THE ECONOMIC JOURNAL

The Economic Journal, **136** (*January*), 335–350 https://doi.org/10.1093/ej/ueaf045
Advance Access Publication Date: 18 June 2025

PATIENCE AND SUBNATIONAL DIFFERENCES IN HUMAN CAPITAL: REGIONAL ANALYSIS WITH FACEBOOK INTERESTS*

Eric A. Hanushek, Lavinia Kinne, Pietro Sancassani and Ludger Woessmann

Decisions to invest in human capital depend on people's time preferences. This paper shows that differences in patience are closely related to substantial subnational differences in educational achievement, leading to new perspectives on long-standing within-country disparities. We use social media data—Facebook interests—to construct novel regional measures of patience within Italy and the United States. The approach is first validated with a cross-country analysis of patience and Facebook interests. We then show that patience is strongly positively associated with student achievement across regions in both countries, accounting for three-quarters of the achievement variation across Italian regions and one-third across US states. The finding is confirmed in an identification strategy employing variation in ancestry countries of the current populations of US states. Results hold for six other countries with more limited regional achievement data.

JEL codes: I21, Z10

Differences in time preferences have been recognised as an important determinant of individual investments in skills since the earliest human capital theory of Becker (1964). More recent analyses show that differences in patience are closely related to cross-country differences in educational outcomes and in the resulting differences in income and growth. But, investigation of the role of such fundamental preferences in explaining historically significant *subnational* variations in education and incomes has been stymied by a lack of representative region-specific measures of time preferences. We combine the massive data available from social media—specifically Facebook interests—with machine-learning algorithms to derive new regional measures of patience. We find that patience has a significant role in accounting for within-country differences in student achievement and, by implication, geographical variations in incomes and economic growth.

* Corresponding author: Ludger Woessmann, ifo Center for the Economics of Education, ifo Institute at the University of Munich, Poschingerstrasse 5, 81679 Munich, Germany. Email: woessmann@ifo.de

This paper was received on 2 April 2024 and accepted on 2 June 2025. The Editor was Steffen Huck.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The authors were granted an exemption to publish parts of their data because access to these data is restricted. However, the authors provided the Journal with temporary access to the data, which enabled the Journal to run their codes. The codes for the parts subject to exemption are also available on the Journal repository. The restricted access data and these codes were also checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.15533368.

We gratefully acknowledge comments from the editor Steffen Huck, two anonymous referees, Elliott Ash, Davide Cantoni, David Figlio, Philipp Lergetporer, Ömer Özak, Solomon Polachek, Paola Sapienza, Uwe Sunde and seminar participants at Cornell University, the University of Rochester, the annual meetings of the American Economic Association and the German Economic Association, the CESifo big data workshop, the IZA education workshop, the Berlin-Munich CRC retreat in Ohlstadt, the IWAEE workshop in Catanzaro, the ifo Center for the Economics of Education and the CESifo Area Conference on the Economics of Education. Alina Meiner provided excellent research assistance in preparing the replication package. This work was supported by the Smith Richardson Foundation. The contribution of L.W. is also part of the German Science Foundation project CRC TRR 190.

Many countries have large, long-standing regional differences in student achievement that in turn affect regional income and growth (Hanushek and Woessmann, 2015; Hanushek *et al.*, 2017). Our analysis here focuses on Italy and the United States—two countries with achievement data for a large number of subnational regions showing persistent historical education and income variations.[1] The differences in math achievement between the top- and bottom-performing US states equal roughly two-thirds of the achievement differences between top- and bottom-performing countries in the OECD and are equivalent to over two years of learning. Similar differences are found between the top- and bottom-performing regions in Italy.

Patience, the relative valuation of present versus future payoffs, appears, not only in individual-level, but also in group-level decisions. Students weigh current gratification such as play time with friends against study time that may lead to deferred rewards. Communities trade off present costs against future benefits when deciding how much to invest in school quality, how strongly to motivate children to learn and whether to design institutions to incentivise learning. Testing any hypothesised contribution of patience in affecting regional differences in educational achievement, however, requires representative regional measures of preferences.

A major innovation of this paper is demonstrating how social media data can be used to derive subnational measures of preferences. The fundamental idea, building on recent international analysis of culture by Obradovich *et al.* (2022), is that social media data contain information that characterises people's underlying preferences such as patience at geographically granular levels. This in turn permits investigation of previously unexplained place-based heterogeneity of skills.

For marketing purposes, Facebook has developed an algorithm to classify the 'interests' of over two billion users. We identify the 1,000 Facebook interests with the largest audiences worldwide and then use Facebook's marketing application programming interface (API) to extract data on the prevalence of these interests by country and subsequently by region. After reducing their dimensionality with a principal component analysis, we employ machine-learning techniques to train an international model that predicts the experimentally validated patience measure of the Global Preference Survey of Falk *et al.* (2018).

We first validate the use of our Facebook-derived measure of patience for characterising educational outcomes through an international analysis that mimics existing investigations of preferences and cross-country achievement (Figlio *et al.,* 2019; Hanushek *et al.,* 2022).

We then use the parameters from the international Facebook analysis along with observed regional Facebook interests to construct subnational patience measures across twenty Italian regions and fifty US states. In both countries, the geographic pattern of the Facebook-derived patience measure coincides with long-standing North–South economic disparities. Regional differences in patience account for over three-quarters of the variation in student achievement across Italian regions and for over one-third across US states. In Italy, a one-SD higher regional patience relates to 1.2–1.6-SD higher math achievement, only slightly lower than the cross-country relationship. Possibly related to the substantial internal mobility, the equivalent US estimate is about one-quarter that for Italy, albeit still statistically significant.

Two analyses, while not conclusive, are in line with a causal interpretation of the descriptive baseline association. First, in a cross-country analysis we consider differential outcomes for migrants within each country based on the Facebook preferences of their origin country. By

[1] The large North–South variation in Italy has raised substantial interest in policy and research (e.g., Putnam, 1993; Ichino and Maggi, 2000; Guiso *et al*., 2004) and has been shown to be related to concepts of social preferences such as trust and social capital (Bigoni *et al.*, 2016), beliefs about cooperativeness (Bigoni *et al*., 2019) and civicness (Michaeli *et al*., 2023).

conditioning on fixed effects for the migrants' countries of residence, we shield against unobserved features of students' residence countries and against reverse causation from education or ability to preferences (as suggested, e.g., by Dohmen *et al.,* 2010 or Benjamin *et al.,* 2013). Results are qualitatively similar to the aggregate cross-country results. Compared to the international analysis, the within-country estimation for Italy and the United States is less prone to confounding from unobserved national traits such as languages, constitutions and institutional factors. Second, a complementary instrumental-variable approach goes further in the subnational analysis to address potential confounders. Ancestors of the current populations of the different US states migrated from different countries, giving rise to variation in patience that is not jointly determined with other current state characteristics. Using the weighted patience level of the ancestry countries as an instrument for states' patience confirms a significant effect of patience on student achievement of the same order of magnitude as the baseline model.

Results are stable in robustness analyses that include using reading achievement, differentiating by gender or assessment wave, adding risk-taking and trust as additional cultural-trait controls, and excluding education-related interests when deriving the patience measure from Facebook interests. Moreover, results are consistent for six additional countries where regional achievement data cover fewer grades or regions. The positive association between regional student achievement and Facebook-derived patience holds in the aggregate and is separately significant in five of the additional countries—Brazil, Canada, Germany, Kazakhstan and Mexico—excluding only Spain.

Our main contribution is showing how the well-established theoretical and empirical relationship between patience and human capital accumulation can explain important parts of subnational differences in skills.[2] Prior analyses have largely left persistent regional differences in skills and income unexplained.[3] Our derivation of the regional patience measures further supports the use of social media data in economic analysis of culture and social networks.[4] With our better quantitative measures of more fundamental preference differences, it becomes possible to understand portions of educational and economic outcomes that were previously listed as unexplained heterogeneity.

1. Methods: Deriving Regional Patience Measures from Facebook Interests

With 2.9 billion monthly active users, Facebook is the world's largest social network. Facebook's core business consists of selling advertising space that provides 97.5% of its revenues. Hence, Facebook's business model depends on its ability to keep users engaged on the platform while advertisers promote their products and services. To this purpose, Facebook puts considerable effort into inferring users' interests (Thorson *et al.*, 2021), which is critical to our analysis.

[2] Time preferences are important for economic development (Galor and Özak, 2016; Sunde *et al*., 2022), and previous analyses have shown that they are an important determinant of individual educational outcomes (Sutter *et al*., 2013; Golsteyn *et al*., 2014; Polachek *et al*., 2015; De Paola and Gioia, 2017; Castillo *et al*., 2019; Galor *et al*., 2020; Angerer *et al*., 2023) and of international achievement differences (Figlio *et al*., 2019; Hanushek *et al*., 2022).

[3] Past studies consider proximate causes of regional skill differences such as family background, school spending and institutional settings (e.g., Hanushek and Raymond, 2005; Woessmann, 2010; Dee and Jacob, 2011). Yet, most stop without providing convincing explanations of more fundamental causes of the substantial on-going geographical variations in outcomes (e.g., Hanushek, 2016).

[4] See, e.g., Wilson *et al*. (2012), Bailey *et al*. (2022), Chetty *et al*. (2022), Obradovich *et al*. (2022) and Marty and Duhaut (2024).

1.1. *Extracting Facebook Interests*

Facebook determines users' interests using a variety of sources, both inside the Facebook platform and on external websites (Cabañas *et al.*, 2018; Obradovich *et al.*, 2022). Sources inside the Facebook platform include personal information that users share on Facebook, as well as users' activity on Facebook, such as page likes, group memberships and content with which users engage. Outside the platform, Facebook tracks users' visited websites, installed apps and purchasing behaviour. Facebook uses these data to deliver content and recommendations based on users' interests and to allow advertisers to target users whose interests are relevant for their products and services.

The hundreds of thousands of interests classified by Facebook are organised in nine main categories: business/industry, entertainment, family/relationships, fitness/wellness, food/drink, hobbies/activities, shopping/fashion, sports/outdoors and technology. Interests can be very broad, such as 'Entertainment' or 'Music', or very narrow, such as 'Caribbean Stud Poker', a casino game. Online Appendix Figure A1 depicts the 1,000 Facebook interests with the largest worldwide audience.

Following Obradovich *et al.* (2022), we retrieve Facebook interest data in two steps. First, we obtain a comprehensive list of Facebook interests by querying Facebook's marketing API, the interface that allows advertisers to configure their advertisement campaigns. Per text query, the API returns a collection of closely related Facebook interests with their estimated worldwide audience. We iteratively feed this function with all 25,322 terms of an English dictionary and 2,000 randomly selected titles of Wikipedia articles. This procedure produces 41,513 unique interests from which we select 1,000 with the largest worldwide audience.[5]

Second, for each of the 1,000 interests, we again use Facebook's marketing API to obtain the estimated audience size separately for each country in which Facebook has a presence, as well as for each Italian region and US state. The audience size reflects the entire user population, rather than any specific subgroup such as parents or students, because we attempt to develop proxies for prevailing national or regional preferences. The population of Facebook users may not be representative for the entire population, but our method builds on the Facebook data's ability to predict a patience measure collected in representative samples. For each geographical entity, the process yields a vector of estimated audiences for each of the 1,000 interests. We standardise the estimated audience to mean zero and SD one across the 1,000 interests in each geographical entity.[6]

1.2. *Predicting Country Patience from Facebook Interests*

We first construct a country-level Facebook measure that we use to validate the overall approach of going from Facebook interests to patience. In the next section, we follow a conceptually similar approach to developing the subnational measures that are the heart of our analysis.

We build our Facebook measure of patience through country-level preferences developed in the Global Preference Survey (GPS) from representative population samples in seventy-six countries (Falk *et al.*, 2018). The GPS collected experimentally validated measures of patience (and other

[5] We use 1,000 interests to make the data collection manageable. During data collection between April 2022 and May 2023, the API allowed a maximum of 300 queries per hour. For example, the over 50,000 queries for the US states take more than seven days of uninterrupted queries.

[6] Dividing the Facebook audience counts by population or Facebook users in each geographic entity yields the same qualitative results.

preferences) by combining a qualitative survey item and a hypothetical choice scenario that were chosen based on their capacity to predict incentivised choices in a laboratory setting.

We begin by reducing the dimensionality of Facebook interests by a principal component analysis (PCA) over the sample of all 216 Facebook countries/entities. The first ten principal components (PCs) capture 70% of the total cross-country variance contained in Facebook interests, twenty PCs capture 80% and 48 PCs capture 90% (Online Appendix Figure A2).

We then train a machine-learning model to characterise the relationship between the country-level PCs of Facebook interests and the GPS measure of patience across the seventy-four countries that have both GPS and Facebook data (Online Appendix Table A1).[7] Using a ten-fold cross-validated least absolute shrinkage and selection operator (LASSO) model, the R^2 of the in-sample prediction of patience by the reduced-dimensionality Facebook interests is quite stable between 0.65 and 0.70, independent of whether 10, 20 or 50 PCs are used (Online Appendix Figure A3).

We use the parameters from the machine-learning model to predict patience for all eighty countries with both Programme for International Student Assessment (PISA) and Facebook data (Online Appendix Table A1 and Online Appendix Figure A4). Given the limited size of the training sample used, we rely on the parsimonious specification with ten PCs for the out-of-sample predictions to avoid overfitting.

The beauty of this approach is its ability to capture latent components of users' underlying degree of patience. The procedure does not lend itself to identifying specific interests as main predictors of the GPS patience measure, because any single Facebook interest can load positively or negatively on different PCs that enter the patience prediction.[8] However, if we correlate our Facebook-derived patience measure with the individual Facebook interests, terms such as dogs, outdoor recreation, adventure, holiday, painting, wildlife and garden consistently show up among the strongest positive correlations, whereas free software, WhatsApp, Facebook Messenger, massively multiplayer online role-playing games and the like show up as negative. Among leisure activities (which are broadly represented by Facebook), the former patience-related interests may indicate valuation of longer-term future rewards, whereas the latter interests may proxy for instant gratification. Still, not all interests with strong correlations are easily interpretable, which aligns with the idea of picking up latent factors that are not well measured by individual Facebook interests. The Facebook-derived measure of patience can obviously be associated with regional differences in all sorts of other measures such as money, intelligence or motivation—but, by construction, only to the extent that these are associated with patience as measured in the GPS.

We perform the same training and prediction models for risk-taking, another preference with relevance for intertemporal decision-making contained in the GPS and previously found to enter international student achievement (Hanushek *et al.*, 2022). The R^2 of the in-sample prediction for risk-taking is lower than for patience (Online Appendix Figure A3), indicating that risk-taking is harder to predict from Facebook interests.[9]

[7] The GPS measure is standardised to have mean zero and SD one across individuals in the GPS countries, so that estimates in our subsequent analyses can be interpreted in terms of SDs.

[8] Obradovich *et al.* (2022) showed that direct prediction of cultural measures from individual Facebook interests generally provides worse predictions than using PCs.

[9] For other GPS preferences, the R^2 of the in-sample prediction is even lower at about 0.35 for trust and at most 0.2 for altruism, positive reciprocity and negative reciprocity (see also Obradovich *et al.*, 2022).

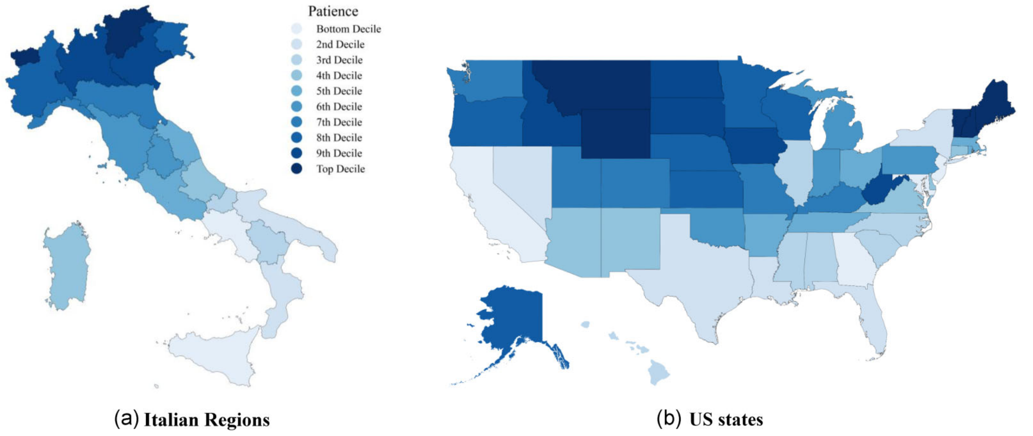

Fig. 1. *Measure of Patience Derived from Facebook Interests for Italian Regions and US States.*
Notes: The figure shows maps of the Facebook-derived measure of patience obtained with four PCs for Italian regions (panel (a)) and US states (panel (b)). Each colour corresponds to a decile of the distribution of patience within each country. Darker colours denote higher levels of patience.

1.3. *Predicting Regional Patience from Facebook Interests*

Our primary analysis builds on the development of subnational variations in patience. In parallel to the cross-country analysis, we reduce the dimensionality of Facebook interests using a PCA fit across regions *within* a given country. Fitting the PCA at the subnational level ensures that the PCs capture country-specific dimensions of Facebook interests. For both Italian regions and US states, the first four PCs capture over 70% of the regional variance in Facebook interests (Online Appendix Figures A5 and A6); 90% of variance is captured by ten PCs in Italy and fifteen PCs in the United States.

Separately for Italy and the United States, we train a ten-fold cross-validated LASSO model to learn the relationship between the GPS measure of patience across countries and the Facebook interests aggregated by the parameters of within-country PCs.[10] A small number of PCs capture a considerable portion of the variation in the GPS patience measure. With ten PCs, the R^2 of the in-sample prediction reaches 0.5 for PCs fitted across Italian regions and over 0.6 for PCs fitted across US states (Online Appendix Figures A7 and A8).

We use the parameter estimates from the two internationally trained models to construct patience measures from the subnational Facebook interests of Italian regions and US states. Figure 1 shows maps of the regional variation of patience in the two countries. In Italy, the regions with the lowest patience measure are Sicily and Campania in the South. The region with the highest level of patience is Trentino-Alto-Adige in the North-East.[11] In the United States, the

[10] We focus on all seventy-four countries with GPS and Facebook data to maximise the sample size for the machine-learning algorithm. However, results are quite similar if we instead only use the twenty-five OECD countries in the training model: the correlation among the two versions is 0.981 for Italian regions and 0.974 for US states, and qualitative results do not change for our main regression analysis.

[11] Interestingly, parts of Trentino-Alto-Adige belonged to Austria and the former Austro-Hungarian empire for long periods of time, and large parts of the population speak German as their first language. The high level of predicted patience in Trentino-Alto-Adige is consistent with the fact that neighbouring Austria has much greater patience than Italy according to the country-level GPS measures, adding qualitative support for the Facebook-derived measure.

states that exhibit the highest level of patience are Vermont and Maine in the North-East. Both countries tend to show a North–South gradient in the Facebook-derived measure of patience.

A similar prediction model for risk-taking performs substantially worse. The R^2 of the in-sample prediction is well below 0.2 for all models with up to ten PCs in both Italy and the United States (see Online Appendix Figures A7, A8 and A9). We include risk-taking as a control variable throughout, given its interrelatedness with patience, but its poor measurement at the regional level means that the estimates for patience are likely lower bounds.[12]

One way of directly validating our Facebook-derived patience measure is to compare it to regional GPS data. The GPS data contain regional identifiers that allow construction of non-representative regional GPS measures of patience (Sunde *et al.*, 2022). These are very noisy due to the small regional GPS sample sizes, averaging fifty individuals per Italian region and twenty per US state. Nonetheless, they are positively correlated with our measure (weighted by the number of GPS observations per region): 0.49 (significant at the 5% level) across Italian regions and 0.23 (significant at the 10% level) across US states.

2. Cross-Country Results

The cross-country patience measures based on Facebook data allow us to validate whether Facebook interests can provide reliable estimates of geographically varying degrees of patience. They also provide some insights into the causal structure of the cross-country pattern of patience and achievement.

2.1. *Cross-Country Validation of Using Facebook Interests to Measure Patience*

To validate our Facebook-derived measures of patience and risk-taking, we estimate their relationship with standardised student math achievement across countries for all seven available waves of the PISA from 2000–18 (see Hanushek *et al.,* 2022):

$$T_{ict} = \beta_1 Patience_c + \beta_2 Risk_c + \alpha_1 \boldsymbol{B}_{ict} + \boldsymbol{\mu}_t + \varepsilon_{ict}.$$

Here T_{ict} is the standardised PISA test score of student i in country c in year t; $\boldsymbol{B}$ is a vector of controls (student gender, age and migration status); the $\boldsymbol{\mu}_t$ are fixed effects for test waves to account for time trends and idiosyncrasies of individual tests; ε_{ict} is an error term; β_1 and β_2 characterise the relationship of patience and risk-taking at the country level with student achievement. OLS regressions are weighted by students' sampling probability, giving equal weight to each country. SEs are clustered at the country level.

The comparison model from Hanushek *et al.* (2022) uses the original GPS measures and shows a strong positive relationship between patience and student achievement and a strong negative relationship with risk-taking (column (1) of Table 1, panel A). Substituting our Facebook-derived preference measures (column (2)) produces slightly larger preference parameters and corroborates the validity of the Facebook-derived measures.[13]

Out-of-sample predictions allow us to extend the analysis of the Facebook-derived measures of patience and risk-taking from forty-eight to eighty countries—all countries that participated

[12] In the cross-country analysis, patience and risk-taking are positively associated, and risk-taking is negatively associated with achievement, leading to negative bias (Hanushek *et al.*, 2022).

[13] The estimates rely on ten PCs of Facebook interests, but results are very similar when using additional (20–50) PCs (Online Appendix Table A2).

Table 1. *Patience, Risk-taking and Student Achievement: Cross-Country Validation Exercise.*

	GPS measure (1)	Facebook measure (ten PCs) Original sample (2)	Extended sample (3)	Non-GPS sample (4)
Panel A. Cross-country analysis				
Patience	1.225***	1.684***	1.722***	1.771***
	(0.132)	(0.135)	(0.119)	(0.210)
Risk-taking	−1.229***	−1.359***	−1.537***	−1.660***
	(0.188)	(0.310)	(0.254)	(0.388)
Control variables	Yes	Yes	Yes	Yes
Observations	1,954,840	1,954,840	2,660,408	705,568
Residence countries	48	48	80	32
R^2	0.200	0.210	0.220	0.241
Panel B. Migrant analysis				
Patience	0.957***	0.805***	0.902***	1.766***
	(0.115)	(0.182)	(0.205)	(0.481)
Risk-taking	−0.315**	−0.677**	−1.221***	−3.531***
	(0.124)	(0.278)	(0.350)	(0.549)
Control variables	Yes	Yes	Yes	Yes
Residence country by wave FEs	Yes	Yes	Yes	Yes
Observations	78,403	78,403	90,983	12,580
Countries of origin	56	56	93	37
Residence countries	46	46	50	34
R^2	0.280	0.272	0.298	0.310

Notes: The dependent variable is the PISA math test score. Least-squares regressions. Panel A: all PISA waves 2000–18; weighted by students' sampling probability. Panel B: waves 2003–18; students with both parents not born in the country where the student attends school; including 180 fixed effects for each residence country by wave cell. Control variables: student gender, age and migration status, imputation dummies and wave fixed effects in panel A; student gender, age, dummy for OECD country of origin and imputation dummies in panel B. Robust SEs adjusted for clustering at the country level (migrant analysis: country of origin) are reported in parentheses. Significance level: *** 1%, ** 5%.
Source: PISA international student achievement test, 2000–18; Falk *et al.* (2018); own elaboration of Facebook data.

in PISA and have Facebook data—encompassing over 2.6 million student observations. Results generalise very well to the extended sample, with increased precision and without significantly different estimates (column (3)). Even in the thirty-two countries that were not part of the original GPS analysis, results are qualitatively the same and statistically highly significant (column (4)).

2.2. *Exploration into Causality: Migrant Analysis*

We also validate our measures with an analysis of migrant students that aims to get closer to a causal interpretation of the cross-country relationship between patience and student achievement. The most significant threats to identification of the preference effects are that the relationships are driven by reverse causation or by other factors of the country of schooling. We restrict the PISA analysis to students with a migrant background and assign them the values of patience and risk-taking of their home countries—an approach that avoids bias from reverse causation (see

Figlio *et al.*, 2019; Hanushek *et al.,* 2022).[14] By observing migrant students from different origin countries, but schooled in the same residence country, we can include fixed effects for residence countries that control for other resident country factors that could bias the baseline cross-country analysis.

Migrant results in panel B of Table 1 show that the prior positive patience relationship and the negative risk-taking relationship again replicate well.[15] When restricting the sample to non-GPS countries (column (4)), estimates become quite imprecise (and larger), indicating limited power of the migrant analysis in the smaller sample.[16]

The cross-country migrant analysis further validates the informational content of the Facebook-derived measures and suggests a causal interpretation.

3. Subnational Results: Italian Regions and US States

Our subnational empirical model follows the cross-country model introduced in the previous section. We think of patience as a deep determinant of student achievement, leading us to employ very parsimonious specifications of achievement differences. From this perspective, proximate inputs often included in education production functions such as parental education or school resources would be bad controls as they are endogenous to a region's patience.

Regionally representative data on student achievement come from student-level INVALSI test data for Italy and from state-level NAEP data for the United States and refer to the last waves before the COVID-19 pandemic (see Online Appendix A for details).[17] We initially focus on math achievement in eighth grade, the oldest cohort available in both countries and closest in age to PISA. To allow for interpretation in terms of SDs, we divide test scores by each country's student-level SD.

3.1. *Achievement across Italian Regions*

The long-standing North–South divide among Italian regions invites investigation of fundamental driving forces. Because the schooling system is regulated mostly at the country level, test score variations across regions are unlikely to be driven by the institutional structure of schools.

Regional differences in patience are strongly and significantly associated with student achievement. Student-level results in panel A of Table 2 show that a one-SD increase in regional patience is associated with an increase in math test scores of 1.40–1.61 SDs, which is close to the cross-country estimates reported in Table 1. Measuring patience with 4, 7 or 10 PCs of Facebook interests has little impact.

Differences in patience account for over three-quarters of the aggregate regional variation in student achievement (panel B of Table 2; see Online Appendix Figure A10).[18] Point estimates of the region-level analysis are very similar, albeit slightly smaller than in the student-level analysis.

[14] Prior analysis suggests that migrant children obtain some of the preferences of their country of origin through family linkages (e.g., Guiso *et al*., 2006; Bisin and Verdier, 2011; Alesina and Giuliano, 2014).

[15] With the Facebook data, we expand the countries of origin considered in the migrant analysis from fifty-six to ninety-three (see Online Appendix Table A3). The destination countries increase only from forty-six to fifty because some PISA countries do not report students' and parents' countries of birth required to determine migrants' country-of-origin preferences.

[16] Results of the migrant analysis are again stable for patience when using 20–50 PCs (Online Appendix Table A4).

[17] INVALSI stands for Istituto Nazionale per la Valutazione del Sistema Dell'Istruzione and NAEP for National Assessment of Educational Progress.

[18] The R^2 is virtually unchanged when wave fixed effects and risk-taking are excluded from the model.

Table 2. *Patience and Student Achievement: Subnational Analysis for Italy and the United States.*

	Four PCs (1)	Seven PCs (2)	Ten PCs (3)
Panel A. Italy (individual level)			
Patience	1.614***	1.398***	1.495***
	(0.173)	(0.107)	(0.107)
Control variables	Yes	Yes	Yes
Wave FEs	Yes	Yes	Yes
Observations	59,034	59,034	59,034
Regions	20	20	20
R^2	0.095	0.100	0.101
Panel B. Italy (regional level)			
Patience	1.370***	1.201***	1.284***
	(0.169)	(0.086)	(0.089)
Wave FEs	Yes	Yes	Yes
Observations	40	40	40
Regions	20	20	20
R^2	0.771	0.843	0.852
Panel C. United States (state level)			
Patience	0.293***	0.172*	0.291**
	(0.089)	(0.096)	(0.131)
Wave FEs	Yes	Yes	Yes
Observations	153	153	153
Regions	51	51	51
R^2	0.360	0.348	0.364

Notes: In panels A and B the dependent variable is the INVALSI eighth-grade math test score in waves 2018 and 2019; in panel C the dependent variable is the NAEP eighth-grade math test score in all NAEP waves 2015–19. Least-squares regressions with wave fixed effects. Unit of observation: student in panel A; region-wave combination in panel B; state-wave combination in panel C. Patience measured at the regional/state level throughout. Columns (1)–(3) use the patience measure computed with 4, 7 and 10 PCs, respectively. Regressions control for risk-taking computed with the equivalent number of PCs. Additional control variables (panel A): student gender, age and migration status; imputation dummies. The Italian region of Trentino-Alto-Adige is represented by the two municipalities of Bolzano and Trento. Robust SEs adjusted for clustering at the regional (state) level are reported in parentheses. Significance level: *** 1%, ** 5%, * 10%.
Source: INVALSI mathematics achievement test, 2017–19; NAEP mathematics achievement test, 2015–19; own elaboration of Facebook data.

The results are also consistent with a cumulative impact process: coefficient estimates at the student level increase continuously from an insignificant 0.31 SDs in second grade to a highly significant 1.88 SDs in tenth grade (Online Appendix Table A5). Region-level estimates are again quite similar. While other mechanisms may additionally be at play, the role of patience suggestively adds up across grades (see also Figlio *et al.,* 2019).

3.2. *Achievement across US States*

Patience is also significantly associated with higher student achievement at the US state level and accounts for slightly more than one-third of the state NAEP variation (panel C of Table 2).[19] A

[19] The sample for Table 2 includes Washington, DC, but results are similar when it is excluded (not shown).

one-SD increase in the Facebook-derived measure of patience is associated with an increase of 0.17–0.29 SDs in test scores.

The cross-state impact of patience is only about one-quarter of that estimated for Italian regions. One possible explanation is that the population in the United States is substantially more mobile and mixed. In 2019, 42% of the US population lived in a state different from their state of birth, suggesting that state cultural differences likely lessen over time. But, the lower explanatory power of patience for US states also suggests a wider range of other factors affecting achievement, such as compensatory state education policies.

The impact of patience in fourth grade is only about half the size as in eighth grade (panel C of Online Appendix Table A5). This is again consistent with a cumulating impact of patience, but a variety of other reasons could also be at play.

3.3. *Robustness Analysis*

Results prove stable for Italy and the United States in a series of robustness analyses. We summarise the various tests here and provide details in Online Appendix B.

Both in Italy and the United States, there are no significant overall gender differences. The impact of patience on reading is very similar to that for math, albeit with slightly smaller point estimates. All prior results hold, not only for the different pre-COVID assessment waves separately, but also for the post-COVID period. Additionally, when we consider other preferences from GPS, only trust is reliably estimated by Facebook interests, but including trust in the models leaves the patience results unchanged.

A potential methodological concern is that the 1,000 Facebook interests used to construct the patience measures contain interests that are directly related to education, possibly introducing endogeneity in the Facebook-derived patience measure to educational outcomes. We exclude seventeen Facebook interests directly related to education from the construction of our patience measure. The alternative subnational patience measures are correlated above 0.998 both for Italian regions and US states, and the estimated impact of measured patience on achievement hardly changes.

Using the individual-level data in Italy, we can confirm that the estimated impact of patience is larger for native students than for migrant students. Results also hold when excluding Trentino-Alto-Adige whose sample is not representative for the entire region and whose German-language population might limit comparability. Italy also participated with a regionally representative sample in the international PISA test in 2012, and we find that results hold equally well with this alternative achievement test. Finally, results are also robust in an analysis of unobservable selection and coefficient stability proposed by Oster (2019).

3.4. *Exploration into Causality: Patience in Ancestry Countries of US State Populations*

The previous analysis of patience and subnational differences in human capital has been descriptive. Concerns about causal interpretation could emerge if patience of the regional population is affected by the region's human capital (reverse causation) or if regional patience is jointly determined with other regional factors that are important for educational achievement (omitted variables). For the cross-country analysis, our investigation of migrants in Section 2.2 addresses these main concerns of endogeneity. The within-country analysis is less prone to bias that may arise from national factors such as languages, laws and institutional settings, and our robustness

Table 3. *Patience and Student Achievement across US States: Using Ancestry Information as Instruments.*

		First stage			Second stage		
	Reduced form (1)	Four PCs (2)	Seven PCs (3)	Ten PCs (4)	Four PCs (5)	Seven PCs (6)	Ten PCs (7)
Ancestry patience	0.945*** (0.349)	3.617*** (0.306)	3.605*** (0.283)	2.143*** (0.229)			
Patience					0.261*** (0.096)	0.221** (0.090)	0.372** (0.156)
Wave FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen–Paap Wald *F*-statistic		139.321	161.733	87.387			
Observations	153	153	153	153	153	153	153
Regions	51	51	51	51	51	51	51
(Centered) R^2	0.518				0.503	0.502	0.469

Notes: In columns (1), (5)–(7) the dependent variable is the NAEP eighth-grade math test score in all NAEP waves 2015–9; in columns (2)–(4) the dependent variable is patience. Unit of observation: state-wave combination. Patience measured at the state level throughout. Column headers indicate the number of PCs used to compute the patience measure. Regressions control for risk-taking computed with the equivalent number of PCs and for the share of missing ancestry information. Robust SEs adjusted for clustering at the state level are reported in parentheses. Significance level: *** 1%, ** 5%.
Source: NAEP mathematics achievement test, 2015–19; ACS; own elaboration of Facebook data.

analysis showed that the subnational results are robust to conditioning on other cultural traits. But, other threats to identification remain.

To explore further the empirical relevance of any remaining endogeneity bias, we look into the ancestry of US residents. The ancestors of current US state populations migrated from different countries, suggesting that time preferences of current residents have varying historical roots (Galor and Özak, 2016; Becker *et al.*, 2020) that may partly be handed down from generation to generation within families (Bisin and Verdier, 2011; Alesina and Giuliano, 2014). This suggests that the levels of patience of the countries from which the ancestors migrated can potentially proxy the patience of each state's current population while neither being affected by nor jointly determined with other current characteristics of US states. We therefore employ an identification strategy that uses patience in the ancestry countries of US states' populations as an instrumental variable (IV) for our Facebook-derived measure of patience in the US states.

Demographic data from the American Community Survey (ACS) include ancestry or ethnic origin information for the 2022 population of each US state (IPUMS, Ruggles *et al.*, 2023). The ancestry-based measure of patience for each US state is developed by assigning respondents the patience level of the respective ancestry country from the GPS. We calculate shares for each ancestry country by US state and weight the ancestry patience levels accordingly.[20]

The reduced-form estimate shows that states whose population migrated from ancestry countries with higher patience levels on average have significantly higher test scores (Table 3, column (1)). In the first stage of the IV model, the ancestry-based patience instrument (measured directly from the GPS data) strongly predicts patience as measured by our Facebook-derived measure,

[20] For about a third of ACS respondents, the ancestry information is either missing or cannot be assigned to a specific country because it refers to an ethnic group (e.g., Kurdish) or a larger region (e.g., Eastern Europe). We assign these respondents the average US patience level from the GPS and control for the share of missing ancestry information in the regressions. Results are very similar when coding missing ancestry information as missing, which is equivalent to assigning the average of the observed ancestries of the respective state (not shown).

Table 4. *Patience and Student Achievement: Subnational Analysis in Eight Countries.*

	Eight countries pooled (1)	Brazil (2)	Canada (3)	Germany (4)	Italy (5)	Kazakhstan (6)	Mexico (7)	Spain (8)	United States (9)
Patience (three PCs)	0.337***	1.556***	0.383**	2.093***	1.527***	0.459**	0.666***	0.060	0.218*
	(0.085)	(0.206)	(0.155)	(0.624)	(0.205)	(0.194)	(0.160)	(0.108)	(0.113)
Grade/age	–	Age 15	Age 15	Grade 9	Grade 8	Age 15	Age 15	Age 15	Grade 8
Wave FEs	Yes	No	Yes	Yes	Yes	No	No	Yes	Yes
Country FEs	Yes	No	No	No	No	No	No	No	No
Observations	381	27	30	32	40	16	32	51	153
Regions	189	27	10	16	20	16	32	17	51
R^2	0.288	0.719	0.726	0.577	0.747	0.178	0.517	0.362	0.300

Notes: The dependent variable is math test scores. Least-squares regressions with wave fixed effects. Unit of observation: region-wave combination. Test and wave information: PISA 2012 for Brazil and Mexico; PISA 2012, 2015 and 2018 for Canada and Spain; IQB 2012 and 2018 for Germany; INVALSI 2018 and 2019 for Italy; PISA 2018 for Kazakhstan; NAEP 2015, 2017 and 2019 for the United States. Regressions control for risk-taking computed with three PCs. Robust SEs adjusted for clustering at the state level are reported in parentheses. Significance level: *** 1%, ** 5%, * 10%.
Source: PISA, IQB, INVALSI and NAEP mathematics achievement tests; own elaboration of Facebook data.

independent of the number of PCs of Facebook interests used (columns (2)–(4)). This first-stage association provides further credence to our derivation of patience measures from Facebook interests.

The second stage of the IV model provides statistically significant estimates of the impact of patience that are of the same order of magnitude as the OLS estimates. For the different numbers of PCs, the IV point estimates range from 0.22 to 0.37 (columns (5)–(7)), quite similar to our baseline model (0.17 to 0.29).

While fully eliminating all concerns of causal interpretation is difficult in settings without experimental manipulation and while the IV model cannot rule out all potential biases, its results provide no indication that the baseline estimates are biased upwards by endogeneity from reverse causation or joint determination with other current state characteristics.

3.5. *Regional Analysis in Additional Countries*

While we have focused on Italy and the United States as two countries with interesting regional variation and consistent test data at different grades for a substantial number of regions, we can assess the stability of our results by extending the analysis to six additional countries with publicly available subnational test data. We leverage regional indicators in the PISA data since 2012 for all countries with at least ten regions: Canada and Spain in 2012, 2015 and 2018, Brazil and Mexico in 2012 and Kazakhstan in 2018. Also, the Institut zur Qualitätsentwicklung im Bildungswesen (IQB) provides regionally representative math achievement data for German ninth-grade students in 2012 and 2018. For each country, we implement the method described in Section 1.3 to obtain regional measures of patience from Facebook interests, consistently using only three PCs because of the small number of regions in some countries.[21]

The consistency of results across these additional countries supports the methodology for investigating achievement differences within countries. In the pooled model of 190 regions in eight countries, the highly significant patience coefficient suggests that a one-SD increase in patience is associated with a 0.34-SD increase in math scores (Table 4, column (1)). Country-specific results are more tentative due to the limited regional information in several countries,

[21] Pooled results are similar using more PCs, but country-specific results are not stable at higher numbers of PCs in Brazil, Canada and Germany.

but separate regressions show a positive regional association between patience and achievement that is statistically significant in each country except Spain (Table 4, columns (2)–(9) and Online Appendix Figure A11). The magnitude of coefficient estimates varies considerably across countries, suggesting that the strength of the relationship might depend on country-specific features, but there are too few country observations to analyse these differences systematically.

4. Conclusions

Time preferences, while clearly important to individual investment decisions, have an even broader impact on education decisions. Aggregate preferences, which are a component of cultural identities, also affect political perspectives and community decisions about educational institutions such as the definition and importance of school quality.

This analysis investigates the importance of patience in determining historically significant, but largely unexplained regional differences in student skills. These skill differences in turn have lasting consequences for incomes and for regional growth patterns.

We use the extensive compilations of social media information by Facebook to estimate preference differences for subnational regions within Italy and the United States. The measures of patience constructed from Facebook interests are validated by international comparisons where direct measures of time preferences are available.

Differences in patience across regions in Italy and across states in the United States provide a powerful explanation of human capital outcomes. This new perspective on student performance helps to explain why, for example, North–South differences in student outcomes in both countries have been very stable over time even in the face of national efforts to equalise performance. Causal identification of the preference-achievement relationship across subnational regions is particularly challenging, but an instrumental-variable approach that exploits historical ancestry variation across US states supports a causal interpretation.

Our findings imply that similar educational inputs can lead to substantially different outcomes due to differences in patience. When addressing within-country differences in student achievement, policymakers might look beyond such proximate factors as school spending or even family educational background to take possible differences in patience into account. Institutional features of schooling, such as reliance on parental choice or test-based accountability, appear less tied to aggregate preferences (Hanushek *et al.*, 2022). Thus, institutional reforms of school systems appear to be a viable policy mechanism for improvement that does not necessarily depend on changing preferences (Woessmann, 2016). Moreover, while cultural traits are considered hard to change (e.g., Guiso *et al.*, 2006; Bisin and Verdier, 2011), evidence shows that traits such as patience are malleable, especially at a young age, and can be improved through specific interventions (e.g., Bird, 2001; Alan and Ertac, 2018; Jung *et al.*, 2021). Hence, policies aimed at increasing patience may also be an avenue for addressing educational investments and regional deficits in student outcomes.

Hoover Institution, Stanford University, NBER, USA, CESifo & IZA, Germany
DIW Berlin, Germany
ifo Institute & University of Munich, Germany
University of Munich, ifo Institute, CESifo, IZA, Germany, Hoover Institution & Stanford University, USA

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Alan, Sule, and Seda Ertac. (2018). 'Fostering patience in the classroom: Results from randomized educational intervention', *Journal of Political Economy*, vol. 126(5), pp. 1865–911.

Alesina, Alberto, and Paola Giuliano. (2014). 'Family ties', in (P. Aghion and S.N. Durlauf, eds.), *Handbook of Economic Growth*, vol. 2, pp. 177–215, Amsterdam: North Holland.

Angerer, Silvia, Jana Bolvashenkova, Daniela Glätzle-Rützler, Philipp Lergetporer, and Matthias Sutter. (2023). 'Children's patience and school-track choices several years later: Linking experimental and field data', *Journal of Public Economics*, vol. 220, 104837.

Bailey, Michael, Drew M. Johnston, Martin Koenen, Theresa Kuchler, Dominic Russel, and Johannes Stroebel. (2022). 'The social integration of international migrants: Evidence from the networks of Syrians in Germany', Working Paper 29925, National Bureau of Economic Research.

Becker, Anke, Benjamin Enke, and Armin Falk. (2020). 'Ancient origins of the global variation in economic preferences', *AEA Papers and Proceedings*, vol. 110, pp. 319–23.

Becker, Gary S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, Chicago: University of Chicago Press.

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. (2013). 'Who is "behavioral"? Cognitive ability and anomalous preferences', *Journal of the European Economic Association*, vol. 11(6), pp. 1231–55.

Bigoni, Maria, Stefania Bortolotti, Marco Casari, and Diego Gambetta. (2019). 'At the root of the North-South cooperation gap in Italy: Preferences or beliefs?', Economic Journal, vol. 129(619), pp. 1139–52.

Bigoni, Maria, Stefania Bortolotti, Marco Casari, Diego Gambetta, and Francesca Pancotto. (2016). 'Amoral familism, social capital, or trust? The behavioural foundations of the Italian North-South divide', Economic Journal, vol. 126(594), pp. 1318–41.

Bird, Edward J. (2001). 'Does the welfare state induce risk-taking?', *Journal of Public Economics*, vol. 80(3), pp. 357–83.

Bisin, Alberto, and Thierry Verdier. (2011). 'The economics of cultural transmission and socialization', in (J. Benhabib, A. Bisin and M.O. Jackson, eds.), *Handbook of Social Economics*, pp. 339–416. Amsterdam: North-Holland.

Cabañas, José G., Ángel Cuevas, and Rubén Cuevas. (2018). 'Facebook use of sensitive data for advertising in Europe', Preprint, Updated February 14, 2018. https://arxiv.org/abs/1802.05030.

Castillo, Marco, Jeffrey L. Jordan, and Ragan Petrie. (2019). 'Discount rates of children and high school graduation', Economic Journal, vol. 129(619), pp. 1153–81.

Chetty, Raj, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, Drew Johnston, Martin Koenen, Eduardo Laguna-Muggenburg, Florian Mudekereza, Tom Rutter, Nicolaj Thor, Wilbur Townsend, Ruby Zhang, Mike Bailey, Pablo Barberá, Monica Bhole, and Nils Wernerfelt. (2022). 'Social capital I: Measurement and associations with economic mobility', *Nature*, vol. 608(7921), pp. 108–21.

De Paola, Maria, and Francesca Gioia. (2017). 'Impatience and academic performance. Less effort and less ambitious goals', *Journal of Policy Modeling*, vol. 39(3), pp. 443–60.

Dee, Thomas S., and Brian A. Jacob. (2011). 'The impact of no child left behind on student achievement', *Journal of Policy Analysis and Management*, vol. 30(3), pp. 418–46.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. (2010). 'Are risk aversion and impatience related to cognitive ability?', *American Economic Review*, vol. 100(3), pp. 1238–60.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. (2018). 'Global evidence on economic preferences', *The Quarterly Journal of Economics*, vol. 133(4), pp. 1645–92.

Figlio, David, Paola Giuliano, Umut Özek, and Paola Sapienza. (2019). 'Long-term orientation and educational performance', *American Economic Journal: Economic Policy*, vol. 11(4), pp. 272–309.

Galor, Oded, and Ömer Özak. (2016). 'The agricultural origins of time preference', *American Economic Review*, vol. 106(10), pp. 3064–103.

Galor, Oded, Ömer Özak, and Assaf Sarid. (2020). 'Linguistic traits and human capital formation', *AEA Papers and Proceedings*, vol. 110, pp. 309–13.

Golsteyn, Bart H.H., Hans Grönqvist, and Lena Lindahl. (2014). 'Adolescent time preferences predict lifetime outcomes', Economic Journal, vol. 124(580), pp. F739–61.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. (2004). 'The role of social capital in financial development', *American Economic Review*, vol. 94(3), pp. 526–56.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. (2006). 'Does culture affect economic outcomes?', *Journal of Economic Perspectives*, vol. 20(2), pp. 23–48.

Hanushek, Eric A. (2016). 'What matters for achievement: Updating Coleman on the influence of families and schools', *Education Next*, vol. 16(2), pp. 22–30.

Hanushek, Eric A., Lavinia Kinne, Philipp Lergetporer, and Ludger Woessmann. (2022). 'Patience, risk-taking, and human capital investment across countries', Economic Journal, vol. 132(646), pp. 2290–307.

Hanushek, Eric A., and Margaret E. Raymond. (2005). 'Does school accountability lead to improved student performance?', *Journal of Policy Analysis and Management*, vol. 24(2), pp. 297–327.

Hanushek, Eric A., Jens Ruhose, and Ludger Woessmann. (2017). 'Knowledge capital and aggregate income differences: Development accounting for U.S. states', *American Economic Journal: Macroeconomics*, vol. 9(4), pp. 184–224.

Hanushek, Eric A., and Ludger Woessmann. (2015). *The Knowledge Capital of Nations: Education and the Economics of Growth*, Cambridge, MA: MIT Press.

Ichino, Andrea, and Giovanni Maggi. (2000). 'Work environment and individual background: Explaining regional shirking differentials in a large Italian firm', *Quarterly Journal of Economics*, vol. 115(3), pp. 1057–90.

Jung, Dawoon, Tushar Bharati, and Seungwoo Chin. (2021). 'Does education affect time preference? Evidence from Indonesia', *Economic Development and Cultural Change*, vol. 69(4), pp. 1451–99.

Marty, Robert, and Alice Duhaut. (2024). 'Global poverty estimation using private and public sector big data sources', *Scientific Reports*, vol. 14(1), 3160.

Michaeli, Moti, Marco Casari, Andrea Ichino, Maria De Paola, Ginevra Marandola, and Vincenzo Scoppa. (2023). 'Civicness drain', Economic Journal, vol. 133(649), pp. 323–54.

Obradovich, Nick, Ömer Özak, Ignacio Martín, Ignacio Ortuño-Ortín, Edmond Awad, Manuel Cebrián, Rubén Cuevas, Klaus Desmet, Iyad Rahwan, and Ángel Cuevas. (2022). 'Expanding the measurement of culture with a sample of two billion humans', *Journal of the Royal Society Interface*, vol. 19(190), 20220085.

Oster, Emily. (2019). 'Unobservable selection and coefficient stability: Theory and evidence', *Journal of Business & Economic Statistics*, vol. 37(2), pp. 187–204.

Polachek, Solomon W., Tirthatanmoy Das, and Rewat Thamma-Apiroam. (2015). 'Micro- and macroeconomic implications of heterogeneity in the production of human capital', *Journal of Political Economy*, vol. 123(6), pp. 1410–55.

Putnam, Robert D. (1993). *Making Democracy Work: Civic Traditions in Modern Italy*, Princeton, NJ: Princeton University Press.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rogers, and Megan Schouweiler. (2023). 'IPUMS USA: Version 14.0 [dataset]', Minneapolis, MN: Integrated Public Use Microdata Series (IPUMS), https://doi.org/10.18128/D010.V14.0.

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim. (2022). 'Patience and comparative development', *The Review of Economic Studies*, vol. 89(5), pp. 2806–40.

Sutter, Matthias, Martin G. Kocher, Daniela Glätzle-Rützler, and Stefan T. Trautmann. (2013). 'Impatience and uncertainty: Experimental decisions predict adolescents' field behavior', *American Economic Review*, vol. 103(1), pp. 510–31.

Thorson, Kjerstin, Kelley Cotter, Mel Medeiros, and Chankyung Pak. (2021). 'Algorithmic inference, political interest, and exposure to news and politics on Facebook', *Information, Communication & Society*, vol. 24(2), pp. 183–200.

Wilson, Robert E., Samuel D. Gosling, and Lindsay T. Graham. (2012). 'A review of Facebook research in the social sciences', *Perspectives on Psychological Science*, vol. 7(3), pp. 203–20.

Woessmann, Ludger. (2010). 'Institutional determinants of school efficiency and equity: German states as a microcosm for OECD countries', *Journal of Economics and Statistics*, vol. 230(2), pp. 234–70.

Woessmann, Ludger. (2016). 'The importance of school systems: Evidence from international differences in student achievement', *Journal of Economic Perspectives*, vol. 30(2), pp. 3–32.

The Economic Journal, **136** (*January*), 335–350 https://doi.org/10.1093/ej/ueaf045
Advance Access Publication Date: 18 June 2025