution and density of ω_i, and let $F_z(\cdot; \tilde{Q})$ denote the cross-sectional distribution of $z_i(\tilde{Q})$. By independence,

$$\mathrm{F}_z(z; \tilde{Q}) = E\left[F_\nu\left(z - \ln \widehat{q_{i,M+1}}(\tilde{Q})\right)\right], \qquad f_z(z; \tilde{Q}) = E\left[f_\nu\left(z - \ln \widehat{q_{i,M+1}}(\tilde{Q})\right)\right] > 0,$$

where the expectation is over initial conditions and all preceding shocks, and positivity follows from $f_\nu > 0$ (Assumption A.1(e)). Hence $F_z(\cdot; \tilde{Q})$ is continuous and strictly increasing, and every seat-clearing cutoff is uniquely defined.

Consider any sequence $\tilde{Q}^n \to \tilde{Q}$. Couple the household problems across n using the same initial conditions and shock histories. Because the coupling holds each household's income path fixed, only continuity in $\tilde{Q}$ at fixed states is used, which Lemma A.1 delivers at every state.

By Lemma A.1(c), the optimal policies converge at every realization at which the limit problem's optimizer is unique, which by Assumption A.1(d) is almost every realization. Continuity of the laws of motion then implies $\ln \widehat{q_{i,M+1}}(\tilde{Q}^n) \to \ln \widehat{q_{i,M+1}}(\tilde{Q})$ almost surely, and the dominated convergence theorem gives $F_z(z; \tilde{Q}^n) \to F_z(z; \tilde{Q})$ at every z.

The updated cutoff between tiers j and $j+1$ is the $(1 - \overline{N}_j)$-quantile of the score distribution,

$$\big[\mathcal{T}(\tilde{Q})\big]_j = \exp\big\{F_z^{-1}(1 - \overline{N}_j; \tilde{Q})\big\}, \qquad j = 1, \ldots, J-1,$$

where $[\,\cdot\,]_j$ denotes the j-th component; F_z is the distribution of the log admission index, so the exponential returns the cutoff in levels, consistent with the units of $\tilde{Q}$ in the main text. Pointwise convergence of distribution functions to a continuous, strictly increasing limit implies convergence of the corresponding quantiles, so $\mathcal{T}(\tilde{Q}^n) \to \mathcal{T}(\tilde{Q})$: the map is continuous. Finally, since $\overline{N}_1 < \overline{N}_2 < \cdots < \overline{N}_{J-1}$, the updated cutoffs are strictly ordered, and by Assumption A.1(f) they lie in $\mathcal{Q}$.

□

Theorem A.3. *A recursive tournament equilibrium exists.*

Proof. The set $\mathcal{Q}$ is nonempty, compact, and convex. By the preceding lemma, the cutoff update map

$$\mathcal{T} : \mathcal{Q} \to \mathcal{Q}$$

is continuous. Brouwer's fixed-point theorem therefore implies that there exists a cutoff vector $\tilde{Q}^* \in \mathcal{Q}$ such that

$$\mathcal{T}(\tilde{Q}^*) = \tilde{Q}^*.$$

Given $\tilde{Q}^*$, the associated value functions and policy rules solve the household dynamic programming problem. Starting from the initial distribution, these policy rules generate the distribution path and the distribution of effective admission scores that reproduce $\tilde{Q}^*$. The resulting value functions, policy rules, distribution path, and cutoff vector therefore constitute a recursive tournament equilibrium. □

Appendix B: Member colleges for each college tier

	List of the member colleges
First Tier	Seoul National, Yonsei, Korea , Sogang, SKKU, Hanyang, KAIST, Pusan, Ewha, Postech
Second Tier	Choongang, Kyunghee, HUFS, University of Seoul, KU, Dongguk,
	Kyongpook, Sookmyung, Ajou, Honggik, Inha, Hangkong, Kookmin
Third Tier	Soongsil, Sejong, Dankook, Kwangwoon, Cheonnam, Seoul Industrial University
	Myongji, Sangmyeong, Catholic, Choongam, Choongbook, Seongshin, Kyeongki
	Kyongwon, Deoksong women, Dongdeok women, Dong-A, Bookyeong
Fourth Tier	The rest of the colleges, 2 year colleges, and high school graduates

Appendix C: Descriptive Evidence

Appendix C.1: Korean Labor and Income Panel Study

The college tier-specific lifetime income is inferred from the Korean Labor Income and Panel Study (KLIPS). KLIPS is a panel dataset of representative Korean households from 1998 to 2021. The dataset provides information on which college each worker graduated from, her major, income history, and other demographic characteristics. Using KLIPS, I generate the average lifetime income of the alumni for each college tier and complement the labor market information of KELS. In fact, KELS also provides individual information on the early labor market outcomes of the sample. Still, both the income data and the participation data have a substantial proportion of missing data compared to KLIPS. Employing KLIPS is more useful in predicting alumni's lifetime income as it contains data on workers of age between 20 and 65.[50]

Appendix C.2: Selection Rules and Effects

Table C.1: Data Selection

	Dropped	Remaining
Original sample (KELS panel)		**6,908**
Arts, music, physical-education track	423	**6,485**
Missing 6th-grade test score	247	
Missing parental education	484	**5,754**
All six income waves missing	36	
Income missing after imputation	373	**5,345**
Tutoring expenditure exceeds income	41	
Implausible tutoring unit price	51	**5,253**
Neither CSAT nor adult earnings observed	1,729	**3,524**

Notes. The estimation sample is reached in three stages: the arts and physical-education track is removed because admission for those students does not run through the CSAT tournament. The final stage keeps a household if *either* a CSAT score *or* adult earnings is observed. Households with missing choice variables are included. The estimator simulates the missing values rather than dropping the observation.

Test-score construction. All test scores enter the model in a common metric anchored to the CSAT. The discrete grade-6 performance measure is mapped to a continuous initial score by a rank-based inverse normal transformation, with ties assigned their mid-rank percentile, and rescaled linearly to the CSAT mean and standard deviation. The grade 7–9 national achievement scores are rescaled linearly to the same CSAT units;

[50]The Lifelong Career Survey (LCS) by the Korea Research Institute for Vocational Education & Training (KRIVET) is an alternative dataset that could be used to generate the proxy of the prize of the tournament (Han, Kang and Lee, 2016). For the purpose of this paper, KLIPS is preferred because it can recover the age-specific income profile.

students observed on a subset of subjects are first rescaled to the full-subject metric. The grade-12 CSAT enters in raw points, summed over the three core sections and two electives.[51]

Appendix C.3: Firm Interview Regression and Income Equation Estimates

Table C.2: Final Interview Request Regression (Probit)

	(1)	(2)	(3)
Tier=1	0.327***	0.737***	0.760***
	(0.038)	(0.103)	(0.113)
Tier=2	0.029	0.140***	0.161***
	(0.044)	(0.030)	(0.035)
Tier=3	-0.123	-0.018	-0.036
	(0.096)	(0.127)	(0.109)
ColGPA	0.006**	0.021***	0.022***
	(0.002)	(0.007)	(0.006)
N	9138	9132	9132
Major	No	Yes	Yes
Company	Yes	No	Yes
ClusterSE	No	Yes	Yes

Source: Confidential hiring data of a major Korean conglomerate, late 2010s (Cho, Kang, Kang and Kim, 2024).
Note: The sample is applicants to the subsidiary firms of the conglomerate over three years. Other explanatory variables include the subsidiary firm's characteristics and the applicant's college major, age, and gender. Tier 4 and below is the omitted category. ColGPA is the standardized college GPA (0–100 scale).

[51]The model treats the test score as the observed contest variable rather than a latent skill (Agostinelli and Wiswall, 2025): the tournament ranks students within a grade, the terminal CSAT enters in raw units.

Table C.3: College Tiers and Earnings: Initial (KELS) and Lifetime (KLIPS)

	(1) Base	(2) +CSAT	(3) +CSAT, Major
Panel A. Initial (log) earnings — KELS, age ≈ 28			
Tier 1	0.342***	0.166***	0.166***
	(0.044)	(0.052)	(0.049)
Tier 2	0.174***	0.071*	0.101***
	(0.032)	(0.037)	(0.036)
Tier 3	0.048	-0.040	-0.003
	(0.038)	(0.040)	(0.039)
Observations	1,199	1,199	1,199
Panel B. Lifetime (log) earnings, tier × age — KLIPS			
Tier 1	-1.856***	-2.561***	-2.924***
	(0.343)	(0.515)	(0.469)
Tier 2	-1.275***	-1.892	-1.608
	(0.334)	(1.389)	(1.346)
Tier 3	-0.812**	-0.414	0.422
	(0.279)	(0.592)	(0.722)
Tier 1 × age	0.065***	0.096***	0.110***
	(0.010)	(0.018)	(0.017)
Tier 2 × age	0.048***	0.077	0.067
	(0.010)	(0.044)	(0.043)
Tier 3 × age	0.029**	0.017	-0.017
	(0.008)	(0.022)	(0.026)
age	0.097***	0.279**	0.273**
	(0.006)	(0.089)	(0.088)
Observations	29,599	685	685
CSAT / entrance score	No	Yes	Yes
Major / field	No	No	Yes

Note: Panel A: OLS of log monthly income at labor-market entry on college-tier dummies (KELS; Tier 4 = baseline; gender controlled in all columns; robust standard errors). Panel B: pooled-OLS of log annual income on tier dummies and their interaction with age (KLIPS, 1998–2018; Tier 4 = baseline; age^2, birth year, and gender controlled in all columns; standard errors clustered at the college-major level). Columns add a CSAT / entrance-score control (Col. 2) and field of study (Col. 3); the KLIPS entrance-score subsample is small ($N = 685$). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C.4: Hours Allocation Regression

Table C.4: The Effects of Parental Background on Hours of Self-Study

	(1) log(1+Study)	(2) log(1+Study)	(3) log(1+Study)
log(Income)	0.238***	0.152***	0.036
	(0.022)	(0.025)	(0.027)
Parental education		0.055***	
		(0.007)	
N	10454	10454	10454
Year	Yes	Yes	Yes
Household FE	No	No	Yes

Source: Korea Educational Longitudinal Study 2005, Korean Educational Development Institute.
Note: log(1+Study) is the log of weekly hours of self-study plus one. I include only households with no missing information on tutoring expenditure, CSAT scores, and household income. Parental education is the average years of the parents' education.

Table C.5: The Effects of Parental Background on Hours of Tutoring

	(4) log(1+Tutoring)	(5) log(1+Tutoring)	(6) log(1+Tutoring)
log(Income)	0.631***	0.546***	0.215***
	(0.018)	(0.019)	(0.025)
Parental education		0.061***	
		(0.006)	
N	18437	18437	18432
Household FE	No	No	Yes
Year FE	No	No	Yes

Source: Korea Educational Longitudinal Study 2005, Korean Educational Development Institute.
Note: log(1+Tutoring) is the log of weekly hours of tutoring plus one. I include only households with no missing information on tutoring expenditure, CSAT scores, and household income. Parental education is the average years of the parents' education.

Parental education soaks up significant variation in hours of self-study, which leaves a relatively small variation with parental income. Tables C.4 and C.5 present the pooled OLS estimates of the regression equation,

$$\ln(1 + y_{it}) = \beta_0 + \beta_1 \log(hhinc_{it}) + \beta_2 m_i + \epsilon_{it} \tag{20}$$

where $hhinc_{it}$ is the income and m_i is parental education of household i. Columns (1) through (3) present the results where y_{it} is hours of self-study, and columns (4) through (6) present the results where y_{it} is hours of tutoring. Columns (1) and (4)

provide the estimates without including the average years of parents' education, and Columns (2) and (5) provide the estimates including the average years of parents' education to equation (20). Overall, hours of tutoring are explained more by parents' income than hours of self-study. Moreover, the covariation between hours of self-study and income falls after controlling for the average years of parents' education, and only the household fixed-effects estimate is small: the income coefficient remains 0.15 with parental-education controls and drops to 0.036 within households.

Appendix C.5: Moment Conditions from KLIPS

A key challenge in estimating the income equation (3) is that KELS follows respondents only into their late twenties: it observes early-career earnings around age 28, which enter the likelihood directly, but not enough of the life-cycle to identify the tier-specific age profiles. I therefore augment the likelihood with moment conditions constructed from the Korean Labor and Income Panel Study (KLIPS).

Construction. Define the mean log income for age bin b and tier j as $\mu_{bj} = \mathbb{E}[\ln y \mid \text{age bin} = b, \text{tier} = j]$, where age is measured in years since labor-market entry and bins are five years wide, with midpoints ranging from about 2 to 27 years after entry. The estimation targets the 21 tier-by-age-bin cell means with sufficient KLIPS observations: six cells each for Tiers 1 and 4, five for Tier 2, and four for Tier 3, with cell sizes between roughly 40 and 600 workers. Each cell enters the penalty weighted by the inverse variance of its mean, n_{bj}/σ_j^2.

Normalization and identification. Because KELS observes the cohort only in its early career, the lifecycle age-profile parameters are identified from KLIPS. The estimation matches mean log income in 21 tier-by-age-bin cells, computed on the model side from the age profile $g_j(a)$ and the simulated cohort's tier-conditional skill means, which are held constant across age bins. The four tier-specific intercepts are estimated, so the overall level of each tier's KLIPS profile is fitted rather than forced onto the profile parameters. The cells then identify two distinct sets of objects: cross-tier differences in the cell levels, net of the model-implied tier differences in terminal skills, pin the intercept gaps that enter the lifetime prize, while the remaining 17 within-tier contrasts over-identify the five parameters of $g_j(a)$ in equation (3) – the common slope, the three tier-specific slope shifts (Tier 4 is the normalization), and the common curvature. KELS separately identifies the entry-level tier premiums τ_j in the early-career earnings equation; the age-profile estimates and the implied lifetime tier gaps are reported in Table E.3.

Appendix D: Likelihood Function Details

The likelihood is built from the structural shocks. For each observed choice bundle I invert the household's optimality conditions to recover the shock realizations that rationalize the observed choices, and transform their joint density with the Jacobian of the inversion. The integrals that remain (over participation shocks, missing choices, and latent test scores) are evaluated by Monte Carlo simulation. Borsch-Supan, Hajivassiliou and Kotlikoff (1992) show that maximum simulated likelihood performs well with a moderate number of draws under a good simulation design; I use antithetic acceleration (Geweke, 1988; Stern, 1997) and common random numbers across parameter evaluations.

Shock Inversion

In period t, a household that participates in tutoring chooses (p_{it}, x_{it}, s_{it}); consumption follows from the period budget $c_{it} = w_{it} - p_{it}x_{it}$. Given the state and the continuation value, the first-order conditions with respect to p_{it}, x_{it}, and s_{it}, together with the test-score equation (5), define a one-to-one mapping between the observed vector $\mathbf{y}_{it} = (p_{it}, x_{it}, s_{it}, q_{i,t+1})$ and the shock vector $\boldsymbol{\eta}_{it} = (\eta^c_{it}, \eta^x_{it}, \eta^s_{it}, \eta^q_{it})$. Each first-order condition balances the marginal utility of the choice (consumption, the taste for tutoring hours, the taste for study, and the leisure term, which enters the hours conditions with weight α_ℓ) against the marginal continuation value of the score. I solve this system for the implied shocks and work with their centered innovations $\tilde{\boldsymbol{\eta}}_{it} \equiv \boldsymbol{\eta}_{it} - \boldsymbol{\mu}_{\eta t}$, where the period-specific mean $\boldsymbol{\mu}_{\eta t}$ is zero in every component except the grade-12 study shock. The contribution of a fully observed tutoring period is

$$f(\mathbf{y}_{it}) = \phi_4(\tilde{\boldsymbol{\eta}}_{it}; \Omega) \cdot \left| \det \frac{\partial \tilde{\boldsymbol{\eta}}_{it}}{\partial \mathbf{y}_{it}} \right|,$$

where $\phi_4(\cdot; \Omega)$ is the mean-zero multivariate normal density. The covariance Ω allows correlation between the hours shocks (ρ_{xs}) and between the consumption shock and each hours shock (ρ_{cx}, ρ_{cs}); the score shock is independent. The Jacobian is computed analytically by differentiating the first-order conditions, with the continuation-value derivatives evaluated on the interpolated value function; the resulting expressions are lengthy but mechanical, and I omit them here (they are contained in the replication package). For households that do not participate in tutoring, the same construction applies to the no-tutoring problem: the observed $(s_{it}, q_{i,t+1})$ are inverted to $(\tilde{\eta}^s_{it}, \tilde{\eta}^q_{it})$ and the consumption and tutoring-taste shocks are integrated out. Zero study hours are handled by censoring: the contribution replaces the study density with the probability that η^s_{it} falls below the threshold at which the optimal s_{it} reaches zero, computed from

the first-order condition for s_{it}.

Tutoring Participation

Participation in tutoring is subject to a fixed utility cost $\chi_{r(i)}$ that depends on the local tutoring-supply quintile $r(i)$, and to a participation shock $\sigma_d \zeta_{it}^d$. A household participates when the value gain from tutoring, net of the cost, exceeds the shock. The likelihood therefore multiplies each period's choice density by the probability of the observed participation status. The participation shock is allowed to load on the hours shocks: conditional on $(\eta_{it}^x, \eta_{it}^s)$, ζ_{it}^d is normal with mean $\rho\, z_{it}^x + b_{ns}\, z_{it}^s$ (in standardized units) and the complementary variance, so selection into tutoring on unobserved tastes is captured. Search costs, quadratic in tutoring quality with quintile-specific curvature, enter the tutoring branch of the value comparison. The participation probability requires the value of the tutoring option at the optimum, which is computed in the same inner optimization used to solve the model.

Terminal Period and Tier Assignment

In the final period the continuation value is the tournament payoff. Because the score slope of log earnings is common across tiers, the individual skill terms $b_{z0} \ln q_{i,M+1} + b_h \xi_{i,M+1}$ are additively separable, and the tournament component of the expected prize takes the telescoped form

$$\ln v_J + \sum_{j=1}^{J-1} \left(\ln v_j - \ln v_{j+1} \right) \Phi\Big(\frac{\ln \widehat{q}_{i,M+1} - \ln \tilde{Q}_j}{\sigma_{\text{total}}} \Big),$$

where $\ln \widehat{q}_{i,M+1}$ is the deterministic component of the final log test score, $\tilde{Q}_j$ are the admission cutoffs, and $\sigma_{\text{total}} = (\sigma_q^2 + \sigma_a^2)^{1/2}$ combines the score shock with the assignment noise σ_a. The derivative of this expression with respect to $\ln q_{i,M+1}$, together with the constant skill slope, enters the terminal first-order conditions. Observed college tiers contribute a separate assignment likelihood: conditional on the observed CSAT score, the probability that student i attends tier j is the normal probability that the score plus assignment noise falls between the cutoffs $\tilde{Q}_j$ and $\tilde{Q}_{j-1}$, which identifies σ_a from students whose observed tier differs from the deterministic cutoff assignment.

Contribution Classes and Missing Data

The estimation sample keeps every household with observed initial conditions and at least one terminal outcome, so households contribute through one of three classes. (i) Households with complete choice and score data contribute the full product of pe-

riod densities. (ii) Households with observed choices but missing test scores contribute the choice densities with the score integrated out: for each simulation draw the missing score is generated from equation (5) and carried forward as the state variable, so the household's subsequent contributions are evaluated on the simulated state. Grade 10 and 11 scores, which the survey does not measure for any student, are treated this way for the whole sample. (iii) Households with missing choice variables have the missing choices simulated from the model: for each draw of the corresponding shocks the household solves its optimization problem, and the optimized choices enter the likelihood and the score transition (Lavy, Palumbo and Stern, 1998; Stinebrickner, 1999; Sullivan, 2009).

Unobserved Heterogeneity and Auxiliary Blocks

Permanent heterogeneity enters through $K = 3$ discrete household types with taste shifters on consumption, study, and tutoring and unrestricted probabilities. The type is integrated out at the household level: the household's likelihood is the probability-weighted sum over types of the product of its period contributions. Three auxiliary blocks complete the likelihood. The non-cognitive measurement system contributes a Gaussian panel likelihood for the six-wave battery, built from the skill transition equation and the measurement normalization described in Section 5.2. The earnings block contributes the density of observed early-career log earnings given the terminal score, non-cognitive skill, and college tier. Finally, the tier-specific lifetime income paths are disciplined by matching simulated tier-by-age income cells to their KLIPS counterparts through a penalized-likelihood term, so the prizes v_j are estimated jointly with the behavioral parameters.

Appendix E: Additional Estimation Results

Table E.1: Parameter Estimates: Shocks and Unobserved Heterogeneity

Panel A. Shock standard deviations and correlations					
Score Shock SD (σ_q)	0.112	(0.002)	Corr(η^x, η^s) (ρ_{xs})	0.280	(0.012)
Consumption Shock SD (σ_c)	0.848	(0.005)	Corr(η^c, η^x) (ρ_{cx})	0.392	(0.010)
Tutoring-Hours Shock SD (σ_x)	0.623	(0.004)	Corr(η^c, η^s) (ρ_{cs})	-0.034	(0.010)
Study Shock SD (σ_s)	0.536	(0.008)			
Panel B. Household types ($K = 3$; Type 1 is the baseline)					
	Type 1	Type 2	Type 3		
Probability (π_k)	0.219	0.337	0.444		
	(0.008)	(0.007)	(0.003)		
Consumption Shifter (λ_k^c)	–	0.405	0.111		
		(0.028)	(0.016)		
Study Shifter (λ_k^s)	–	0.528	0.000		
		(0.012)	(0.000)		
Tutoring Shifter (λ_k^x)	–	0.359	0.335		
		(0.007)	(0.005)		

Note: Standard errors in parentheses: household-block bootstrap with re-estimation in each replicate. Panel A: the score shock is independent of the choice shocks; the choice-shock correlations $(\rho_{xs}, \rho_{cx}, \rho_{cs})$ complete the covariance Ω of the likelihood. Panel B: permanent unobserved heterogeneity follows a three-type mixture. Type 1 is the baseline; the shifters scale the tastes proportionally, $\alpha_z e^{\lambda_k^z}$ for $z = c, x, s$. The study component is estimated as an additive offset to α_s and reported here as the equivalent log-multiplier, with its standard error computed from the same bootstrap replicates on the converted scale.

Table E.2: Non-Cognitive Measurement System: Loadings and Signal Shares

			Signal share by wave (%)					
Measure	λ_m	σ_m	Gr. 7	Gr. 8	Gr. 9	Gr. 10	Gr. 11	Gr. 12
Self-concept batteries								
Social (anchor)	1.000	0.478	46.5	51.5	51.9	—	—	49.5
Family	0.845	0.697	22.6	26.3	26.6	—	—	24.7
Body	0.921	0.696	25.8	29.9	30.2	—	—	28.2
Peer-relations parcel								
Peer relations	1.634	0.415/0.394	—	—	—	64.1	69.8	—
sd of the measured stock ξ_{t+1}			0.445	0.493	0.496	0.339	0.367	0.473

Note: The signal share is $\lambda_m^2 \mathrm{Var}(\xi_{t+1}) / [\lambda_m^2 \mathrm{Var}(\xi_{t+1}) + \sigma_m^2]$, the fraction of each measure's variance attributable to the latent non-cognitive skill. The wave-t battery is administered after the period's investments and measures the stock they produce, ξ_{t+1}; the grade-12 battery therefore measures the terminal skill ξ_{M+1} that enters the earnings equation. The factor is normalized once, on the grade-7 social self-concept battery ($\lambda_{\text{soc}} = 1$, $\mu_{\text{soc}} = 0$); the location and scale of ξ_{t+1} in later waves are identified by the assumption that this battery, administered with identical items and scales in every wave in which it appears (grades 7, 8, 9, and 12), is age-invariant, which equates its loading and intercept to their grade-7 values (Agostinelli and Wiswall, 2025). Measurement-error variances are identified from contemporaneous cross-measure covariances within wave. $\mathrm{Var}(\xi_{t+1})$ follows from the estimated initial dispersion, grade-specific self-productivities $\gamma_{\xi,1,t}$, and innovation variance; investment-input contributions are negligible ($< 10^{-3}$ of the total) and omitted. The peer-relations parcels of grades 10 and 11 share a common loading with wave-specific error dispersions (0.415 and 0.394); their signal shares differ mainly because the grade-10 self-productivity ($\gamma_{\xi,1,4} = 0.49$) compresses the latent variance measured in grade 10. Dashes indicate waves in which the measure is not collected.

Table E.3: Age-Profile Parameters and Implied Lifetime Tier Gaps

Panel A. Age profile $g_j(a)$ in equation (3)	
Common slope (b_{a1})	0.0271
	(0.0005)
Tier 1 slope shift ($b_{a1,1}$)	0.0314
	(0.0008)
Tier 2 slope shift ($b_{a1,2}$)	0.0244
	(0.0013)
Tier 3 slope shift ($b_{a1,3}$)	0.0193
	(0.0015)
Common curvature (b_{a2})	-0.0009
	(<0.0001)
Panel B. Tier intercepts of the KLIPS profile	
Tier 1 intercept (a_1)	2.993
	(0.015)
Tier 2 intercept (a_2)	2.891
	(0.018)
Tier 3 intercept (a_3)	2.922
	(0.016)
Tier 4 intercept (a_4)	2.683
	(0.012)
Panel C. Implied gap in log lifetime income relative to Tier 4	
$\ln v_{i1} - \ln v_{i4}$	0.700
	(0.017)
$\ln v_{i2} - \ln v_{i4}$	0.504
	(0.018)
$\ln v_{i3} - \ln v_{i4}$	0.470
	(0.020)

Note: Panel A: parameters of the tier-specific age profile $g_j(a) = (b_{a1} + b_{a1,j})\tilde{a} + b_{a2}\tilde{a}^2$, where $\tilde{a} = a - a^*$ is years since labor-market entry and $b_{a1,4} = 0$; identified by the 21 KLIPS tier-by-age mean-income cells described in Appendix C.5. Panel B: the tier-specific intercepts a_j, which fit the level of each tier's KLIPS profile; their cross-tier differences, not the KELS entry premiums τ_j of equation (3), determine the lifetime gaps in Panel C, while the τ_j are identified separately from KELS early-career earnings. Panel C: the gap in log lifetime income, equation (2), between tier j and Tier 4, computed at $\beta = 0.95$ over the thirty-one post-entry years $\tilde{a} = 0, \ldots, 30$. Because equation (3) is additively separable in the skill terms, these gaps are common across households; the level of v_{ij} varies with $q_{i,M+1}$ and $\xi_{i,M+1}$. Standard errors in parentheses, from the same household-block bootstrap replicates as Table 2.

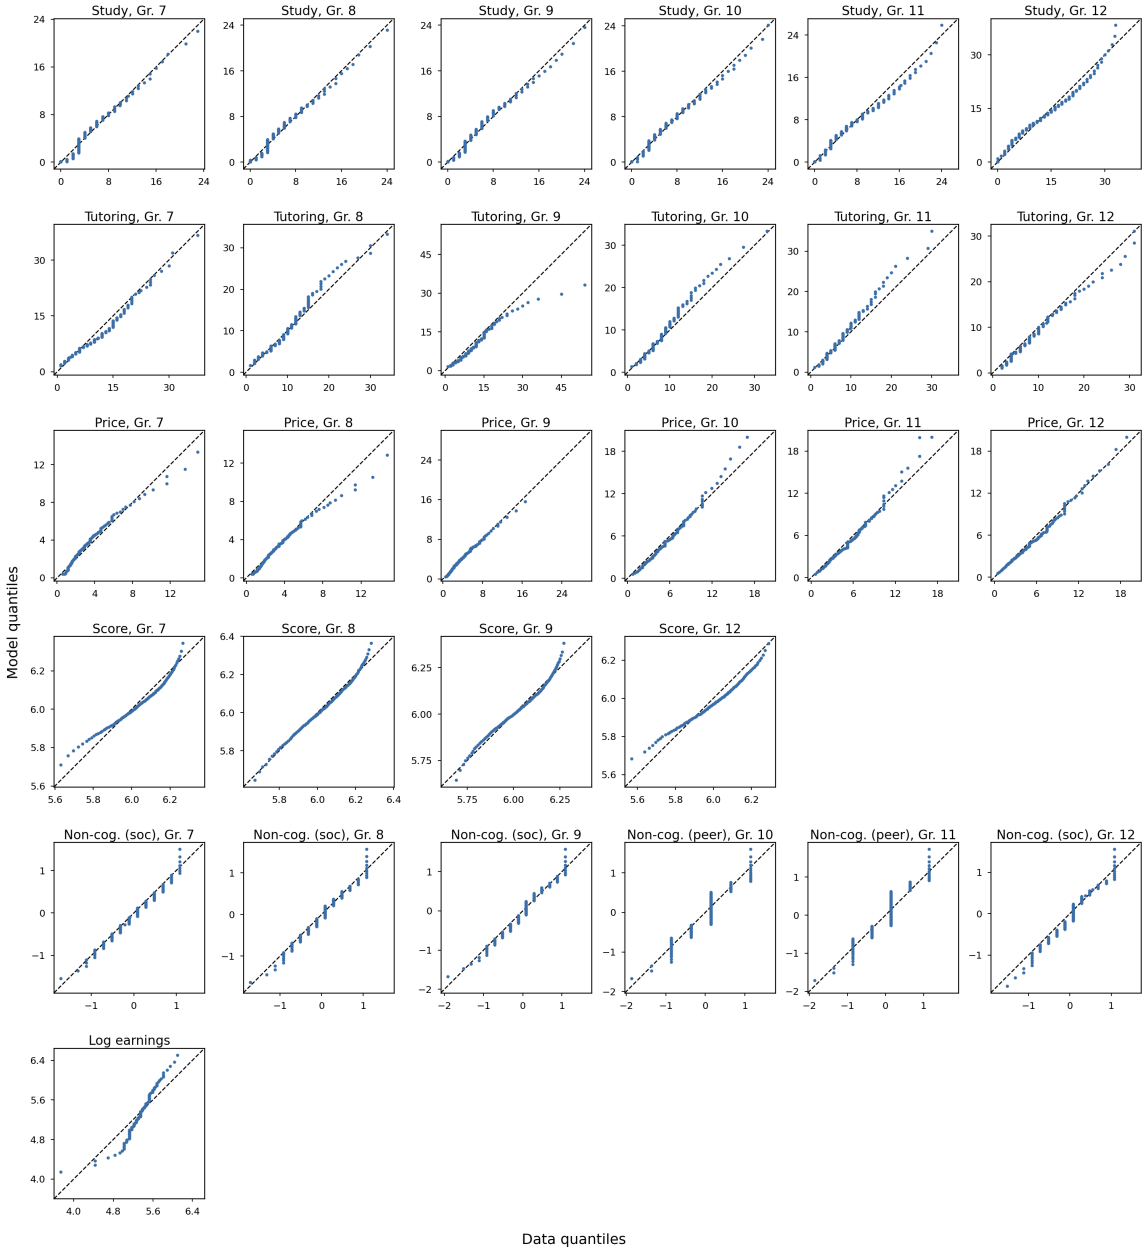
Appendix F: Model Fit

Table F.1: Model Fit: Means and Standard Deviations

	Gr. 7	Gr. 8	Gr. 9	Gr. 10	Gr. 11	Gr. 12
Tutoring participation (rate)						
Data	0.780	0.790	0.710	0.740	0.693	0.417
	(0.414)	(0.407)	(0.454)	(0.439)	(0.461)	(0.493)
Model	0.771	0.785	0.731	0.745	0.685	0.523
	(0.420)	(0.411)	(0.444)	(0.436)	(0.464)	(0.499)
Unit price of tutoring (participants, 10,000 KRW)						
Data	3.75	3.89	4.61	6.25	6.31	6.43
	(2.98)	(2.90)	(3.30)	(3.48)	(3.51)	(3.94)
Model	3.87	3.73	4.60	6.00	6.02	6.02
	(2.79)	(2.73)	(3.17)	(3.93)	(3.99)	(3.93)
Tutoring hours per week (participants)						
Data	12.73	11.32	12.89	9.51	8.97	9.93
	(7.74)	(6.62)	(9.12)	(5.93)	(5.78)	(6.99)
Model	11.76	11.85	11.20	10.48	10.09	8.82
	(7.44)	(7.39)	(7.33)	(7.12)	(7.09)	(6.50)
Self-study hours per week (all households)						
Data	5.48	5.93	6.30	7.34	8.19	13.51
	(5.39)	(5.20)	(5.41)	(5.70)	(6.11)	(9.31)
Model	5.57	6.00	6.42	7.46	8.13	13.81
	(5.33)	(5.25)	(5.41)	(5.54)	(5.75)	(9.11)
Log test score (all households)						
Data	6.034	6.032	6.033	–	–	5.999
	(0.153)	(0.155)	(0.155)			(0.170)
Model	6.035	6.033	6.035	5.996	5.997	5.985
	(0.133)	(0.148)	(0.154)	(0.187)	(0.162)	(0.146)

Note: Model moments come from the same baseline simulation as the decomposition and policy tables (Tables 4 and 5): a one-cohort forward simulation at the estimated parameters with admission cutoffs re-cleared so that tier enrollment equals the observed seat counts. Test scores are carried period to period and tutoring participation is drawn from the model hurdle. Model moments use the full simulated cohort; data moments are computed over the estimation sample's choice-usable observations, the rows that enter the likelihood, so small compositional differences between the columns are expected, and the data columns can differ from the descriptive statistics of Table 1, which use all available responses in each wave. Standard deviations in parentheses. Tutoring and study hours are weekly; the tutoring price is in units of 10,000 KRW ($\approx$ USD 9 at the 2005–2010 average exchange rate); test scores are in logs, and the grade-12 model dispersion includes the test-day shock (σ_q). Grade 10 and 11 test scores are latent, so no data moments exist there (–). Statistics for tutoring hours and the tutoring price are computed over participants (data: observed participants; model: simulated participants); participation, study hours, and test scores use all households.

Figure F.1: Model Fit: Quantile–Quantile Plots



Note: Each panel plots the quantiles of the simulated distribution (vertical axis) against the quantiles of the empirical distribution (horizontal axis); the 45-degree line indicates an exact distributional fit. Simulated distributions come from the one-cohort forward simulation of the estimated model described in Section 6.2: test scores are carried period to period and tutoring participation is drawn from the model hurdle. Panels cover weekly self-study hours, weekly tutoring hours, and tutoring unit prices by grade (the tutoring-hours and price panels condition on participation), log test scores in the grades where scores are observed, the non-cognitive measures by wave, and log early-career earnings. Axes are in the natural units of each variable; the grade-9 price axis is truncated at extreme quantiles for readability.

Appendix G: Additional Counterfactual Analyses

Table G.1: Tier-1 Seat Expansion: The Tutoring Response over Expansion Size and Premium Elasticity

	Elasticity σ of the Tier-1 premium response					
Expansion a	1.4	2.1	3.0	4.0	6.0	No dilution
+2.5%	−0.9	−0.5	−0.3	−0.1	−0.0	+0.3
+5%	−1.9	−1.0	−0.6	−0.3	−0.0	+0.5
+10%	−3.8	−2.1	−1.2	−0.6	−0.1	+1.1
+20%	−7.5	−4.3	−2.3	−1.2	−0.1	+2.1
+50% (main)	—	**−10.9**	−6.1	−3.3	−0.5	+5.1
+100%	—	—	—	−7.0	−1.6	+9.4

Note: Each cell reports the percentage change in mean tutoring expenditure (grades 7–12, unconditional, baseline 29.4 in units of 10,000 KRW per month, ≈ USD 9 at the 2005–2010 average exchange rate) when Tier-1 capacity expands by a with admission cutoffs re-clearing, Tier-2 and Tier-3 seat masses held fixed, and the Tier-1 log lifetime premium reduced by $\ln(1+a)/\sigma$. The bold cell is the main experiment of Table 5. Dashes mark cells excluded because the implied premium decline exceeds the Tier 1–Tier 2 gap ($\ln(1+a)/\sigma \geq 0.196$), which would invert the prize ordering; no parameter is adjusted to restore excluded cells. Conventions as in Table 5.