

Assessments in Education*

Hans Henrik Sievertsen (VIVE, University of Bristol, and IZA)

January 2023

Summary

Assessments like standardized tests and teacher evaluations are central elements of educational systems. Assessments affect the behaviour of students, teachers, parents, schools, and policymakers through at least two channels: The information channel and the incentive channel. Students use the information to adjust study effort and to guide their course selection. Schools and teachers use information from assessments to evaluate teaching quality and the effectiveness of the applied methods. Educational programs use information from assessment results to sort students in educational programs and employers use the results as signals of productivity in their hiring decisions. Finally, policymakers use assessments in accountability systems to reward or penalize schools, and parents use information from assessment results to select schools. The incentive channel is a natural consequence of the information channel: Students are incentivized to work hard and do well in assessments to get access to educational programs and jobs. Teachers and schools are incentivized to do well to receive rewards or avoid punishments in accountability systems. The information channel is important for ensuring the most efficient human capital investments: students learn about the returns and costs of effort investments and about their abilities and comparative advantages. Teachers and schools learn about the most effective teaching methods. However, because of the strong incentives linked to assessments, both students and teachers might focus on optimizing assessment results at the cost of learning. Students might for example select tracks that maximize

*Contact: University of Bristol, School of Economics. Address: Priory Road Complex, Priory Road, Bristol, BS8 1TU, United Kingdom. E-mail: h.h.sievertsen@bristol.ac.uk. I thank Ea Hoppe Blaabæk, Simon Burgess, Stefania Simion, and two anonymous referees for helpful comments. The usual disclaimer applies.

their grades instead of selecting tracks aligned with their interests and comparative advantages. Understanding the implications of assessments for the behaviour of students, parents, teachers, and schools is therefore necessary to achieve the overall goals of the educational system. Because education affects lifetime earnings, health, and well-being and assessments play an important role in individuals' educational careers, assessments are also important for efficiency and equity across domains. Biases in assessments and the heterogeneity in access to assessments are sources of inequality in education according to gender, origin, and socioeconomic background. Finally, because assessment results also carry important consequences for individuals' educational opportunities and in the labor market, they are a source of stress and reduced well-being.

Keywords: education, assessments, human capital, inequality

Why Should Economists Care About Assessments in Education?

Assessments such as standardized tests and teacher evaluations of students' classroom participation are central elements of most educational systems. As Figure 1(A) illustrates, in both 2000 and 2018, more than 70 percent of schools across 42 countries used assessments to inform parents about their child's progress. Nevertheless the use of assessments varies widely across time and space. For instance, as depicted in Figure 1(B), the frequency at which assessments were used to evaluate teacher effectiveness ranged broadly across countries from less than ten percent to more than 90 percent across countries. And Figure 1(C) shows that most countries increased the use of assessments for national comparisons over the last two decades. The great degree of heterogeneity in the use of assessments and the changes over time indicate that no consensus has been reached yet on the best assessment system. The effects that assessments have on individuals and schools must first be determined before an optimal assessment system can be identified. While research in assessments have a longer history in education and sociology compared to economics, economists have also provided theoretical and empirical contributions to the topic. This chapter aims at summarising the research on assessments from the economist's viewpoint.

Why should economists care about assessments in education? The short answer is that assessments affect efficiency, equity, and well-being. Assessments are central components in school choice and accountability systems that are commonly used to maintain and improve school effec-

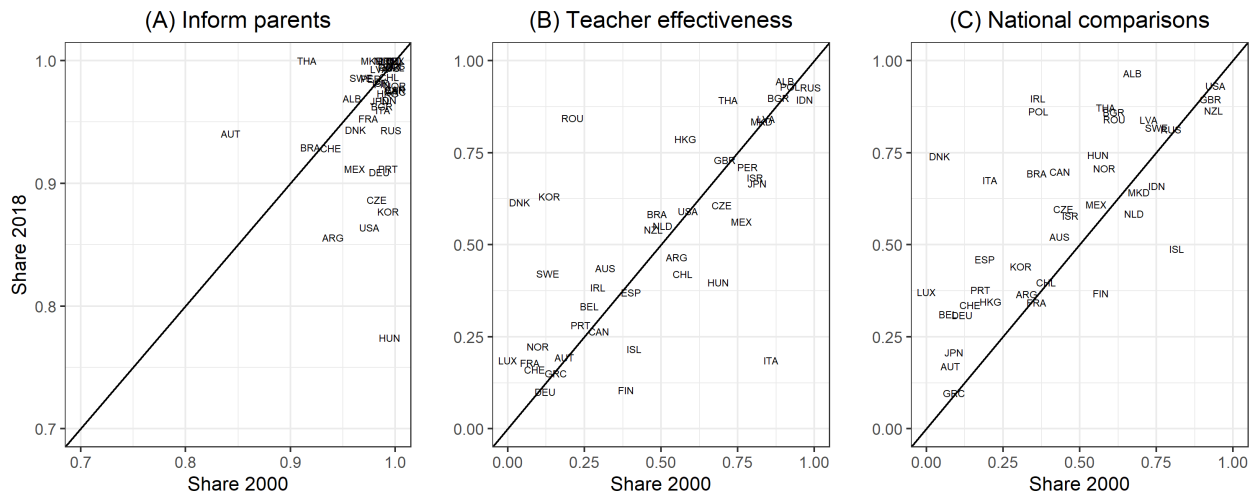


Figure 1: Share of schools using assessments for 15-year-old students to (A) inform parents about their child’s progress, (B) make judgments about teacher effectiveness, and (C) compare school to district or national performance. Source: Own calculations based on data from PISA 2000 and PISA 2018.

tiveness [Burgess and Greaves, 2021]. They affect drop-out decisions and course choices [Stinebrickner and Stinebrickner, 2012, Bar et al., 2009, Ahn et al., 2019] and are used to sort individuals into educational programs. Hence, providing access to (fair) assessments can therefore reduce inequality in access to education [Hyman, 2017]. However, popular assessment types suffer from biases based on gender, origin, and socioeconomic background; furthermore assessments cause increased stress and affect mental health negatively [Bach and Fischer, 2020, Heissel et al., 2021]. Understanding how assessments affect individuals is, therefore, important to understanding human capital development, equity, and well-being.

The longer explanation for why economists should care about assessments in education relates to the core economic principles underlying many components of assessments and to economists’ tools being well-suited for studying assessments. For example, how assessments affect students, teachers, schools, and parents can be studied in light of a principal-agent model [Bergbauer et al., 2021], where parents act as principals and schools and teachers act as agents. However, the school and teachers also act as principals where the agent is the student. The teacher cannot perfectly observe the effort of the student, and the parent cannot perfectly observe the effort of the teacher or school. The solution is a “contract” based on assessments, the results of which may affect the school’s funding directly in an accountability system or indirectly through school competition and

user choice. For students, the “contract” means that the assessment results are important for their future educational track and labor market success, which gives them incentive to study but also has consequences for their mental health and well-being.

The trade-off between the benefits of assessments for school effectiveness and equity and the potential costs in terms of mental health and well-being illustrates that no easy solution exists, and that detailed research across the social sciences that includes evidence from educational sciences is needed. This demand for research is even clearer considering the increased policy attention being placed on assessments. In the 2018 World Development Report, increased use of assessments is included as one of three key strategies to improve global human capital development [The World Bank, 2018]. However, the increased use of assessments has not been without controversy; for example, critics say it is harmful to learning [The Guardian, 2014]. The worry that assessments, and especially standardized testing, may distort students’ incentives is not new. Indeed, in the 1880s, Harvard University replaced a finer grading system with the six letters, A-F, to reduce the discretionary effect of grades on student motivation [Grant and Green, 2013]:

“The Faculty last year did away with the minute percentage system of marking, and substituted a classification of the students in each course of study in five groups, the lowest of which includes those who have failed in the course. It is hoped that this grouping system will afford sufficient criteria for the judicious award of scholarships, honorable mention, and the grade of the Bachelor’s degree, while it diminishes the competition for marks and the importance attached by students to College rank in comparison with the remoter objects of faithful work.” (Annual Report of the President of Harvard 1885)¹

The controversy over the use of assessments calls for more research on the consequences of assessments and the circumstances under which assessments might be more harmful than beneficial. And the toolbox of social scientists is excellently suited for that. The chapter is divided into two main parts. The first section includes information on the skills and knowledge measured by assessments and issues about measurement, such as biases, exogenous factors, manipulation, and cheating. The second section presents an overview of the effects that assessments have on individuals, including the importance for learning, well-being, and long-term outcomes.

¹This quote originally appeared in Smallwood [1969], but here it is taken from [Grant and Green, 2013].

Several aspects of assessments must be clarified before discussing this topic. First, in the context of this paper, all attempts to evaluate student learning and understanding are considered assessments. These attempts include standardized tests, oral examinations, teacher assessments of classroom participation, essays, and many other forms of evaluations. Unless otherwise stated, a standardized test refers to a centrally designed and evaluated assessment and grades are teacher evaluations of student learning. Second, internal assessments may be teacher assessments of classroom participation or internal tests designed and conducted by the teacher, while external assessments are often centrally designed and evaluated to be comparable across classrooms, schools, and even countries. Third, assessments can be blind or non-blind. In blind assessments, the assessor has no information about the student, but the assessor either directly observes the student, knows the student, or knows characteristics of the student in non-blind assessments. Fourth, assessments may be internally or externally evaluated. An assessment might be blind, but still be evaluated internally within the same school, in contrast to external assessments that are assessed by staff at other schools or by central institutions. Fifth, assessments can have varying degrees of importance for students and schools, also called stakes. The outcomes of low-stakes assessments have no direct consequences for students or schools; however, the outcomes of high-stakes assessments may determine whether a student can enroll in the educational program of choice or may affect school funding.

This chapter is written from an economist's perspective. As such, the literature covered is primarily based on contributions from economists. While research from the social sciences discipline is also addressed, this chapter does not do justice to the contributions that scholars in the social sciences and other disciplines have made to this topic. As another example, educational research has studied assessment for decades, long before economists cared about the topic. However, studies from this area of research are not covered in this overview either, not because they are not important but to focus the content on the economist viewpoint. Moreover, insights from research that focuses specifically on the measurement, interpretation, and use of assessments and psychometrics also are not covered detail. Nevertheless, several insights contributions by economists are built on the shoulders of insights from these disciplines.

Measurement

The core intention of an assessment is typically to evaluate students' learning and understanding in order to give accurate feedback to the student, parents, teachers, the school, and possibly the accountability system. However, the results of an assessment may also depend on factors other than student learning and understanding, such as the amount of effort the student invested in the assessments, the leniency of the assessor, biases, manipulation and cheating, and other external factors like noise and air quality. The following subsections identify various assessments and explain the skills and knowledge they evaluate, supported by evidence from the literature.

What Skills Assessments Measure

In the context of education, the objective of an assessment is typically to evaluate the learning of the student in contrast to measuring an innate and constant level of skill. Moreover, at least traditionally, the focus of assessments in education has been measuring cognitive skills in contrast to non-cognitive skills, such as grit, locus of control, motivation, interpersonal and social skills. Borghans et al. [2016] used four datasets representing individuals in the United States, the United Kingdom, and the Netherlands to shed light on what assessments measure. The four datasets included results from Intelligent Quotient (IQ) measures, standardized tests, school grades, and personality tests. Borghans et al. [2016] examined the correlation between the four measures within each sample and showed that the lowest correlations were found between IQ and grades and IQ and personality. The strongest correlations were found between IQ and results from standardized test and between grades and standardized tests. The correlation between grades and personality was also reasonably high. The authors concluded that these concepts measure different aspects of performance: While grades seem to capture elements of both of cognitive skills (strong correlation with standardized tests) and non-cognitive skills (strong correlation with personality), the standardized tests seem to capture mostly cognitive skills (weaker correlation with personality). Borghans et al. [2016] also studied how assessments predict future assessment outcomes; the findings are in line with the takeaways from the correlations. Finally, the authors also study how the various assessments predict later life outcomes: Grades and standardized test scores predicted later life outcomes, such as wages, health (BMI), and life satisfaction, better than IQ, possibly be-

cause grades also capture personality [Borghans et al., 2016]. The conclusion that school grades capture both cognitive and non-cognitive skills while standardized tests mainly capture the former agree with the conclusions in reviews in the psychological and educational and pedagogical literature (see Almlund et al. [2011] for a detailed review). The finding that grades and standardized tests capture slightly different aspects may explain why many educational systems still apply both types of assessments.

There is also some research evidence on what test scores do not capture, or more specifically, how teachers affect student behaviors and non-cognitive skills. For example Jackson [2018], uses data from North Carolina to quantify measures of behavior such as absences, suspensions, and grade repetition, and found that they were stronger predictors of later life outcomes than test scores. Importantly, Jackson [2018] showed that teacher value added in terms of test scores were only weakly correlated with teacher value added in behaviors. In other words, teachers who are good at leading their students to improved test scores are not necessarily good at helping them to improve their non-cognitive skills. The conclusions by Jackson [2018] are supported by results of Petek and Pope [forthcoming], who, based on data from Los Angeles, also concludes that teacher value-added on test-scores only are weakly correlated with teacher value-added for student behaviour, but that both matter independently for long-run outcomes.

Effort

Assessments will only give a good indication of students' understanding if the students actually attempted to show their level of comprehension of the material. However, when a test is not important to students, they may choose to invest little effort into demonstrating their understanding. This concern is especially relevant for low-stakes tests that carry little importance for the students themselves and for which the invested effort depends on the students' intrinsic motivation. This has led to the criticism of using low-stakes assessment as a tool to assess learning, especially across contexts where intrinsic motivation levels may vary across the groups compared. One example is the Programme for International Student Assessment (PISA) that is administered in more than 60 countries by the Organisation for Economic Co-operation and Development (OECD) to measure mathematics, reading, and science skills among 15-year-old students. Using the PISA scores to compare the learning across countries relies on assumptions of similar investments in

test effort and, thus, similar levels of intrinsic motivation across these countries. Gneezy et al. [2019] studied this by randomly adding extrinsic financial incentives to do well to students conducting a similar test in the United States and in China. The authors reported that the test scores for the Chinese students did not improve as a result of the added extrinsic incentives, but the financial rewards caused an improvement in the American students' performance. This pattern aligns with the hypothesis that the Chinese students had already maximized their effort because of their high level of intrinsic motivation to do well on the low-stakes test. In other words, adding another incentive for these students did not improve the performance because they already were at the limit. The American students, on the other hand, did not exert their maximum test effort because they had lower levels of intrinsic motivation, so they had room to improve their score in response to the incentive. Thus, different levels of intrinsic motivation across countries may explain some of the variation in PISA results across countries.

Another approach to studying the role of effort in assessment is to link students' responses to other behavior in the same setting. For example, the PISA assessment is linked to an extensive survey. Fuchsman et al. [2020] investigated inconsistent answers to the survey questions. For instance, a student may indicate liking school in response to one survey question but may indicate not liking school in response to a subsequent question that is worded slightly differently; this outcomes would suggest a lack of effort and lack of attention to the survey on the student's part. Using this measure of disengagement with the PISA survey as well as a measure based on item non-responses to the survey, Fuchsman et al. [2020] concluded that variation in effort can explain between 32 and 38 percent of the variation in PISA test scores across countries.

Grading Standards

Standardized tests are fairly straightforward to ensure consistent standards across time and contexts. These tests may even be multiple choice assessments with computer-based evaluations. However, as described previously, these tests do not measure all aspects of student learning and assessments involving human judgment, such as teacher grades, play an important role in many contexts. Due to the subjectivity and human aspect of these assessments, the grade a student receives for a given performance may vary across assessors, time, and contexts. Related issues, such as grade inflation, grade compression, and general variation in grading standards, are discussed

in the following subsection.

Aggregate statistics provide clear evidence of so-called grade inflation in many countries and settings. This is typically seen as an increase in grade point averages or an increasing share of a cohort receiving a specific degree classification. For example, in the United Kingdom, the share of students receiving the highest degree classification (a "first" or "distinction") increased from 15.7 percent in 2011 to 37.9 percent in 2021 [OFS, 2022]. In the United States, the share of high school students with top grades increased from 39 percent in 1998 to 47 percent in 2016 [Buckley et al., 2018]. While these changes may reflect changes in student composition, evidence suggests that this is the result of more lenient marking. [Denning et al., 2022] for instance, concludes that rising college GPAs in the United States cannot be explained by changes in student composition or actual student learning and exam performance, but can only be explained by more lenient grading standards. These findings are also confirmed by Buckley et al. [2018], who show that the standardized SAT scores of high school students in the United States did not improve similarly to the increase in grades from 1996 to 2016 [Buckley et al., 2018]. Using data on grades from 1982 to 2001 at Clemson University, Hernandez-Julian and Looney [2016] showed that about one-quarter of the grade inflation over that period was driven by students choosing subjects that were more leniently marked, another quarter was driven by improved student quality, but almost half the grade inflation was unexplained and, therefore, likely due to more lenient grading standards over time within the same subject.

The analogy to price inflation is obvious: for the same performance the student receives a higher grade today, compared to what the student would have received a year ago. This is just like consumers paying a higher price for the same product today compared to what they paid a year ago. However, the analogy is not perfect: There is no upper limit on how much nominal prices can increase. Grades, on the other hand, can only increase until everyone gets the best grade. Thus, the term *grade compression* has been suggested as a replacement for the term *grade inflation* [Babcock, 2010]. However, the term *grade compression* is silent about the direction of the change in the grading. It could, in principle, be compression in the lower part of the grade distribution. Furthermore, the theoretical implications of grade inflation and grade compression may be very different [Zubrickas, 2015]. Finally, grade compression is, to some extent, studied separately in terms of importance between finer and coarser grading systems.

Why does grade inflation arise? One source of grade inflation is the incentives at the individual assessor level to give a generous grade. Individual students are happy to receive good grades because those grades may help them gain admission into a good educational program or career and reflect a positive signal of their human capital to future employers. The individual teacher and school will also view good grades positively because they can reflect well on the teacher's instruction and may even have a direct financial impact on them. Moreover, the funding of the school or the salary of the teacher might be tied to student performance, either directly through accountability systems or indirectly by attracting future students [Diamond and Persson, 2016, Nordin et al., 2019]. In line with these predictions, Wikström and Wikström [2005] showed that in Sweden, grade inflation is higher in municipalities with more competition between schools.

Chan et al. [2007] formalized a signaling theory of grade inflation and showed that schools with many good students have an incentive to apply a more lenient grading policy to help their few mediocre students. However, employers have rational expectations, and if a school has too many weak students, the signaling value of good grades gets diluted. Moreover, Babcock [2010] showed that inflationary grading policies are strategic complements: if one school applies an inflationary policy, other schools will have little difficulty following them, and as a result, grade inflation becomes contagious. Ehlers and Schwager [2016] extended the model of Chan et al. [2007] by adding a reputation aspect: after applying a generous grading policy, the next cohort suffers because employers have learned about the lenient grading policy and will consider this when evaluating future graduates from that school. They showed that this dampens the degree of grade inflation. Zubrickas [2015] formally modeled the grading policy on the student-teacher interaction level. In his model, teachers use grades to incentivize student effort. By applying a more lenient grading policy, the teacher can motivate students farther down the grading distribution at the cost of those at the top of the distribution. In the dynamic version of the model, employers learn about this and realize that that top grades are less informative than they initially thought, which, in turn, gives the teacher an incentive to give the best grade even farther down the ability distribution. While most theories predict that high ability students tend to bear the individual costs of grade inflation, there are some exceptions. Schwager [2012] set up a model, according to which employers use signals other than grades if grade inflation makes the signaling value of grades too weak. As a consequence, they use social class as a signal of ability, which is, of course, especially costly for

high ability workers from disadvantaged backgrounds. The winners are high achieving students from the higher class, because they avoid the competition from high achieving students from the lower class.

Why is grade inflation an issue of concern? On an individual level the theories on optimal grading policies predict that it harms the high skilled students [Babcock, 2010, Zubrickas, 2015]. Assessing the individual consequences of grade inflation, Nordin et al. [2019] studied the impact of institution level grade inflation in Swedish upper secondary schooling. The authors found that students attending institutions with high rates of grade inflation have higher earnings later in life, driven by their post-secondary schooling choices. However, this comes at the cost of high achieving students at institutions that do not apply such high rates of grade inflation. An unusual type of grade inflation happened in France in 1968 as a consequence of a student protest. The protest interrupted the exam period, and as a consequence, passing thresholds were lowered. Students at the margin of passing benefited from this treatment by gaining access to educational programs to which they otherwise would not have access. The students at the margin also received higher wages later in life and their offspring also had better educational outcomes [Maurin and McNally, 2008]. While the individual students clearly benefited from more lenient marking in this case, it may very well be that these outcomes would be very different if this policy was applied permanently. For example in this case, the expectation of an exam still gave the students an incentive to learn. If exams were cancelled permanently, this would not be the case. Another potential consequence of grade inflation is that it may affect major program of study and drop-out decisions. Ahn et al. [2019] developed a model in which the higher level of observed study effort by women is due to higher perceived benefits of grades. If some programs or institutions have less lenient grading policies, which they document for STEM degrees, women may be discouraged from entering these fields. Bar et al. [2009] studied the consequences of the "Cornell experiment" in the mid 1990s when Cornell started to publish median grades for courses. The authors reported that this led to an increased demand for the courses with higher median grades, which led to a higher overall grade point average and grade inflation, in line with the findings of Hernandez-Julian and Looney [2016] discussed previously and also in line with the findings of Sabot and Wakeman-Linn [1991]. The move to the more leniently graded courses was mainly driven by weaker students. Another cost of grade inflation is that if teachers are very lenient, students may not have a strong incentive

to study hard, so they may reduce their study effort and as a result learn less. Indeed, [Babcock, 2010] found a negative association between the expected mark (leniency) and study effort. This correlation is in line with the empirical finding that higher grading standards lead to more learning, especially for high ability students [Betts and Grogger, 2003, Figlio and Lucas, 2004]. At the aggregate level, the grade inflation may have important welfare costs in terms of worse matches on the labor because of the noisy signaling value [Fredriksson et al., 2018].

Are there any benefits to grade inflation? The main cost is the distorted and noisy signal, which may lead employers to disregard the value of the signal. Boleslavsky and Cotton [2015] modeled grading standards and school quality in a setting where investment in school quality also depends on grades. As a consequence of more lenient grading standards, schools have an incentive to invest more in quality to improve the distinction of their graduates. Although students might face an incentive to shirk under more lenient grading policies, Boleslavsky and Cotton [2015] showed that these moral hazard concerns are dominated by the students' anticipation of higher investments by the school.

Biases

At least two sources of bias can be found in assessments. First, the design of the assessment and questions asked may favour some groups of students. Second, the evaluation of the performance in the assessment might be consciously or unconsciously biased towards subgroups of the population. For example, if a girl receives a lower mark for the same performance in an assessment compared to a boy, this would be classified as a bias discriminating against women. In the following discussion, I will focus on the second source of bias. However, there is some evidence on the first source of bias as well. For example, there is evidence on gender differences in behaviour when stakes are high, suggesting that an assessment system focusing on fewer and higher stakes assessments, would affect girls and boys differently [Azmat et al., 2016]. It is very difficult to empirically identify the first source of biases, as discussed by Sackett and Kuncel [2018], which might explain why empirical research has focused more on the second source of bias.

Blind assessments, in which the performance of the student is evaluated without any knowledge about the individual student's characteristics, cannot be biased with respect to the background of the student. However, if the assessment is non-blind, such that the assessor of the

performance either observes and possibly knows the student or can infer the name, gender, or race of the student, the feedback may be influenced by that information. For example, teachers may be more lenient in grading the completed exams of students who typically do well on tests. Alternatively, teachers who give feedback to hundreds of students at the end of the term, who may have difficulty remembering how well individual students did on prior exams, may unconsciously incorporate the group average into the assessment. It is difficult to remember how well each individual student did, and the teacher might unconsciously include group average in the assessment. For example, if boys historically have done well in mathematics, all boys may receive a slightly higher mark, as a result of stereotyping behavior or statistical discrimination [Phelps, 1972].

The observation that blind assessments per definition do not include any biases driven by observable student characteristics has been used to empirically test for evidence of bias in non-blind assessments. The approach was to exploit the fact that students often are exposed to both types of assessments in the same subject and that any systematic differences between these assessments may point toward a gender bias in the spirit of assessments of gender discrimination in orchestra auditions [Goldin and Rouse, 2000]. Alternative explanations exist for the systematic differences between non-blind and blind assessments, which typically also differ on other dimensions. For example, the blind assessment may be a written and centrally assessed exam, the results of which are then compared to the results of a blind assessment of the classroom performance by the teacher. However, the literature has gone a long way to rule out these alternative explanations for systematic differences between blind and non-blind assessments. Among the first studies to show evidence of systematic differences in assessments was Lavy [2008], who compared two scores of high school exams in Israel. The first score was based on externally and anonymously assessed state exams. The second score was from school internal exams that were assessed by the student's own teacher. Comparing these two scores, Lavy found evidence of bias against male students, which increased the gender gap, as girls already outperform boys. Interestingly, Lavy found that this bias varied by teacher characteristics and there is little evidence of stereotyping or statistical discrimination.

In England, Burgess and Greaves [2013] showed that black Caribbean pupils were more likely to receive a poorer assessment by their teacher relative to the blind final test score, than their white peers. In contrast to Lavy [2008], these authors found empirical support that indicated stereotyp-

ing was the most likely reason some groups systematically receive more lenient teacher assessments. More specifically, if students belonged to a group that usually did well at a specific school, they were more likely to receive a more optimistic teacher assessment [Burgess and Greaves, 2013].

Similar to the findings from Israel [Lavy, 2008], Falch and Naper [2013] found that Norwegian teachers gave higher marks for girls across all subjects, compared to the blind exam scores, at the end of compulsory schooling. Interestingly, Falch and Naper studied the hypothesis that the difference was driven by girls doing worse than boys in competitive environments as suggested by laboratory experiments [Gneezy et al., 2003]. The authors exploited that in some regions the results impacted enrollment in subsequent schooling, and in others it did not. However, they concluded that the competitiveness could not explain the gender gap between blind and nonblind grades. Instead, like in the Israeli context, they found suggestive evidence that the gap was related to teacher characteristics.

Studying the context of the last year of compulsory schooling in Denmark, Rangvid [2015] discovered that boys, children whose parents have lower educational background, and migrants received lower teacher scores than girls, children of highly educated parents, and natives, with similar blind exam performance. The largest difference was found by parental background, and the smallest difference was by found to be by origin. Studying high school students in Denmark and comparing teacher assessments to blind written assessments, Burgess et al. [2022] find that compared to girls, boys were more generously marked by their teacher in mathematics, but less generously marked by their teacher in Danish.

In the context of 16-year-old Swedish students, Lindahl (2016) investigated whether differences between teacher assessments and national exam results related to the teacher and student being the same gender or sharing a minority status. Contrary to expectations based on evidence from educational literature in general [Dee, 2005], Lindahl concluded that the differences for girls and non-native students declines when the teacher was female and of non-native origin, respectively. The author interpreted these findings as an indication that teachers were trying to actively compensate for having a similar background to students and the potential (unconscious) biases that might involve.

Terrier [2020] compared blind and non-blind assessments for middle school children in France and showed that girls received more generous marks from their teacher in mathematics but not

in literacy. The finding that girls are more generously marked in mathematics is in line with the earlier evidence from France by [Breda and Ly, 2015], who compare blind and non-blind assessments in higher education in France, and find that girls are more generously marked in non-blind assessments in male-dominated fields. Finally, Lavy and Megalokonomou [2019] found large variation in the gender gap between blind and non-blind assessments across subjects for high school students in Greece.

A particularity of the university system in the United Kingdom is that students apply to programs based on their predicted exam grades, and they receive offers that are either unconditional or conditional on actual exam grades. These predictions are made by their teachers, and any potential biases in predictions may therefore have long-term consequences for these students. Indeed, Murphy and Wyness [2020] showed that only 16 percent of the grades were accurately predicted. Most of the teacher predictions were too optimistic, with 75 percent of the predictions being higher than the actual exam result. A finding of concern reported by the authors was the existence of systematic differences by student background and school type. High achieving low socioeconomic students received a teacher prediction that was considerably lower than that of their high achieving peers that are not from low socioeconomic households. At the same time, high achieving students in state schools received less optimistic teacher predictions than high achieving students in private schools. In addition to teacher bias, Murphy and Wyness listed differences in responses to predictions or a different trend in achievement growth as explanations for these patterns and the showed that these predictions have implications for university application behavior. This study by Murphy and Wyness [2020] is slightly different to the earlier cited papers comparing blind and non-blind assessments, as they study differences in predicted results and actual results. One important difference is that these predictions might be intentionally optimistic to encourage student investments and improve student confidence. However, any systematic differences between predicted grades and actual results may come from the same source as in previously cited papers.

A few studies have applied alternative approaches to identify biases in non-blind assessments. Hanna and Linden [2012] conducted an experiment in India in which they randomly varied the characteristics of children on their exam sheets before they were marked by local teachers. They found no evidence of discrimination on the basis of age or gender, but discrimination against children from a lower caste was observed. Rangvid [2019] exploits a flaw in a grading reform

whereby names were replaced by student identifiers to ensure an unbiased assessment. Because the student identifier includes the first four letters of the students' names, this made the gender identifiable for some names but not for others. The study found no clear evidence of significant gender differences between the scripts in which the gender could be identified and the scripts for which the gender was not identifiable.

To summarise, evidence has verified systematic gender differences between blind and non-blind assessments favouring girls from Israel, Norway, France, Denmark, Greece, United Kingdom, and Sweden. This difference varied across subjects. Similarly, evidence has uncovered systematic differences between blind and non-blind assessments that favors natives and work against students with immigrant status in England and Denmark. Finally, evidence from Denmark, India, and the United Kingdom suggests that students from disadvantaged backgrounds are less generously marked by their teachers. The consequences of these systematic differences between external and internal assessments are discussed later.

Notably, in most of the studies cited, standardized blind tests were used as the unbiased benchmark. However, statistics showing mean differences in standardized blind tests by gender, socioeconomic background, and origin are often used as a basis to hypothesize that these assessments also suffer from systematic biases. While biases cannot be inferred from such mean differences, disproving that these assessments are biased is also difficult. One imperfect approach to study this question is to assess the correlation between blind test scores and other (real) outcomes. If the assessment systematically discriminates against a group, the correlations would under-predict the outcomes for this group. Another strategy is to compare the change in the explanatory power of test scores with and without controlling for some of these group characteristics. Ideally, the predictive power should be unaffected by adding controls for student background. All these tests are far from perfect, but Sackett and Kuncel [2018] discussed the various findings and approaches found in research on biases by gender, socioeconomic background, and origin and concluded that the "overwhelming conclusion across decades of research is that tests are not biased against women and racial/ethnic minority group members in their use in predicting subsequent academic performance" and that "there is now extensive evidence countering both the claim that admissions tests measure nothing but SES and the claim that the apparent predictive power of admissions tests is an artifact that disappears when SES is controlled" [Sackett and Kuncel, 2018].

Manipulation and Cheating

Assessment outcomes carry important consequences for the students who may not be able to enroll in the desired educational program if their test score does not meet the minimum required for enrollment, for the teacher who may miss out on a bonus or even lose a job, and for the school that might lose funding as a consequence of unsatisfactory student performance or may simply face lower demand from prospective students. These high stakes mean that students, teachers, and the school have strong incentives to cheat or manipulate test scores. Research has found evidence of such a response, ranging from "soft" manipulation like targeting the teaching toward what is tested ("teaching to the test") to more direct evidence of cheating and manipulation by actively retrospectively changing the answers on students' completed exams.

Linking assessments closely to rewards in accountability systems provides strong incentives to improve performance in the subjects included in the accountability system, but this may then come at the cost of students' performance in non-tested subjects or grades [Holmstrom and Milgrom, 1991]. Studying the accountability system in Chicago Public Schools in the mid 1990s, Jacob [2005] showed that the system led to improvements in math and reading achievement compared to earlier trends and compared to development in other regions. However, non-tested subjects and cohorts saw no improvements, and improvement was mainly driven by test-specific skills and student effort on the test. As noted previously, focusing effort on subjects that are tested and on specific test-taking skills is called teaching to the test. Whether strong incentives that lead to teaching test are bad for learning depends on the costs of monitoring learning and the cost of learning and on the design of the test. If these costs are very high, strong incentives that lead to teaching to the test may still be the most efficient solution [Lazear, 2006]. Likewise, if the test is well-designed and covers important aspects, focusing on these topics might be worth it. It might be the case that an assessment provides an ideal measure of the underlying human capital that the policy makers care about, and focusing efforts on this assessment could therefore be aligned with the the social goals. However, as discussed in Neal [2013], assessments are often only designed to achieve one goal, but used for two goals. For example when low-stakes assessments are designed to measure student learning, but are also used in accountability systems, strong teacher incentives might mean that the assessment is not achieving any of the goals.

Closely related to teaching to the test is the concept of teaching to the rating. Accountability systems may lead to incentives to improve the test scores of certain students. For example, if the accountability is linked to average student performance, the school has an incentive to target the teaching toward students associated with the lowest marginal cost for improvements to assessment results. In some cases, such as in the No Child Left Behind policy in the United States, the accountability system is linked to several cutoffs in indicators that incentivize targeting resources to students who score right around these cutoffs. Reback [2008] used data from Texas to provide evidence of such behavior. Importantly, they showed that targeting the resources to students around incentivized cutoffs improves their performance, but that this happens at the cost of students at other parts of the distribution who improve less than they otherwise would. This type of response is closely linked to the concept of "cream-skimming" [Heckman et al., 2002] and multi-tasking where agents focus on the task that is measured at the expense of other tasks that are not measured or counted [Holmstrom and Milgrom, 1991]. Related to teaching to the rating and teaching to the test, Jacob [2005] also finds evidence of the strategic placement of students in terms of placing students in special needs programs to avoid including their performance in the accountability system.

Teachers often face strong incentives to manipulate their students' assessment results, and various approaches to detecting manipulation have uncovered solid evidence that teachers do engage in such practices. For instance, Jacob and Levitt [2003] studied unusual blocks of correct answers in Chicago Public Schools in the 1990s to exploit that the easiest way for a teacher to change test scores is to adjust blocks of answers rather than randomly adjusting answers for individual students. All students in one classroom having identical, correct answers to a set of questions, such as 8 through 15, but different answers to all other questions would be an indication of such a manipulation, especially if the blocks of correct answers vary from classroom to classroom. Also, evidence of students in a classroom experiencing unusual test score gains, with their scores suddenly improving far beyond what their regular trajectory would predict and then returning to their old trajectory, was also considered an indication of manipulation among the school children in Chicago [Jacob and Levitt, 2003]. Another sign of manipulation is unusual test score distributions that is aligned with teacher or school incentives. For example a test score distribution that shows unusual low mass below incentive cutoffs and unusual high mass just above incentive cutoffs, suggest that their

results were pushed above the thresholds. This approach was used to detect teacher manipulation in Diamond and Persson [2016] and Dee et al. [2019].

Using the first two strategies described, Jacob and Levitt [2003] uncovered evidence of test score manipulation in four to five percent of all classrooms in Chicago Public schools. As these strategies rely on unsophisticated manipulation, they are easy to detect using statistical rules. In the setting of the New York State Regents examinations, teachers locally mark their own students' exams. Dee et al. [2019] found that the test score distributions in the Regents exams exhibited discontinuities in line with students being pushed above important cutoffs in the marking. Such sharp discontinuities are unlikely to happen due to teaching to the rating. The authors further revealed that more than 40 percent of scores below important cutoffs were manipulated and moved above the cutoffs, accounting for six percent of all tests. Dee et al. also showed that this type of manipulation disappeared once the marking became part a centralised system. Manipulation by teachers has important consequences for students who have a greater chance chances of graduating from high school. However, the manipulation also has harmful effects as some students do not complete advanced courses as a consequence of being manipulated above a cutoff. This points to the importance of grades for updating beliefs about investments and skills. Students who are pushed above a threshold might be misled to believe that their understanding of the material is better than it actually is, and they might therefore not invest sufficiently in terms of subsequent study effort.

Research shows that teacher manipulation also occurs in settings with softer incentives than in the New York setting studied by Dee et al. [2019]. For example, in Sweden, where no direct incentives to manipulate test scores exist, Diamond and Persson [2016] show evidence of test score manipulation in high schools. The incentive for the schools and teachers in this setting was mainly to attract future students by documenting high passing rates and high grade point averages. While test score manipulation for financial reasons may go against the teacher's professional norms, the authors demonstrated that the test score manipulation mainly occur for students who do worse than predicted given prior performance. They did not find any evidence that teachers were more likely to manipulate the performance of students with certain background characteristics. In other words, teachers appear to be trying to correct for a bad test day and for potentially being exposed to the negative factors outside the control of the students, that I discuss in the next section. As in

the case of the United States, test score manipulation has long-term consequences for the students in Sweden. The students who had their grades manipulated upward were more likely to enroll in a university degree program, complete more years of education, were less likely to have a child as a teenager, and had higher earnings at age 23.

Lin and Levitt [2020] investigated student cheating on exams by comparing the number of identical answers for neighboring students to the number of identical answers for students sitting farther away. They found that when students sat next to each other, they had more shared incorrect answers and more shared correct answers, compared to if two students were sitting more than two seats apart. The increase was only found among students sitting directly next to each other. There was no significant increase in the number of shared answers for students sitting behind each other or farther than two seats apart.

External Factors

Most assessment situations involve a number of factors that affect student performance that are outside the control of the student and often also outside the control of the teacher or school [Kane and Staiger, 2002]. Such external factors, like temperature and air quality, tiredness, and well-being at the time of the assessment may affect students' condition and ability to perform well. Evidence of the impact of some of these factors is provided in this section.

Air Quality and Pollution

Evidence from various settings has shown that individuals' performance is affected by air quality and pollution, for example, impacting productivity on the labor market [Graff Zivin and Neidell, 2012] and the performance of chess players [Künn et al., forthcoming]. Importantly, evidence has also shown that pollution affects students' performance in important assessments. Ebenstein et al. [2016] studied the performance of students completing high-stakes exit exams in Israeli secondary schools, known as the Bagrut. Students' performance on these exams has implications for their chances to enroll in university programs. The authors assessed how air pollution affects the students' performance in these tests by exploiting day-to-day variation in PM_{2.5} levels due to forest fires and sandstorms. Because the students sit for the Bagrut across several days and sit for several tests, the authors were able to compare performance across different levels of pollution exposure

for the same students. Importantly, the students could not self-select into an exam time and, instead follow a pre-determined schedule unrelated to individual and school characteristics. The main finding was that a one standard deviation increase in air pollution caused a four percent of a standard deviation drop in performance, with larger effects for boys, low-performing students and students from disadvantaged backgrounds. As these exams are important for subsequent university enrollment, the pollution also impacted later life outcomes. Indeed, Ebenstein et al. [2016] found that higher pollution levels at the time of the test causes lower levels of completed education and lower earnings.

Tiredness

Individual health and tiredness at the time of the assessment is likely to also affect performance. However, obtaining accurate measures of health and tiredness at the time of assessment to enable these factors to be studied is very difficult, as is finding variations in these factors that are not related to other student characteristics. For example, students from disadvantaged backgrounds may, on average, be more tired at the time of the assessment because of lower quality housing than their peers from other households, and this variation in tiredness will then be confounded by other factors that vary across these students, such as general educational support and nutrition. To circumvent these issues and still shed light on the importance of tiredness, Sievertsen et al. [2016] exploited variation in time of the day the individual was assessed. The authors relied on the fact that the timing of the test depends on a combination of the regular schedule and the school level computer availability and is, thus, unrelated to individual characteristics. Similar to Ebenstein et al. [2016], Sievertsen et al. [2016] compared the test performance of the same individual sitting for several tests and found that for every hour later in the day, that the test was taken the test performance decreases by 0.9 percent of a standard deviation. Interestingly, the authors found that breaks can mitigate these effects.

Distractions

A dog barking outside the testing room or neighbouring students making distracting noises will likely also affect student performance on tests; however, there is limited evidence on causal effects of such distractions. One study has, however, studied the effect of wider distractions in terms of

the value of leisure around the test-taking and test preparation time. Metcalfe et al. [2019] investigated the effect of having assessments in times of large football tournaments, such as the European Championships and the World Cup, for students in England sitting the General Certificate in Secondary Education or GCSE achievement test. These tournaments are very popular and involve broadcasting of matches, parties, and other distractions that are likely to raise the value of leisure and increase the opportunity costs of studying for exams. The authors uncovered that exposure to this distraction mainly is driven by year of birth, as these tournaments only take place every other year, and found that students exposed to this distraction do worse in the assessments. The effects are largest for boys from disadvantaged backgrounds. While the study by Metcalfe et al. [2019] provides insights into how distractions affect learning in general, given that the distraction is centered around assessment, it also points to the importance of the environment around the assessments.

Stress and Nutrition

While general stress levels and nutrition quality might depend on the students' background, in the sense that they are correlated with parental background, stress levels and nutrition quality are often outside the control of the individual child. Heissel et al. [2021] show that children in poorer neighbourhoods have higher stress levels and that higher levels of stress are associated with lower test scores. Studying the impact of nutrition and food, Bond et al. [forthcoming] investigated the as good as random timing of the assignment of the Supplemental Nutrition Assistance Program (SNAP) benefits in the United States. The authors exploited that in some states the monthly timing of the benefits depends on the first letter of the surname, combined with information on names with the timing of SAT college admission exam. The results show that students who attend this high-stakes assessment just before they would receive new benefits (i.e., at the time of highest food insecurity) score six percent of a standard deviation lower on the exam, compared to those sitting for the exam just after receiving the benefits. Because the SAT results are important for university acceptance, these results carry over into a lower probability of attending a four-year university program for the low-income students.

There is also evidence of systematic differences in how students respond to the stakes of assessments. Azmat et al. [2016] showed that girls outperformed boys in tests, but the gap declined

as the stakes of the test increased. One possible explanation for this pattern is that boys exhibit low test effort when stakes are low. An alternative explanation is that girls perform worse under pressure than boys. One piece of evidence going against the first explanation is that girls actually perform worse on low stakes tests. The authors hypothesized, therefore, that differences were due to gender differences in response to pressure [Azmat et al., 2016].

Assessments and Individual Behavior and Outcomes

How do assessments affect the individual? Assessments provide information about the learning progress. The individual might use this information to update beliefs about the returns on study effort and beliefs about their own abilities. Assessments also provide incentives. A good result in an exam may provide easier access to the desired educational program, a scholarship, or a dream job. Assessments, therefore, often also provide an incentive for the student to invest in study effort and maybe select courses in which they know they will get a better mark, which affects their learning and human capital development directly. In this section, research on how assessments, holding underlying human capital constant, affect individual behavior and outcomes is discussed.

Learning

A return to the initial discussion on assessment types and objectives in this chapter will be useful to the discussion on how assessments affect learning. Standardized testing is often introduced in connection to accountability systems. The contrasts to these schemes are internal tests that are designed and conducted by the teacher. While the former has the advantage of producing results that are comparable across classrooms, schools, and maybe countries, the latter has the advantage of being more targeted toward the specific classroom teaching. While pedagogical research points to standardized tests being ineffective for (or even detrimental to) learning, because these assessments provide no detailed feedback that the teacher and student can use to improve learning and because the assessment is not tailored to what is taught in that classroom [Guskey, 2007], still such assessments might also affect learning by acting as incentives for students, teachers, and schools to invest effort. The incentive mechanism is closely linked to the information mechanism. Assessments provide information to teachers and students about their progress. The students (or the

parents) can learn about the returns on study effort and adjust effort levels accordingly or change study strategies. Likewise, teachers might adjust teaching methods or target specific areas of the curriculum following an assessment.

Looking at the big picture, Bergbauer et al. [2021] compared PISA attainment data for two million students in 59 countries over the period 2000 to 2015 together with information about the use of assessment systems in the schools. Based on a cross-country panel-regression exploiting changes across time within countries, but controlling for aggregate time trends and country level differences, they found that standardized testing raised attainment in PISA test scores. Importantly, they found that the benefit of testing is seen in low-performing countries, whereas they found no evidence of a significant association between standardized testing and PISA results in countries that already do well in the PISA assessment. Another important finding of Bergbauer et al. [2021] is that they only found a significant positive effect from standardized testing, but no evidence of any impact of internal testing and reporting. One threat to these conclusions is that national standardized testing might be introduced simultaneously with other policies that can explain part of the effect. Bergbauer et al. [2021] ruled out and control for the impact of number of alternative policies and explanations; their findings were also supported by a micro-level study focusing on an external shock to testing in Denmark [Andersen and Nielsen, 2020]. In this study, the authors investigated the subsequent performance of students who had missed sitting for a national standardized assessment in Danish schools due to an IT crash. They discovered that these students performed worse on tests in subsequent years, an effect that was stronger for students for disadvantaged backgrounds. One downside of these approaches is that the outcome is closely related to the treatment. Students may do worse in subsequent assessments because they had less practice in taking tests. Future research can link the absence of testing to other outcomes as well.

It is not entirely clear how the assessments affected learning in the case of Bergbauer et al. [2021] and Andersen and Nielsen [2020]. One potential mechanism is that not completing a test means that students miss out on a chance to get (unbiased) feedback on their learning. In the context of Andersen and Nielsen [2020], the next assessment was three years after the IT-crash, illustrating that a missed opportunity for feedback and reacting to this feedback might mean three years of not getting the extra support needed or not making required changes to study strategies. To shed light on how test feedback promotes learning Beuchert et al. [2020] identified the effect

of scoring just below a cutoff (e.g., scoring "Below average" instead of "Average"). The authors found that children receiving negative feedback in mathematics do better in subsequent maths assessment compared to students who only did marginally better at the initial assessment. They also found evidence of a positive spillover to other subjects. These findings suggest that test feedback is used actively and affects future learning. This conclusion is in line with earlier findings on students in higher education. Using variation in exam feedback in universities in the United Kingdom, Bandiera et al. [2015] found that students who received feedback on their exam results in the previous term performed better than those who did not receive this information. This effect was mainly driven by students with less initial information, suggesting that assessments and feedback play an important role in informing the students about the learning process.

As previously discussed, evidence from several countries has shown that non-blind teacher assessments differ systematically from blind assessments. A number of the aforementioned studies documenting these biases also assess the consequences for students. Terrier [2020] calculated individual teacher favoritism for girls and exploited variation in this measure, and under the assumption of the random assignment of teachers to students in sixth grade, she demonstrated that this causes a widening gender learning gap in mathematics and French. In mathematics this widening was driven by girls performing better (and boys not being affected) and in French the gap was driven by boys doing worse (and girls not being affected). Furthermore, for girls, being assigned to a teacher who favours girls increases the likelihood of choosing a science track in high school four years after being exposed to that teacher showing favoritism towards girls. Similar conclusions have been reached in other settings; for example, Lavy and Megalokonomou [2019] provided evidence that assessments matter for subsequent learning, potentially through affecting students' expectations and subsequent investments in study effort. These findings are in line with the findings of Burgess et al. [2022] that indicated that a semi-blind assessment in mathematics at the end of high school reduces the gender gap in enrolling and graduating in subsequent university STEM degrees.

A growing area of the literature addresses ways that assessments motivate student learning, especially if the result of the assessment has important consequences for the individual. Hvidman and Sievertsen [2021] examined a grading reform that required high-stakes grades to be recoded. As students knew that the change in their grades was due to these recodings and not due to their indi-

vidual characteristics or performance, the only reason for them to react to this reform was that the grade point average was important for them. Hvidman and Sievertsen [2021] showed that students who were negatively affected did better in subsequent assessments to compensate for the negative recoding. This reaction had long-term effects on their enrollment in subsequent educational programs. Interestingly, the response was driven by girls. There was no response in the short run or in the long run by boys. Another setting that provided evidence of grades acting as incentives to increase learning is Greece, where the chance to retake an important high-stakes assessment was shown to improve learning considerably [Bizopoulou et al., 2022]. Somewhat contradicting these strong incentive responses, a study on university students in Texas found no evidence that students who were close to receiving a better letter grade (A-F) did better than students who were farther away, despite their stronger incentive [Grant and Green, 2013], which might be explained by differences in the importance of the assessment.

Belief Updating, Course Choice, and Drop-Out Decisions

Belief Updating and Choices

Beyond using assessments to adjust study effort, students can use them to update their beliefs about their skills. Stinebrickner and Stinebrickner [2012] studied how students in a Canadian university used feedback from assessments in their decisions about study effort and drop-out decisions. They found that at university entry, students were too optimistic about how they would do, which was mostly explained by being too optimistic about their own abilities (in contrast to being too optimistic about their study effort). However, students then updated their beliefs actively based on assessment results, with some heterogeneity depending on views of explanations for assessment results. The take-away message is that 40 percent of all university drop-outs result from students' learning about their abilities from assessments. Studying the same setting on university students in Canada, but focusing on major choice, Stinebrickner and Stinebrickner [2011] found that grades play an important role in explaining why fewer students major in STEM subjects than beliefs at university entry suggest, in line with the previously discussed findings by [Ahn et al., 2019] related to how gender differences in perceived benefits of grades possibly explaining gender differences in pursuing and completing STEM majors. Anaya et al. [2022] found

evidence that girls have lower self-perceived mathematics ability than boys, even given the same mathematics test performance. Anaya et al. [2022] also showed that the self-perceived mathematics ability, given math performance, was different if parents worked in STEM occupations. On the other hand, Andrew and Hauser [2011] documented limited evidence of updating expectations and beliefs in response to changes in grade point averages. Responses were modest and only in reaction to large changes. The topics of belief formation and expectations in education, including the role of feedback from assessments, has been studied by sociologists for decades. For instance, Holm et al. [2019] show that assessment results affect college enrollment and completion, and that students from disadvantaged backgrounds react more to these signals. Birkelund [2020] show that children of immigrants react less to assessment results in their decision to enrol in high school.

Grades affect individuals' choices through information about the individuals' own abilities but also by giving them incentives to attend courses that are more leniently assessed if they care about their overall grade point average, as discussed in the section on grade inflation [Bar et al., 2009, Ahn et al., 2019].

Relative Feedback

How students use assessments as information to update their beliefs depends on both their prior beliefs and how the feedback is provided. One key aspect of the latter is whether the assessment feedback provides information on the relative performance. To study the importance of relative feedback, Azmat et al. [2019] carried out an experiment in a Spanish classroom where students in the control group only received information about their own performance, while students in the treatment group also received information about their relative performance. In subsequent assessments, students in the treatment group were less likely to pass, and they received lower grades. Using survey data the authors also show that the treated students were more satisfied. The response to relative performance information agrees with a theory that students, on average, underestimate their relative rank. This suggests that students care not only about their absolute scores, but also about their relative rankings.

Teacher Bias and Student Outcomes

As discussed previously, Terrier [2020] showed that teacher favoritism in non-blind assessments affect high school track choice. This was also confirmed by Burgess et al. [2022], who showed that semi-blind assessments in mathematics at the end of high school reduced the gender gap in enrolling and graduating from university STEM degree programs. The counter-factual treatment would be a semi-blind assessment in another subject (like Danish) and mainly a non-blind teacher assessment in mathematics. They also all sit for a fully blind mathematics exam. As they show that teachers are more lenient to boys in their assessments, compared to the exams, this points to the importance of assessments for track choice.

Students who were under-predicted in the United Kingdom [Murphy and Wyness, 2020] were also less likely to apply to a high-tier university and more likely to be overqualified for the programs to which they apply. Taken together, the systematic differences in prediction error by student background and the implications of the prediction error for university application behavior suggests that the system of using predicted grades for university applications may have implications for getting high achieving students from disadvantaged backgrounds matched with appropriate university programs.

Heterogeneous Belief Updating

Some evidence of gender differences in response to grades and grading policies has been reported, which explains part of the gender difference in enrolling and majoring in STEM degree programs [Ahn et al., 2019, Anaya et al., 2022]. At the aggregate level, Bergbauer et al. [2021] concluded that increased standardized testing is especially useful in low-performing countries. Likewise, at the micro level, Andersen and Nielsen [2020] uncovered evidence that children from disadvantaged backgrounds suffered the most from not having access to standardized testing. On the other hand, Beuchert et al. [2020] found no difference in the response to test feedback by parental background. While there is evidence that students from disadvantaged backgrounds have lower expectations about their future educational careers than students from more affluent backgrounds who do equally well in assessments, Karlson [2019] found no evidence that students from disadvantaged backgrounds were less responsive in their updating of expectations, which somewhat

contrasts the findings from Holm et al. [2019], but also suggests that heterogeneity in responses will depend on the setting. Finally, Birkelund [2020] that children of immigrants were less response to grades in their high school enrollment decision.

Rank Effects

A rapidly growing literature has examined the role of ordinal rank in the classroom setting. These ranks are often determined based on assessments. There is solid evidence that students' relative rank in the classroom affects their later outcomes. Exploiting the natural variation in peer achievement in British primary schools, Murphy and Weinhardt [2020] show that having a higher rank improves student performance and confidence in corresponding subjects during secondary school. Subsequently, concurrent rank effects have been found to impact academic outcomes and risky behavior in high school [Elsner and Isphording, 2017, 2018], and academic achievement at university [Elsner et al., 2021]. The effects of a students relative position in elementary school have been shown to have lasting impacts on subsequent educational choices and investments, ultimately impacting earnings [Denning et al., 2018].

Mental Health and Well-Being

As discussed, the importance of an assessment for future outcomes can motivate a student to study harder and learn more; it may also have negative consequences on the well-being of the individual. Using data on children aged eight to 15 in New Orleans, Heissel et al. [2021] studied the link between high-stakes testing and stress levels. The authors documented 18 percent higher levels of saliva-based measures of cortisol, a measure of the stress hormone capturing how the body's stress system is functioning, in weeks with high-stakes testing (before taking the test) compared to the same time in non-testing weeks. The effect was mainly driven by boys, and while the sample was relatively disadvantaged, they also found suggestive evidence that the response was largest for children living in poorer neighbourhoods. Importantly, Heissel et al. also showed that large responses in stress levels were associated with worse performance.

In Germany, children are assigned to an academic or non-academic track around the age of ten to 11. Bach and Fischer [2020] studied how well children do in schools where the track recommendation is linked to the performance on assessments compared to the children in the schools where

the track recommendation were more vague. In line with evidence that the stakes of the assessment incentivize student effort [Sievertsen et al., 2016], Bach and Fischer discovered that students in the schools where track choice is tightly linked to assessments show stronger growth in assessments. However, they also found that these children had worse mental health and were more likely to be stressed. These results point to an important downside of high-stakes testing. In Beuchert et al. [2020], the authors also studied the effect on well-being. While the data was imperfect for this test, the authors found no evidence that scoring just below a cutoff affects mental health negatively. Taking a macro-perspective and exploiting variation in the exposure to national high-stakes testing across countries and across time, Högberg and Horn [2022] find that being exposed to more high-stakes testing is associated with more self-reported school stress.

It is worth highlighting that it is more difficult to quantify the impact on mental health than on learning, because the latter, almost by definition, is quantified and measured automatically, and the latter is rarely captured and depends on the availability of survey data. An important area for future research would be the persistence of effects of testing on stress and mental health.

Long-Term Outcomes

Access to Degrees

Assessments in school and education may affect long-term outcomes of the students through several channels. Firstly, assessment results are used to sort students into subsequent educational programs, and exogenous factors and the design of these assessments will, therefore, affect the educational programs in which the students can enroll. Secondly, assessment results might be used by employers to screen applicants and, as such, will affect the match on the labor market, their earnings, and their employment chances. This section presents the evidence on these two factors. Thirdly, assessments provide students with information about the returns on studying and about their skills. Students may respond to this information by selecting specific tracks and increase or decrease their study effort.

As noted previously, evidence on teacher manipulation shows that students who benefited from the manipulation in the sense that their scores were improved because it also (in most cases) had better later life educational outcomes [Diamond and Persson, 2016, Dee et al., 2019]. Likewise,

as discussed in the case of more lenient grading standards, the higher grade given was shown to benefit students at the margin of passing who obtained access to degrees they otherwise would not have had access to [Maurin and McNally, 2008]. One direct explanation for this is that being pushed above a threshold gives access to more educational programs. This mechanism was confirmed by a study on national exams for all students at the end of compulsory schooling in England [Machin et al., 2020]. Comparing the subsequent educational trajectories, students in England who barely failed to receive a “good pass” in English in these high-stakes exams were less likely to enroll in upper-secondary academic and vocational tracks and start tertiary education, three years after the exams.

Prior grades and admission exams are often used to sort students into subsequent tracks and educational programs. For example, in the United States, the ACT or SAT admission exams are required for admission to most four-year colleges. These admission tests were historically taken outside school testing by students interested in attending college. This gives rise to several sources of inequality. Firstly, students and their parents need to know about this system and should also know how these tests work in order to do well. Secondly, direct costs are involved in terms of fees paid to the test centers. Thirdly, indirect costs are also involved in terms of travel to the test center outside school days. Hyman [2017] studied the consequences of these admission tests for access to college by exploiting policy variation through the introduction of mandatory ACT testing in Michigan in 2007. Following this policy, the testing was moved the school day at no costs to the students or the schools. This policy should, thus, alleviate all three concerns about inequality: information, direct and indirect costs. The overall response to the policy was a two percent increase in college enrollment, with substantially larger effects in schools with high poverty shares. Importantly, Hyman found no evidence that the marginal student pushed into college through this policy is likely to graduate from college. A take-away figure from the study is that for every ten poor students who get at least a college ready score on the ACT prior to the policy, another five would have scored above that threshold but did not attend the test before the reform. Studying the introduction of mandatory SAT testing in Maine, Hurwitz et al. [2015] revealed similar overall results to Hyman [2017] in a setting where the tests took place outside the school schedule but at no cost and with transport support.

Further evidence that financial costs constitute a barrier for low-income students is provided by

Goodman et al. [2020], who found that low-income students were less likely to retake the SAT university admissions test, but that low-income students who have waivers for test fees have higher retake test rates than their low-income peers without waivers. In this study, the authors observed that students were more likely to retake the SAT test if they were close to a score that was a multiple of 100. In other words, students were significantly more likely to retake if they scored 2099 than if scored 2100. Importantly, retaking the test caused higher scores, and lower retake rates among low-income students may, therefore, be another source of inequality, especially because Goodman et al. [2020] showed that retaking was especially beneficial for disadvantaged students.

Job-Market Signaling

Under the assumption that individuals with higher labor market productivity also have lower costs from completing educational degrees, the job-market signaling model suggests that employers use years of completed education to screen workers [Arrow, 1973, Spence, 1978]. The key insight is that workers who have completed more education will receive a higher wage, not necessarily because the education made them more productive but simply because the education allows employers to extract information about the individuals' innate productivity. It is straightforward to extend the job-market signaling model from level of education to educational credits in terms of how well the students did in school. If students with higher labor market productivity are assumed to also have lower costs from doing well on assessments, the employers can use assessment results to screen workers.

Grades may not only affect individuals' choices through information about their own abilities but also by giving them incentives to attend courses that are more leniently assessed if they care about their overall grade point average, as discussed in the section on grade inflation [Bar et al., 2009, Ahn et al., 2019]. To test whether employers use assessment to sort workers, Piopiunik et al. [2020] conducted an experiment where they sent fictitious job applications to German employers, with randomly varied elements on the applicants' curriculum vitae, including their grade point average. The researchers reported that a higher grade point average increased the likelihood of a job interview, and that the effect was stronger for college graduates than for high school graduates. They uncovered no differences by applicant gender, but they found that larger firms cared more about the grades than smaller firms and that older human resource managers cared less about

grades. The findings by Piopiunik et al. [2020] aligns with earlier evidence from Koedel and Tyhurst [2012] who conducted a similar experiment in the United States, that showed that a signal of high mathematics skills on the curriculum vitae increased the likelihood of a positive response by the employer. In a similar study conducted in Germany, Protsch and Solga [2015] documented evidence that employers use certain grade point average thresholds in their recruitment decisions.

Looking beyond the experimental variation in fictitious applicants and studying actual job market outcomes, Clark and Martorell [2014] used data on high school students in Texas and compared labor market outcomes of students who barely passed the high school exit exam to those who barely failed the exam. Both groups of students received a certificate, the former a certificate for obtaining the high school degree, and the latter a certificate for completing high school. As students very close to the cutoff should be essentially identical in their skills, any difference in labor market outcomes should reflect the signaling value of actually passing the degree. However, the authors found no evidence of this signal on subsequent earnings. While this finding goes against the findings of the fictitious application studies, it is worth highlighting that Piopiunik et al. [2020] found that the effects were considerably smaller among high school applicants than college applicants and that Protsch and Solga [2015] found suggestive evidence that employers use specific cutoffs in the grade point averages to screen potential workers. Using a design similar to that of Clark and Martorell [2014], Jepsen et al. [2016] compared labor market outcomes for individuals barely passing the General Educational Development (GED) certification to the outcomes of individuals who barely failed it. They found no effects on subsequent employment for women but larger earnings in year two after earning the certificate.

Another possible reason for why the two studies produced weak evidence of the signaling value of grades and passing exams, while the survey experiments suggested large labor market effects, is that these two studies focused on local threshold effects for a rather particular degree type. Looking across the ability distribution, Feng and Graetz [2017] finds sizable returns to receiving a first (vs upper second) degree and to receiving an upper second (vs lower second) degree in the British university setting. Hansen et al. [2021] used a grading reform to study the effect of as-good-as random variation in grade point averages for graduates from the two largest universities in Denmark. The research exploited the same recoding as in Hvidman and Sievertsen [2021] which has the advantage that it affects the entire grade distribution. They found that an exogenous

increase in grade point average caused significantly higher earnings in the first years after graduation, but that the effect disappeared within two to four years. This suggests that that employers initially use the grades to screen potential workers, but that they quickly learn about the workers' skills.

Taking Stock

This section summarizes the evidence covered in this paper related to the impact assessments have on efficiency, equity, and well-being but first presents a discussion on the impact of assessments on schools and beyond.

Effects on Schools

Assessments may have a direct impact on the school's finances and on individual teachers' pay and employment when directly linked to accountability incentives like in case of the NCLB policy in the United States [Jacob, 2005]. Schools might also be indirectly affected by assessments, if assessment results are publicly available and parents and students use this information to decide what school to attend [Diamond and Persson, 2016, Nordin et al., 2019]. As discussed, such incentives might affect schools' and teachers' behavior such as teaching to the test placing students on special needs programs [Jacob, 2005], teaching to the rating [Reback, 2008], teacher test score manipulation [Jacob and Levitt, 2003, Diamond and Persson, 2016, Dee et al., 2019], and lenient grading policies [Nordin et al., 2019].

There is evidence that school's prioritization of grades based on the expectation that they matter for parents is justified. In the school choice literature, the evidence that parents value school quality is solid [Black, 1999, Imberman and Lovenheim, 2016, Burgess et al., 2015, Borghans et al., 2015, Harjunen et al., 2018]. School average test scores are a product of both the quality of schools and the student composition, whereas value added is meant to capture only the school quality. A school that prioritizes doing well on assessment either through improved teaching or strategic responses (e.g., teaching to the test or rating) will improve both raw test scores and value added. To mention a few examples, one strategy to assess the degree to which parents care about test scores, is to study the revealed preferences in terms of their willingness to pay. Looking at house

prices in Massachusetts around school district boundaries, Black [1999] concluded that a five percent increase in school test scores caused an increase in the willingness to pay for housing by 2.1 percent. The boundary idea is based on the premise that houses close to the school district boundary should be essentially identical in terms of all other amenities except for the school, which is allocated based on the district. However, similar effects have also been found in Finnish settings where the assessment results were not publicly available [Harjunen et al., 2018]. Imberman and Lovenheim [2016] used data from Los Angeles and confirmed the finding that assessment results affects house prices. Value added measures, on the other hand, do not, which suggest that parents care about test scores, but maybe less about the origin of test score improvements. There is some evidence for heterogeneity in how parents value school quality. Studying school choice in England, Burgess et al. [2015] showed that while most parents had preferences for academic performance of the schools, some socioeconomic gradient existed in these preferences. For example, the preference for academic quality was stronger for high socioeconomic parents. However, this may also be due to constraints, either actual access constraints or information constraints. In sum, there is clear evidence that a school's assessment result affect parents' demand and that they are willing to pay for a school that receives better results. Less clear is if the parents care about whether the good assessment comes directly from school quality, or through better students and, thus, better peers.

Assessments and Inequality

To summarize the research evidence covered with respect to how assessments affect inequality, it is useful to separate the mechanisms into four channels. Firstly, there is evidence that some groups suffer from biases in assessments [Lavy, 2008, Burgess and Greaves, 2013, Falch and Naper, 2013, Rangvid, 2015, Terrier, 2020]. Such biases may have longer-run effects on students educational outcomes, because students use assessments to update their beliefs about returns to studying and their abilities [Stinebrickner and Stinebrickner, 2012], which is also directly documented in various settings [Terrier, 2020, Murphy and Wyness, 2020, Burgess et al., 2022]. Thirdly, heterogeneity in access to and response to assessments affect test taking behavior and subsequent access to education Hyman [2017], Holm et al. [2019], Andersen and Nielsen [2020], Goodman et al. [2020], Birkelund [2020]. Fourthly, there is evidence of socioeconomic differences in how students are affected by and exposed to external shocks in test taking [Ebenstein et al., 2016]

Assessment and Efficiency

The research evidence covered also clearly shows that assessments affect efficiency in terms of human capital decisions. The decision of the optimal level (and later type of) schooling has been considered an investment at least since some of the early models on optimal human capital decisions [Schultz, 1961, Becker, 1964]. In these models, schooling raises productivity and we invest in it until marginal returns equal the marginal costs. However, as the evidence shows, individuals' educational decisions are affected by the leniency in the assessments [Bar et al., 2009], likely because employers use the grades as a screening device [Piopiunik et al., 2020, Clark and Martorell, 2014, Protsch and Solga, 2015, Jepsen et al., 2016, Hansen et al., 2021]. Moreover, students react to the high-stakes incentives in assessments by increasing their effort [Sievertsen et al., 2016], and such strategic responses are actually also observed by schools [Reback, 2008] that focus on the students who count more in the accountability system. In sum, because assessments are used as a signaling device in accountability systems, students' and schools' investment in education might differ from an optimal decision in terms of maximizing learning and, thus, human capital accumulation.

Mental Health and Well-Being

It is clear from the evidence discussed, that our knowledge about how testing affects mental health and well-being is relatively limited. One explanation for this is that the outcome, mental health and well-being, is both harder to measure and less frequently measured automatically, compared to, for example, learning and later life outcomes. However, the existing evidence points to effects on stress [Heissel et al., 2021], and also on mental health in very young ages [Bach and Fischer, 2020]. Given the strong evidence on biases on how assessments affect incentives to study hard, drop-out decisions, course choices, access to subsequent educational programs, and labor market outcomes, that there may effects on mental health and well-being is not surprising. However, much more research is needed to answer questions on how to balance the benefits of standardized testing for equity [Hyman, 2017] and potential detrimental effects on well-being.

Key references

- Annika B Bergbauer, Eric A Hanushek, and Ludger Woessmann. Testing. *Journal of Human Resources*, pages 0520–10886R1, 2021
- Lex Borghans, Bart HH Golsteyn, James J Heckman, and John Eric Humphries. What grades and achievement tests measure. *Proceedings of the National Academy of Sciences*, 113(47):13354–13359, 2016
- Simon Burgess and Ellen Greaves. Test scores, subjective assessment, and stereotyping of ethnic minorities. *Journal of Labor Economics*, 31(3):535–576, 2013
- Jeffrey T Denning, Eric R Eide, Kevin J Mumford, Richard W Patterson, and Merrill Warnick. Why have college completion rates increased? *American Economic Journal: Applied Economics*, 14(3):1–29, 2022
- David N Figlio and Maurice E Lucas. Do high grading standards affect student performance? *Journal of Public Economics*, 88(9-10):1815–1834, 2004
- Joshua Hyman. Act for all: The effect of mandatory college entrance exams on postsecondary attainment and choice. *Education Finance and Policy*, 12(3):281–311, 2017
- C Kirabo Jackson. What do test scores miss? the importance of teacher effects on non-test score outcomes. *Journal of Political Economy*, 126(5):2072–2107, 2018
- Rema N Hanna and Leigh L Linden. Discrimination in grading. *American Economic Journal: Economic Policy*, 4(4):146–68, 2012
- Nathan Petek and Nolan G. Pope. The multidimensional impact of teachers on student outcomes. forthcoming
- Todd Stinebrickner and Ralph Stinebrickner. Learning about academic ability and the college dropout decision. *Journal of Labor Economics*, 30(4):707–748, 2012

Further References

- Thomas Ahn, Peter Arcidiacono, Amy Hopson, and James R Thomas. Equilibrium grade inflation with implications for female interest in stem majors. Technical report, National Bureau of Economic Research, 2019.
- Mathilde Almlund, Angela Lee Duckworth, James Heckman, and Tim Kautz. Personality psychology and economics. In *Handbook of the Economics of Education*, volume 4, pages 1–181. Elsevier, 2011.
- Lina Anaya, Frank Stafford, and Gema Zamarro. Gender gaps in math performance, perceived mathematical ability and college stem education: The role of parental occupation. *Education Economics*, 30(2):113–128, 2022.
- Simon Calmar Andersen and Helena Skyt Nielsen. Learning from performance information. *Journal of Public Administration Research and Theory*, 30(3):415–431, 2020.
- Megan Andrew and Robert M Hauser. Adoption? adaptation? evaluating the formation of educational expectations. *Social Forces*, 90(2):497–520, 2011.
- Kenneth J Arrow. Higher education as a filter. *Journal of Public Economics*, 2(3):193–216, 1973.
- Ghazala Azmat, Caterina Calsamiglia, and Nagore Iriberry. Gender differences in response to big stakes. *Journal of the European Economic Association*, 14(6):1372–1400, 2016.
- Ghazala Azmat, Manuel Bagues, Antonio Cabrales, and Nagore Iriberry. What you don’t know... can’t hurt you? a natural field experiment on relative performance feedback in higher education. *Management Science*, 65(8):3714–3736, 2019.
- Philip Babcock. Real costs of nominal grade inflation? new evidence from student course evaluations. *Economic inquiry*, 48(4):983–996, 2010.
- Maximilian Bach and Mira Fischer. Understanding the response to high-stakes incentives in primary education. *ZEW-Centre for European Economic Research Discussion Paper*, (20-066), 2020.

- Oriana Bandiera, Valentino Larcinese, and Imran Rasul. Blissful ignorance? a natural experiment on the effect of feedback on students' performance. *Labour Economics*, 34:13–25, 2015.
- Talia Bar, Vrinda Kadiyali, and Asaf Zussman. Grade information and grade inflation: The cornell experiment. *Journal of Economic Perspectives*, 23(3):93–108, 2009.
- Gary S Becker. Human capital: A theoretical and empirical analysis, with special reference to education. 1964.
- Annika B Bergbauer, Eric A Hanushek, and Ludger Woessmann. Testing. *Journal of Human Resources*, pages 0520–10886R1, 2021.
- Julian R Betts and Jeff Grogger. The impact of grading standards on student achievement, educational attainment, and entry-level earnings. *Economics of Education Review*, 22(4):343–352, 2003.
- Louise Beuchert, Tine Louise Mundbjerg Eriksen, and Morten Visby Krægpøth. The impact of standardized test feedback in math: Exploiting a natural experiment in 3rd grade. *Economics of Education Review*, 77:102017, 2020.
- Jesper Fels Birkelund. Aiming high and missing the mark? educational choice, dropout risk, and achievement in upper secondary education among children of immigrants in denmark. *European Sociological Review*, 36(3):395–412, 2020.
- Aspasia Bizopoulou, Rigissa Megalokonomou, and Stefania Simion. Do second chances pay off? evidence from a natural experiment with low-achieving students. 2022.
- Sandra E Black. Do better schools matter? parental valuation of elementary education. *The quarterly journal of economics*, 114(2):577–599, 1999.
- Raphael Boleslavsky and Christopher Cotton. Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–79, 2015.
- Timothy N Bond, Jillian B Carr, Analisa Packham, and Jonathan Smith. Hungry for success? snap timing, high-stakes exam performance, and college attendance. *American Economic Journal: Economy Policy*, forthcoming.

- Lex Borghans, Bart HH Golsteyn, and Ulf Zölitz. Parental preferences for primary school characteristics. *The BE Journal of Economic Analysis & Policy*, 15(1):85–117, 2015.
- Lex Borghans, Bart HH Golsteyn, James J Heckman, and John Eric Humphries. What grades and achievement tests measure. *Proceedings of the National Academy of Sciences*, 113(47):13354–13359, 2016.
- Thomas Breda and Son Thierry Ly. Professors in core science fields are not always biased against women: Evidence from france. *American Economic Journal: Applied Economics*, 7(4):53–75, 2015.
- Jack Buckley, Lynn Letukas, and Ben Wildavsky. Measuring success: Testing, grades, and the future of college admissions. *Johns Hopkins University Press*, 2018.
- Simon Burgess and Ellen Greaves. Test scores, subjective assessment, and stereotyping of ethnic minorities. *Journal of Labor Economics*, 31(3):535–576, 2013.
- Simon Burgess and Ellen Greaves. School choice and accountability. In *Oxford Research Encyclopedia of Economics and Finance*. 2021.
- Simon Burgess, Ellen Greaves, Anna Vignoles, and Deborah Wilson. What parents want: School preferences and school choice. *The Economic Journal*, 125(587):1262–1289, 2015.
- Simon Burgess, Daniel Sloth Hauberg, Beatrice Schindler Rangvid, and Hans Henrik Sievertsen. The importance of external assessments: High school math and gender gaps in stem degrees. *Economics of Education Review*, 88:102267, 2022.
- William Chan, Li Hao, and Wing Suen. A signaling theory of grade inflation. *International Economic Review*, 48(3):1065–1090, 2007.
- Damon Clark and Paco Martorell. The signaling value of a high school diploma. *Journal of Political Economy*, 122(2):282–318, 2014.
- Thomas S Dee. A teacher like me: Does race, ethnicity, or gender matter? *American Economic Review*, 95(2):158–165, 2005.

- Thomas S Dee, Will Dobbie, Brian A Jacob, and Jonah Rockoff. The causes and consequences of test score manipulation: Evidence from the new york regents examinations. *American Economic Journal: Applied Economics*, 11(3):382–423, 2019.
- Jeffrey T Denning, Richard Murphy, and Felix Weinhardt. Class rank and long-run outcomes. *The Review of Economics and Statistics*, pages 1–45, 2018.
- Jeffrey T Denning, Eric R Eide, Kevin J Mumford, Richard W Patterson, and Merrill Warnick. Why have college completion rates increased? *American Economic Journal: Applied Economics*, 14(3): 1–29, 2022.
- Rebecca Diamond and Petra Persson. The long-term consequences of teacher discretion in grading of high-stakes tests. *National Bureau of Economic Research - Working Paper*, (22207), 2016.
- Avraham Ebenstein, Victor Lavy, and Sefi Roth. The long-run economic consequences of high-stakes examinations: Evidence from transitory variation in pollution. *American Economic Journal: Applied Economics*, 8(4):36–65, 2016.
- Tim Ehlers and Robert Schwager. Honest grading, grade inflation, and reputation. *CESifo Economic Studies*, 62(3):506–521, 2016.
- Benjamin Elsner and Ingo E Isphording. A big fish in a small pond: Ability rank and human capital investment. *Journal of Labor Economics*, 35(3):787–828, 2017.
- Benjamin Elsner and Ingo E Isphording. Rank, sex, drugs, and crime. *Journal of Human Resources*, 53(2):356–381, 2018.
- Benjamin Elsner, Ingo E Isphording, and Ulf Zölitz. Achievement rank affects performance and major choices in college. *The Economic Journal*, 131(640):3182–3206, 2021.
- Torberg Falch and Linn Renée Naper. Educational evaluation schemes and gender gaps in student achievement. *Economics of Education Review*, 36:12–25, 2013.
- Andy Feng and Georg Graetz. A question of degree: the effects of degree class on labor market outcomes. *Economics of Education Review*, 61:140–161, 2017.

- David N Figlio and Maurice E Lucas. Do high grading standards affect student performance? *Journal of Public Economics*, 88(9-10):1815–1834, 2004.
- Peter Fredriksson, Lena Hensvik, and Oskar Nordström Skans. Mismatch of talent: Evidence on match quality, entry wages, and job mobility. *American Economic Review*, 108(11):3303–38, 2018.
- Dillon Fuchsman, Tim R Sass, and Gema Zamarro. Testing, teacher turnover and the distribution of teachers across grades and schools. *Education Finance and Policy*, pages 1–61, 2020.
- Uri Gneezy, Muriel Niederle, and Aldo Rustichini. Performance in competitive environments: Gender differences. *Quarterly Journal of Economics*, 118(3):1049–1074, 2003.
- Uri Gneezy, John A List, Jeffrey A Livingston, Xiangdong Qin, Sally Sadoff, and Yang Xu. Measuring success in education: the role of effort on the test itself. *American Economic Review: Insights*, 1(3):291–308, 2019.
- Claudia Goldin and Cecilia Rouse. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American economic review*, 90(4):715–741, 2000.
- Joshua Goodman, Oded Gurantz, and Jonathan Smith. Take two! sat retaking and college enrollment gaps. *American Economic Journal: Economic Policy*, 12(2):115–58, 2020.
- Joshua Graff Zivin and Matthew Neidell. The impact of pollution on worker productivity. *American Economic Review*, 102(7):3652–73, 2012.
- Darren Grant and William B Green. Grades as incentives. *Empirical Economics*, 44(3):1563–1592, 2013.
- Thomas R Guskey. Using assessments to improve teaching and learning. *Ahead of the curve: The power of assessment to transform teaching and learning*, pages 15–29, 2007.
- Rema N Hanna and Leigh L Linden. Discrimination in grading. *American Economic Journal: Economic Policy*, 4(4):146–68, 2012.
- Anne Toft Hansen, Ulrik Hvidman, and Hans Henrik Sievertsen. Grades and employer learning. Technical report, IZA Discussion Papers, 2021.

- Oskari Harjunen, Mika Kortelainen, and Tuukka Saarimaa. Best education money can buy? capitalization of school quality in finland. *CESifo Economic Studies*, 64(2):150–175, 2018.
- JJ Heckman, C Heinrich, and J Smith. The performance of performance standards. *Journal of Human Resources*, 37(4):778–811, 2002.
- Jennifer A Heissel, Emma K Adam, Jennifer L Doleac, David N Figlio, and Jonathan Meer. Testing, stress, and performance: How students respond physiologically to high-stakes testing. *Education Finance and Policy*, 16(2):183–208, 2021.
- Rey Hernandez-Julian and Adam Looney. Measuring inflation in grades: An application of price indexing to undergraduate grades. *Economics of Education Review*, 55:220–232, 2016.
- Björn Högberg and Daniel Horn. National high-stakes testing, gender, and school stress in europe: a difference-in-differences analysis. *European Sociological Review*, 2022.
- Anders Holm, Anders Hjorth-Trolle, and Mads Meier Jæger. Signals, educational decision-making, and inequality. *European Sociological Review*, 35(4):447–460, 2019.
- Bengt Holmstrom and Paul Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.*, 7:24, 1991.
- Michael Hurwitz, Jonathan Smith, Sunny Niu, and Jessica Howell. The maine question: How is 4-year college enrollment affected by mandatory college entrance exams? *Educational Evaluation and Policy Analysis*, 37(1):138–159, 2015.
- Ulrik Hvidman and Hans Henrik Sievertsen. High-stakes grades and student behavior. *Journal of Human Resources*, 56(3):821–849, 2021.
- Joshua Hyman. Act for all: The effect of mandatory college entrance exams on postsecondary attainment and choice. *Education Finance and Policy*, 12(3):281–311, 2017.
- Scott A Imberman and Michael F Lovenheim. Does the market value value-added? evidence from housing prices after a public release of school and teacher value-added. *Journal of Urban Economics*, 91:104–121, 2016.

- C Kirabo Jackson. What do test scores miss? the importance of teacher effects on non-test score outcomes. *Journal of Political Economy*, 126(5):2072–2107, 2018.
- Brian A Jacob. Accountability, incentives and behavior: The impact of high-stakes testing in the chicago public schools. *Journal of Public Economics*, 89(5-6):761–796, 2005.
- Brian A Jacob and Steven D Levitt. Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *The Quarterly Journal of Economics*, 118(3):843–877, 2003.
- Christopher Jepsen, Peter Mueser, and Kenneth Troske. Labor market returns to the ged using regression discontinuity analysis. *Journal of Political Economy*, 124(3):621–649, 2016.
- Thomas J Kane and Douglas O Staiger. The promise and pitfalls of using imprecise school accountability measures. *Journal of Economic perspectives*, 16(4):91–114, 2002.
- Kristian Bernt Karlson. Expectation formation for all? group differences in student response to signals about academic performance. *The Sociological Quarterly*, 60(4):716–737, 2019.
- Cory Koedel and Eric Tyhurst. Math skills and labor-market outcomes: Evidence from a resume-based field experiment. *Economics of Education Review*, 31(1):131–140, 2012.
- Steffen Künn, Juan Palacios, and Nico Pestel. Indoor air quality and cognitive performance. *Management Science*, forthcoming.
- Victor Lavy. Do gender stereotypes reduce girls’ or boys’ human capital outcomes? evidence from a natural experiment. *Journal of Public Economics*, 92(10):2083–2105, 2008. ISSN 0047-2727. doi: <https://doi.org/10.1016/j.jpubeco.2008.02.009>. URL <https://www.sciencedirect.com/science/article/pii/S0047272708000418>.
- Victor Lavy and Rigissa Megalokonomou. Persistency in teachers’ grading bias and effects on longer-term outcomes: University admissions exams and choice of field of study. Technical report, National Bureau of Economic Research, 2019.
- Edward P Lazear. Speeding, terrorism, and teaching to the test. *The Quarterly Journal of Economics*, 121(3):1029–1061, 2006.
- Ming-Jen Lin and Steven D Levitt. Catching cheating students. *Economica*, 87(348):885–900, 2020.

- Erica Lindahl. Are teacher assessments biased?—evidence from sweden. *Education economics*, 24(2): 224–238, 2016.
- Stephen Machin, Sandra McNally, and Jenifer Ruiz-Valenzuela. Entry through the narrow door: The costs of just failing high stakes exams. *Journal of Public Economics*, 190:104224, 2020.
- Eric Maurin and Sandra McNally. Vive la revolution! long-term educational returns of 1968 to the angry students. *Journal of Labor Economics*, 26(1):1–33, 2008.
- Robert Metcalfe, Simon Burgess, and Steven Proud. Students’ effort and educational achievement: Using the timing of the world cup to vary the value of leisure. *Journal of Public Economics*, 172: 111–126, 2019.
- Richard Murphy and Felix Weinhardt. Top of the class: The importance of ordinal rank. *The Review of Economic Studies*, 87(6):2777–2826, 2020.
- Richard Murphy and Gill Wyness. Minority report: the impact of predicted grades on university admissions of disadvantaged groups. *Education Economics*, 28(4):333–350, 2020.
- Derek Neal. The consequences of using one assessment system to pursue two objectives. *The Journal of Economic Education*, 44(4):339–352, 2013.
- Martin Nordin, Gawain Heckley, and Ulf Gerdtham. The impact of grade inflation on higher education enrolment and earnings. *Economics of Education Review*, 73:101936, 2019.
- OFS. Analysis of degree classifications over time - changes in graduate attainment from 2010-11 to 2020-21. Technical report, Office for Students, 2022.
- Nathan Petek and Nolan G. Pope. The multidimensional impact of teachers on student outcomes. forthcoming.
- Edmund S Phelps. The statistical theory of racism and sexism. *American Economic Review*, 62(4): 659–661, 1972.
- Marc Piopiunik, Guido Schwerdt, Lisa Simon, and Ludger Woessmann. Skills, signals, and employability: An experimental investigation. *European Economic Review*, 123:103374, 2020.

- Paula Protsch and Heike Solga. How employers use signals of cognitive and noncognitive skills at labour market entry: insights from field experiments. *European Sociological Review*, 31(5):521–532, 2015.
- Beatrice Schindler Rangvid. Systematic differences across evaluation schemes and educational choice. *Economics of Education Review*, 48:41–55, 2015.
- Beatrice Schindler Rangvid. Gender discrimination in exam grading? double evidence from a natural experiment and a field experiment. *The BE Journal of Economic Analysis & Policy*, 19(2), 2019.
- Randall Reback. Teaching to the rating: School accountability and the distribution of student achievement. *Journal of Public Economics*, 92(5-6):1394–1415, 2008.
- Richard Sabot and John Wakeman-Linn. Grade inflation and course choice. *Journal of Economic Perspectives*, 5(1):159–170, 1991.
- Paul R Sackett and Nathan R Kuncel. Eight myths about standardized admissions testing. *Measuring success: Testing, grades, and the future of college admissions*, pages 13–39, 2018.
- Theodore W Schultz. Investment in human capital. *The American economic review*, 51(1):1–17, 1961.
- Robert Schwager. Grade inflation, social background, and labour market matching. *Journal of Economic Behavior & Organization*, 82(1):56–66, 2012.
- Hans Henrik Sievertsen, Francesca Gino, and Marco Piovesan. Cognitive fatigue influences students’ performance on standardized tests. *Proceedings of the National Academy of Sciences*, 113(10): 2621–2624, 2016.
- Mary Lovett Smallwood. *An historical study of examinations and grading systems in early American universities*. Harvard University Press, 1969.
- Michael Spence. Job market signaling. In *Uncertainty in economics*, pages 281–306. Elsevier, 1978.
- Todd Stinebrickner and Ralph Stinebrickner. Learning about academic ability and the college dropout decision. *Journal of Labor Economics*, 30(4):707–748, 2012.

Todd R Stinebrickner and Ralph Stinebrickner. Math or science? using longitudinal expectations data to examine the process of choosing a college major. Technical report, National Bureau of Economic Research, 2011.

Camille Terrier. Boys lag behind: How teachers' gender biases affect student achievement. *Economics of Education Review*, 77:101981, 2020.

The Guardian. Letter by academics - oecd and pisa tests are damaging education worldwide-academics. *Guardian Education*, 2014.

The World Bank. *World Development Report -Learning to realize education's promise*. World Bank Group, 2018.

Christina Wikström and Magnus Wikström. Grade inflation and school competition: an empirical analysis based on the swedish upper secondary schools. *Economics of education Review*, 24(3): 309–322, 2005.

Robertas Zubrickas. Optimal grading. *International Economic Review*, 56(3):751–776, 2015.