Comments and Discussion

COMMENT BY

ERIK HURST This paper documents a very interesting set of facts. In particular, the paper shows spatial variation in the extent to which adult residents have a deficit in what Hoxby defines as "advanced cognitive skills." The paper then speculates that advanced cognitive skill deserts may arise because of differential local investments in these skills during adolescence. Overall, I expect that the findings in this paper will stimulate a large amount of future research.

The paper is mostly descriptive. As a result, I do not have many substantive comments on the paper's message. However, three things entered my mind as I was reading it. First, I wondered about other ways to measure advanced cognitive skills within a local area. Second, I wondered whether information on within-county variation in cognitive skills would complement the paper's cross-county analysis. Finally, I wondered whether other correlates may be useful to readers with respect to understanding the causes of advanced cognitive skill deserts. I expand on each of these comments below.

HOW TO MEASURE COGNITIVE SKILL DESERTS In the paper, advanced cognitive skills are defined broadly as those skills needed to perform higher order reasoning. In particular, the paper refers to advanced cognitive skills as those that "require a capacity to solve problems through logic, think in the abstract, engage in critical thinking, and derive general principles from a set of facts." To create maps of cognitive skill deserts, the paper uses data from the Program for the International Assessment of Adult Competencies (PIAAC). The PIAAC is a large survey that assesses respondents' numeracy and literacy skills. Respondents are binned into six levels of skills after taking each test. The paper defines respondents as having advanced cognitive skills if they are classified in the three highest levels of skill on the numeracy portion of the PIAAC. By this definition, about 37 percent of the US adult population is classified as having advanced cognitive skills. Figure 3 of the paper shows the share of respondents in each county who have advanced cognitive skills by this metric. The cognitive skill deserts documented in the paper using this metric are concentrated almost exclusively in the South census region. College towns throughout the South show higher levels of advanced cognitive skills. But, for the most part, the rest of the South—including much of Appalachia—has more counties with low levels of advanced cognitive skills, particularly compared to other regions.

The US Census puts out maps showing variation across counties in the share of residents over the age of 25 with a bachelor's degree or higher (McElrath and Martin 2021). In my discussion, I showed one of these maps (figure 1). The Census Bureau creates this map using individual-level data pooling together the 2015–2019 waves of the American Community Survey. The patterns of the spatial variation in the share of residents with at least a bachelor's degree is nearly identical to the spatial variation in the advanced cognitive skills as measured by the PIAAC documented in the paper. In particular, in the United States there is a "bachelor's degree desert" in the South. Within the South, there are pockets of counties with a higher share of bachelor's degrees. These counties often include college towns and are the same counties in the South that have a larger amount of residents with advanced cognitive skills. The similarity in spatial patterns across the two measures begs the question of whether cognitive skill measures from the PIAAC are just proxying for lower levels of accumulated schooling. It would have been nice to have a plot in the paper correlating county-level share of high cognitive skills using the PIAAC numeracy measures with the county-level share of residents with a bachelor's degree. Are the numeracy measures proxying for something above and beyond low levels of accumulated education? Are the cognitive skill deserts highlighted in the paper simply places where education levels of adult residents are low? What are high scores on numeracy exams from the PIAAC measuring that is distinct from obtaining a bachelor's degree? The paper is silent on these questions. Going forward, it may be useful to flush out whether cognitive skill deserts are something distinct relative to places with lower levels of accumulated schooling.

WITHIN-COUNTY VERSUS CROSS-COUNTY VARIATION IN ADVANCED COGNITIVE SKILLS My second comment on the paper centers on what is the correct level of aggregation to think about cognitive skill deserts. The paper uses variation across counties. In doing so, it suggests that certain counties may invest less in developing advanced cognitive skills relative to other counties.

Figure 1. Percentage of People 25 Years and Older with a Bachelor's Degree or Higher, 2015–2019

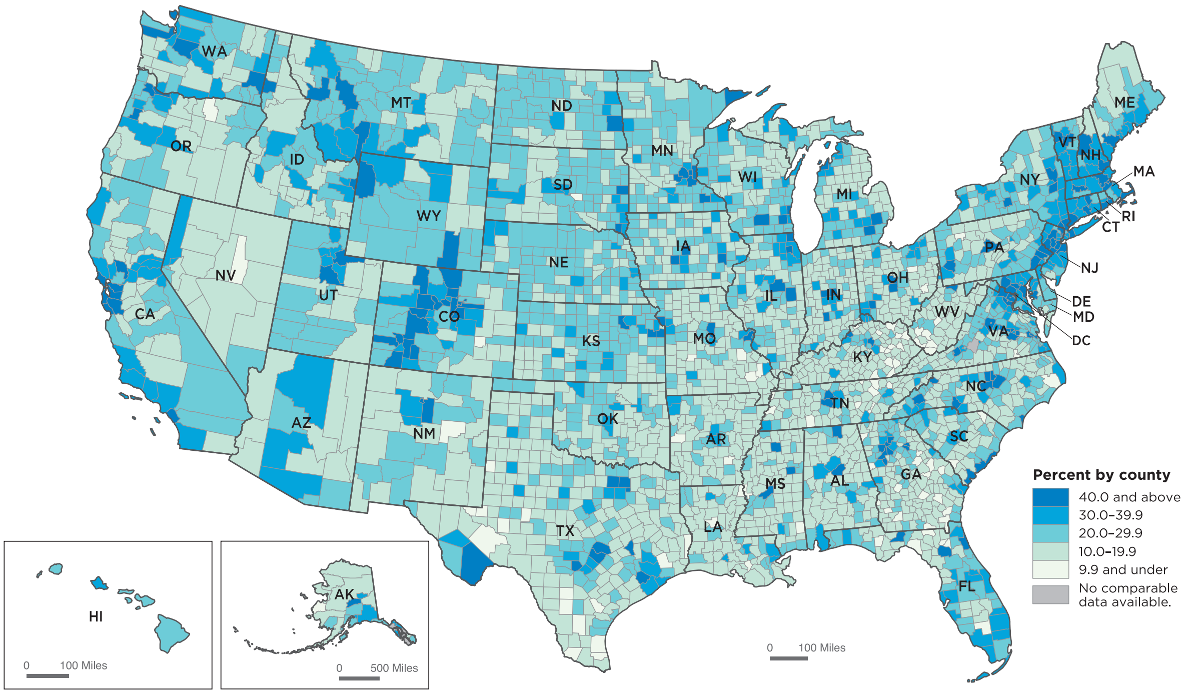

Sources: US Census Bureau, table S1501, 2010–2014 American Community Survey, five-year estimates; https://www.census.gov/content/dam/Census/library/publications/2021/acs/acsbr-009.pdf. Used with permission.

Notes: Due to the county boundary changes that occurred within the 2005–2009 and 2010–2019 time periods, the following county equivalents appear in the "No comparable data available" category: Petersburg, Alaska; Prince of Wales-Hyder, Alaska; Hoonah-Angoon, Alaska; Bedford, Virginia; and Bedford City, Virginia.

It would be interesting to think about within-county variation as well. Consider, for example, the city of Chicago. Chicago is comprised of dozens of different neigbhorhoods. I would conjecture that the variation across neighborhoods within Chicago with respect to the paper's measure of advanced cognitive skills is as large as the variation across counties. There is some evidence to back this up. The *Chicago Tribune* (2017) reports average student SAT scores by Chicago high school. The lowest average SAT math score from students in a Chicago high school was 360 while the highest average SAT math score was 686. The schools with lower SAT math scores were in geographically different areas within Chicago than the schools with higher scores. There are parts of Chicago that would look like they were in a cognitive skill desert (measured by SAT math scores) relative to other parts of Chicago.

Going forward, it may be useful to explore the extent to which within-county variation in measures of cognitive skills are useful in helping us learn about the causes of cross-county variation in measures of cognitive skills. There is large spatial variation in measures of schooling or test scores even within a large city like Chicago. Why is it more interesting to focus on cross-county differences in cognitive skills relative to focusing on within-county spatial differences? Future work can shed light on these issues.

WHAT EXPLAINS THE EXISTENCE OF COGNITIVE SKILL DESERTS? My third and final comment centers on potential explanations for the spatial variation in advanced cognitive skills. The paper focuses on a handful of potential explanations for the cognitive skill deserts. The first discusses early childhood factors, and the second focuses on influences during adolescence. The paper shows that advanced cognitive skills measures of adults (the PIAAC data) in a given county correlate strongly with test score measures of adolescents in that location. However, the adult measures of cognitive skills in a given location are only weakly correlated with test score measures of younger children. The paper then concludes that advanced cognitive skills are mostly engrained in adolescence. That conclusion rests on the extent to which test scores of young children actually measure well a child's cognitive skills. If test scores measure cognitive skills with error and that error is larger for younger children than for older childer, we would expect more spatial correlation between the test scores of adults and the test scores of teenagers than we would between the test scores of adults and the test scores of younger children.

Going forward, it would be useful to explore other demographic and socioeconomic correlates of spatial differences in measures of skills. For example, how does spatial variation in PIAAC cognitive skill measures vary with spatial differences in parental education, adult income, adult industry mix, and other demographic variables (such as race and ethnicity)? These correlations can help shed light on some of the mechanisms underlying the spatial variation in measures of skills.

REFERENCES FOR THE HURST COMMENT

Chicago Tribune. 2017. "SAT Scores in Illinois: See How Your High School Compares." November 9. https://www.chicagotribune.com/suburbs/ct-school-report-card-sat-scores-2017-htmlstory.html.

McElrath, K., and M. Martin. 2021. "Bachelor's Degree Attainment in the United States: 2005 to 2019." American Community Survey Brief ACSCR-009. Washington: US Department of Commerce, US Census Bureau. https://www.census.gov/content/dam/Census/library/publications/2021/acs/acsbr-009.pdf.

COMMENT BY

BRIAN A. JACOB In this paper Hoxby examines the variation in cognitive ability across geographic locations in the United States. She documents three important facts. First, adults with advanced cognitive skills are clustered disproportionately in certain places. Specifically, urban and coastal areas have a particularly high proportion of adults with advanced skills. Examples include northern New England (such as Boston), large metropolitan areas in California, and selected counties in the Midwest that Hoxby refers to as the "Lutheran Belt." Appalachia, the Ozarks, and areas of the inland South (Louisiana, Mississippi, Georgia) are "skill deserts," with very few adults possessing advanced cognitive skills.

Second, cognitive skills among children are distributed more evenly across geography compared with cognitive ability of adults. While children in urban and coastal areas and parts of the Midwest tend to outperform those in Appalachia and the Southeast, the differences are much less stark than in adults.

Third, there is a correlation between the geographic distribution of adult skills and the analogous distribution of child skills. Importantly, the magnitude of the correlation increases as children get older, particularly as they enter adolescence. That is, the correlation between adult skills and the achievement level of high school students is larger than the correlation between adult skills and elementary school achievement.

Hoxby argues that these facts, in combination with other evidence, have important implications. Throughout the paper, she emphasizes two related themes. One involves the salience of advanced cognitive skills. She first argues that skill-biased technological change and related economic forces have increased the importance of such skills in today's labor market. The second theme is the importance of adolescence as an "age of opportunity." Referencing brain science research indicating that adolescence is the time during which advanced cognitive skills develop, she suggests that it might be particularly beneficial to target educational interventions during adolescence. Based on the geographic skill distribution, policymakers should target skill deserts in particular for such interventions.

There is a lot to like in Hoxby's analysis. First, the attention on adolescence is a useful antidote to the policy community's intense focus on early childhood over the past two decades.[1] This is not to say that educators should avoid intervening early in children's lives, but rather that the intense focus on this time period risks neglecting effective strategies for older children. Second, the focus on geography is consistent with other recent work, such as the analysis of intergenerational mobility by Chetty and Hendren (2018a, 2018b) and Chetty, Hendren, and Katz (2016). In particular, Hoxby highlights the challenges faced by rural communities, which are sometimes neglected as policymakers have focused (understandably) on the struggles of those in urban areas.

In this comment, I seek to make several points. To begin, I raise some methodological issues that complicate Hoxby's analysis. Second, I provide some supplementary evidence to support Hoxby's contention that students in the United States struggle during adolescence. Finally, I discuss what evidence we have on potential interventions for adolescents in skill deserts.

THE CHALLENGE OF ASSESSING AND INTERPRETING COGNITIVE ABILITY I wholeheartedly agree with Hoxby's contention that cognitive skills are a more useful measure of an individual's capacity to function in contemporary labor markets than educational attainment. Unfortunately, assessing individual skills presents a number of challenges.[2]

First, standardized test scores are noisy measures of true ability, which fluctuate for many reasons, ranging from whether an individual was sick or distracted during the test to which particular items were asked on the assessment. Moreover, there are reasons to believe that the degree of measurement error may vary based on factors such as the age or gender of the test taker.

1. As an example, enrollment in state-funded prekindergarten programs has risen dramatically in the past two decades, from roughly 3 percent (14 percent) of three-year-olds (four-year-olds) in 2002 to 6 percent (34 percent) in 2020. See Friedman-Krauss and others (2021, 9).

2. Jacob and Rothstein (2016) discuss various challenges with assessing student ability and using such assessments in research.

Second, unlike income or temperature, cognitive ability does not have any natural metric. Test scores are reported on different and arbitrary scales. Moreover, there is no reason to believe that the scores reported from standardized assessments have an interval property—that is, a one-unit change having the same meaning at every point on the scale. For example, it is unlikely that an increase from 400 to 450 on the Scholastic Aptitude Test (SAT) represents the same improvement in student knowledge as an increase from 700 to 750. Like utility, measured achievement is best thought of as ordinal, not cardinal. Bond and Lang (2013) illustrate how the unavoidably arbitrary nature of test scaling can influence empirical analysis. They show that the change in the Black-white test score gap between kindergarten and third grade can be as small as zero or as large as 0.6 standard deviations depending on the assumptions made about how to scale standardized assessment results.

Third, the use of standardized scores (subtracting the mean and dividing by the standard deviation) is not a magic bullet. As Jacob and Rothstein (2016) explain, standardized scores are no more comparable across tests or samples than raw or scale scores because standardization is relative to some norming population, which in practice can be small and nonrepresentative. Consider a common empirical result that interventions aimed at younger children tend to have larger effects on standardized test scores than do those aimed at older children. Cascio and Staiger (2012) point out that this pattern may be attributable to the fact that the variance in individual ability increases with age. Given that older children have been exposed to more out-of-school influences as well as more opportunities to learn (or not) in school, it is quite plausible that the true variance of ability increases with age. In this case, one would expect to see the pattern of declining effects with age even in the absence of any true relationship.

These issues complicate the analysis Hoxby proposes. For example, if test scores of young children have more measurement error than adult scores, the correlation between child and adult scores in a region could increase with the child's age even if the relationship between the underlying ability of children and adults remained the same. Even more broadly, I would argue that it is extremely difficult to determine how to measure advanced skills in common standardized assessments, much less create common measures across assessments targeted at different ages. With the limitations imposed by the available data—from the SAT, Early Childhood Longitudinal Study, Kindergarten Class of 1998–99 (ECLS-K), National Education Longitudinal Study (NELS), and Program for the International Assessment of Adult Competencies (PIAAC)—it is even harder to do so.

The approach Hoxby takes of using percentiles of the distribution is certainly reasonable, but far from ideal. Moreover, I would argue that the focus on advanced skills is not necessary. In practice, especially in the aggregate that Hoxby studies, the correlation between the average test scores and the proportion of individuals scoring advanced is quite high. The point is that cognitive skill is important and that the geographic distribution of such skill is increasingly unequal in the United States.

In sum, while there are compelling reasons to believe that skill is distributed more evenly (in terms of geography) among young children than adults, and that the association between child and adult skill increases during adolescence, I think this is a hard case to make statistically and I think the magnitude of the changes is entirely unclear.

THE EDUCATIONAL DECLINE IN ADOLESCENCE While I have some quibbles with how Hoxby uses cognitive assessments, I am willing to accept that adolescence is a critical period in a child's education. In addition to the neurological changes in the brain that children experience as adolescents, teenagers (at least in the United States) live in a unique culture that does not prioritize education. As the famous sociologist of education James Coleman wrote in his 1965 classic, *Adolescents and the Schools*, the "adolescent culture . . . shows little interest in education and focuses the attention of teen-agers on cars, dates, sports, popular music, and other matters just as unrelated to school" (72).

For years education researchers have recognized that children experience a relative drop in academic performance during adolescence. The US Department of Education administers an assessment to a nationally representative set of students in grades four, eight, and twelve every other year. The assessment, known as the National Assessment of Educational Progress (NAEP), is designed to consistently measure student learning across grades and over time. Moreover, a collection of educators and psychometricians work hard to link test scores to more meaningful measures of skill mastery. Based on this well-developed link between performance on test items and skill mastery, NAEP reports the percentage of students scoring in each of four categories: below basic, basic, proficient, and advanced. A consistent finding across subjects and years is the decline in student performance over the school years. On the 2019 NAEP mathematics assessment, for example, 19 percent of fourth graders scored below basic compared with 40 percent of twelfth graders. While 9 percent of students were designated advanced in fourth grade, only 3 percent demonstrated this level of mastery by twelfth grade. A similar pattern is apparent when comparing the United States to other countries. In a 2015 international assessment of fourth graders, the United States ranked thirteenth out of forty-three countries. In a similar assessment given to 15-year-olds across the globe in 2018, the United States ranked twenty-ninth.[3]

POTENTIAL SOLUTIONS The evidence presented by Hoxby emphasizes the importance of reaching adolescents in skill deserts—typically poor, rural communities that lack a critical mass of highly skilled adults. What do we know about strategies to serve this population? Despite the chaos and poor quality that have characterized many children's experiences with online learning during the COVID-19 pandemic, educational technology offers some promise for helping boost achievement of children in disadvantaged, rural communities.

In discussing educational technology, it is important to distinguish between the use of virtual instruction as a supplement and as a substitute for face-to-face learning. A large body of research shows that student outcomes are substantially lower in online environments compared with traditional brick-and-mortar schooling. This is true at both the K-12 and the post-secondary level (Figlio, Rush, and Yin 2013; Hart, Friedmann, and Hill 2018; Heppen and others 2017; Bettinger and others 2015; Woodworth and others 2015).

However, research also points to several ways in which educational technology can enhance learning. First, there is evidence that technology can expand access to high-quality content and instruction. Students in under-resourced schools tend to have fewer advanced placement (AP) offerings, elective courses, and foreign language courses compared with their peers (Barker 1985). Similarly, high-poverty schools are also less likely to offer summer school, where students can retake a course they failed during the year (Watson and Gemin 2008). The best evidence on whether simply improving access to different courses through virtual schooling affects students' academic outcomes comes from a large-scale random assignment study carried out in Maine and Vermont (Heppen and others 2012). Sixty-eight schools that had not historically offered Algebra I to eighth graders were randomly assigned to either a treatment group, which was given access to an online Algebra I course, or a control group, which did not receive access. Algebra-ready students in treated schools showed

3. Data for fourth graders come from the 2015 Trends in International Mathematics and Science Study (TIMSS) and data for 15-year-olds come from the 2018 Programme for International Student Assessment (PISA). Assessment results can be accessed through the National Center for Education Statistics, https://nces.ed.gov/surveys/international/ide/. Forty-three countries reported scores on both of these exams and are thus used to generate the calculations reported here.

improvements on test scores and took more advanced courses in high school. Goodman, Melkers, and Pallais (2019) illustrate this potential at the post-secondary level. They study the Georgia Institute of Technology's online MS in computer science. Using a regression discontinuity design that exploits the admissions threshold, they show that the online option substantially increases overall enrollment.[4]

There is also compelling evidence that so-called intelligent tutoring systems, which provide instruction, practice, and feedback tailored to the needs of individual students, can improve student achievement. Some of the most compelling evidence comes from large randomized trials conducted in developing countries such as India (Banerjee and others 2007; Muralidharan, Singh, and Ganimian 2019). However, there is also evidence that intelligent tutoring is effective in the United States (Escueta and others 2020). As important, existing research suggests important lessons for developers and practitioners. For example, teachers in the United States often face challenges in effectively implementing computer-aided technology in a classroom setting (Drummond and others 2011; Pane and others 2010, 2013). In addition, experience to date suggests that computer-aided learning alone—in the absence of personal interaction between an adult and child—is not particularly effective, and the most effective programs are "blended," meaning they include some group-based instruction along with some individual student work with a personalized learning technology (Muralidharan, Singh, and Ganimian 2019).

The spread of online instruction driven by pandemic lockdowns spurred renewed interest in technology-aided education resources. The recently established National Student Support Accelerator provides a comprehensive set of resources for school districts and communities interested in implementing high-intensity, technology-supported tutoring programs.[5] Researchers and educators are taking a careful look at the potential of these strategies. If this work paves the way to reach at-risk adolescents in skill deserts, then it would truly be a silver lining of the pandemic.

4. There is even greater evidence in the developing economy context. For example, Bianchi, Lu, and Song (2020) find that the Chinese government's push to expand computer-assisted learning in rural communities substantially improved educational attainment.

5. Annenberg Institute for School Reform at Brown University, "National Student Support Accelerator," https://studentsupportaccelerator.com/.

REFERENCES FOR THE JACOB COMMENT

Banerjee, Abhijit, Shawn Cole, Esther Duflo, and Leigh Linden. 2007. "Remedying Education: Evidence from Two Randomized Experiments in India." *Quarterly Journal of Economics* 122, no. 3: 1235–64.

Barker, Bruce. 1985. "Curricular Offerings in Small and Large High Schools: How Broad Is the Disparity?" *Research in Rural Education* 3, no. 1: 35–38.

Bettinger, Eric, Lindsay Fox, Susanna Loeb, and Eric Taylor. 2015. "Changing Distributions: How Online College Classes Alter Student and Professor Performance." Working Paper 15–10. Stanford, Calif.: Stanford Center for Education Policy Analysis.

Bianchi, Nicola, Yi Lu, and Hong Song. 2020. "The Effect of Computer-Assisted Learning on Students' Long-Term Development." Working Paper 28180. Cambridge, Mass.: National Bureau of Economic Research.

Bond, Timothy N., and Kevin Lang. 2013. "The Evolution of the Black-White Test Score Gap in Grades K–3: The Fragility of Results." *Review of Economics and Statistics* 95, no. 5: 1468–79.

Cascio, Elizabeth U., and Douglas O. Staiger. 2012. "Knowledge, Tests, and Fadeout in Educational Interventions." Working Paper 18038. Cambridge, Mass.: National Bureau of Economic Research.

Chetty, Raj, and Nathaniel Hendren. 2018a. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *Quarterly Journal of Economics* 133, no. 3: 1107–62.

Chetty, Raj, and Nathaniel Hendren. 2018b. "The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates." *Quarterly Journal of Economics* 133, no. 3: 1163–228.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106, no. 4: 855–902.

Coleman, James S. 1965. *Adolescents and the Schools.* New York: Basic Books.

Drummond, Kathryn, Marjorie Chinen, Teresa Garcia Duncan, H. Miller, Lindsay Fryer, Courtney Zmach, and Katherine Culp. 2011. *Impact of the Thinking Reader [R] Software Program on Grade 6 Reading Vocabulary, Comprehension, Strategies, and Motivation: Final Report*. Washington: National Center for Education Evaluation and Regional Assistance.

Escueta, Maya, Andre Joshua Nickow, Philip Oreopoulos, and Vincent Quan. 2020. "Upgrading Education with Technology: Insights from Experimental Research." *Journal of Economic Literature* 58, no. 4: 897–996.

Figlio, David, Mark Rush, and Lu Yin. 2013. "Is It Live or Is It Internet? Experimental Estimates of the Effects of Online Instruction on Student Learning." *Journal of Labor Economics* 31, no. 4: 763–84.

Friedman-Krauss, Allison H., W. Steven Barnett, Karin A. Garver, Katherine S. Hodges, G. G. Weisenfeld, and Beth Ann Gardiner. 2021. *The State of Preschool 2020*. New Brunswick, N.J.: National Institute for Early Education Research, Rutgers.

Goodman, Joshua, Julia Melkers, and Amanda Pallais. 2019. "Can Online Delivery Increase Access to Education?" *Journal of Labor Economics* 37, no. 1: 1–34.

Hart, Cassandra, Elizabeth Friedmann, and Michael Hill. 2018. "Online Course-Taking and Student Outcomes in California Community Colleges." *Education Finance and Policy* 13, no. 1: 42–71.

Heppen, Jessica B., Nicholas Sorensen, Elaine Allensworth, Kirk Walters, Jordan Rickles, Suzanne Stachel Taylor, and Valerie Michelman. 2017. "The Struggle to Pass Algebra: Online vs. Face-to-Face Credit Recovery for At-Risk Urban Students." *Journal of Research on Educational Effectiveness* 10, no. 2: 272–96.

Heppen, Jessica B., Kirk Walters, Margaret Clements, Ann-Marie Faria, Cheryl Tobey, Nicholas Sorensen, and Katherine Culp. 2012. *Access to Algebra I: The Effects of Online Mathematics for Grade 8 Students*. Washington: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, and US Department of Education.

Jacob, Brian, and Jesse Rothstein. 2016. "The Measurement of Student Ability in Modern Assessment Systems." *Journal of Economic Perspectives* 30, no. 3: 85–108.

Muralidharan, Karthik, Abhijeet Singh, and Alejandro J. Ganimian. 2019. "Disrupting Education? Experimental Evidence on Technology-Aided Instruction in India." *American Economic Review* 109, no. 4: 1426–60.

Pane, John F., B. A. Griffin, D. McCaffrey, R. Karam, L. Daugherty, and A. Phillips. 2013. "Does an Algebra Course with Tutoring Software Improve Student Learning?" Santa Monica, Calif.: RAND Corporation.

Pane, John F., Daniel F. McCaffrey, Mary Ellen Slaughter, Jennifer L. Steele, and Gina S. Ikemoto. 2010. "An Experiment to Evaluate the Efficacy of Cognitive Tutor Geometry." *Journal of Research on Educational Effectiveness* 3, no. 3: 254–81.

Watson, John, and Butch Gemin. 2008. *Using Online Learning for At-Risk Students and Credit Recovery*. Vienna, Va.: North American Council for Online Learning.

Woodworth, James L., M. E. Raymond, K. Chirbas, M. Gonzalez, Y. Negassi, W. Snow, and C. Van Donge. 2015. *Online Charter School Study 2015*. Stanford, Calif.: Center for Research on Educational Outcomes.

GENERAL DISCUSSION Carol Graham emphasized two points that were made in both Caroline Hoxby's presentation and Erik Hurst's discussion: adolescents are likely giving up if they are not on a trajectory toward college and low-skill jobs are disappearing. Graham noted that these patterns line up with the main findings from her work on deaths of despair. In her research, Graham finds that deaths of despair are high among non-college-educated white people and are highly correlated with low levels of hope. Graham observed that many of the areas with low cognitive skills that Hoxby highlights in her paper are predominantly white areas with low levels of civic education, low levels of trust in science, and high rates of opioid use. Graham suggested that these trends can create a vicious feedback loop in which prime-age males experience low levels of participation in the labor force and high levels of despair.

Graham then argued that these patterns present a bleak outlook for future generations given that we do not see large movements out of these areas of despair and into areas of opportunity. However, Graham noted that one source of optimism is the declining gap in educational attainment by race, which Hurst noted in his discussion as well. Graham pointed to survey data which show that both Black and Hispanic people are more likely to believe in the value of a college education than low-income white people, even if it may be harder for them to get one. Finally, Graham concluded that hope plays a central role in helping adolescents overcome barriers to receiving a college degree and performing well in the labor market.

Richard V. Reeves brought up Melissa Kearney and Phillip Levine's research on teenage pregnancy rates across the country.[1] Reeves suggested this work could complement Hoxby's geographic analysis of cognitive skill attainment and lead to some insights.

Janice Eberly asked Hoxby to comment on gender differences in cognitive skill levels. Eberly referred to Hurst's mention of gender differences in educational attainment and wondered if the same patterns hold in Hoxby's data.

Taking a step back, Hoxby explained that the paper she presented is the third in a series of three papers looking at cognitive skill patterns in the United States, where the first paper focuses on the long-term effects of not making the transition to advanced cognitive skills in early adolescence and the second paper analyzes natural experiments in cognitive skill interventions. The first paper finds that failing to transition to advanced cognitive skills in early adolescence does have long-term consequences, and the second one finds that successful interventions are more productive if done during adolescence than if done before or afterward, leading Hoxby to conclude that adolescence is a particularly malleable period in cognitive skill development.

Addressing Eberly's comment on gender differences, Hoxby noted that male and female educational trajectories look quite different from one another. Up until grade three, the trajectories are mostly similar, but afterward

1. See, for example, Melissa S. Kearney and Phillip B. Levine, "Why Is the Teen Birth Rate in the United States So High and Why Does It Matter?" *Journal of Economic Perspectives* 26, no. 2 (2012): 141–63.

they diverge, Hoxby explained. This is due to differences in timing of cognitive brain development in males and females. Boys usually fall about a year behind in terms of cognitive development and that gap remains up until the end of high school. Hoxby argued that the cognitive development lag that males experience during this critical period has longer-term effects, such as lower college degree attainment levels for males.

Jim Stock expressed a concern that some of the correlation patterns in Hoxby's work may be misinterpreted as causal relationships. Specifically, Stock referred to the associations between areas with low cognitive skills and trust in scientists' views on climate change. Stock expressed the need for caution in interpreting these associations so that causal relationships are not inappropriately attributed to complex correlations. He then asked Hoxby to comment on how her work can be used to better understand the causal mechanisms that may underlie these correlations.

In response to Stock's comments, Hoxby reiterated that her paper is careful to distinguish between correlations and causal relationships. She claimed that there may be several, non-mutually exclusive causal mechanisms having an impact on cognitive skill development, such as the teen pregnancy rates that Reeves noted earlier. However, Hoxby highlighted two associations that she believes to be the important takeaways from her paper. The first is that there is a lot of movement in cognitive development that happens in early adolescence. In other words, one's cognitive trajectory is not determined by the third grade. Hoxby argued that economists of education are often too fatalistic about children's potential to develop advanced cognitive skills later in life even if they fall behind in early childhood. The fact that adolescents are highly influenced by their environments is a reason to be optimistic about potential cognitive skill interventions that target adolescents. The second point Hoxby underscored is that there is a strong association between the cognitive skills of children and the skills of the adults around them. Hoxby then clarified that this does not necessarily imply that adults' cognitive skills causally affect their children's cognitive skills development. However, she reiterated that if we do want to get closer to understanding the underlying causal mechanisms, we need to carefully investigate the existing strong correlations.

Erik Hurst pondered whether the mechanisms that cause skill deserts within cities differ from the mechanisms that lead to skill deserts across cities.

In response to Hurst's question, Hoxby suggested that the same mechanisms may be at play in a large city as the ones occurring across cities. She used Chicago as an example of a city in which the differences in school quality and other environmental factors across neighborhoods is similar to the differences found across cities in the country. To accentuate her point, Hoxby discussed the wide variation in middle school quality across the country. She argued that public middle schools are among the most neglected schools with the most teacher vacancies, leading to large differences in quality from one neighborhood to the next.