

NATIONAL
ACADEMY
of
EDUCATION

Workshop to Examine Current and Potential Uses of
NCES Longitudinal Surveys
by the Education Research Community

Building Better Longitudinal Surveys (on the Cheap) Through Links to Administrative Data

Susan Dynarski
University of Michigan



National Academy of Education

***Workshop to Examine Current and Potential Uses of NCES Longitudinal Surveys
by the Education Research Community***

**Building Better Longitudinal Surveys (on the Cheap)
Through Links to Administrative Data**

Susan Dynarski
University of Michigan

December 2014

This paper was prepared for the National Academy of Education's *Workshop to Examine Current and Potential Uses of NCES Longitudinal Surveys by the Education Research Community*, held on November 5-6, 2013, in Washington, DC. The workshop was organized to provide input to the National Center for Education Statistics (NCES) on how its longitudinal survey program could be enhanced to serve the changing needs of the education research community in light of (1) changing technological opportunities, (2) changing data availability (with an increasing focus on the classroom and learning outcomes), and (3) a changing U.S. population. The collection of commissioned papers reflects the views of the authors and not necessarily those of the National Academy of Education or the U.S. Department of Education. This paper was supported with funding from the Institute of Education Sciences.

INTRODUCTION

The charge of this National Academy of Education workshop is to evaluate the current and potential uses of the longitudinal surveys fielded by the National Center for Education Statistics (NCES) of the U.S. Department of Education (ED). The present paper is focused on how administrative data can be used to improve these surveys. Other papers in this workshop have done a thorough job of describing in detail the current surveys and their contents. Rather than duplicate their work, I discuss the strengths and weaknesses of the existing NCES surveys in broad terms. I discuss in more detail the administrative data that could be used to increase the power, scope, and utility of the NCES surveys.

The scale and scope of education data has changed dramatically over the past few decades. Until recently, surveys were the main source of information on student outcomes such as enrollment, test scores, skills, educational attainment, employment, and earnings. Education researchers relied heavily on the Panel Study of Income Dynamics (PSID) and the National Longitudinal Surveys (NLS, of the Bureau of Labor Statistics [BLS]). The longitudinal surveys of the NCES joined this set of surveys in the 1970s.¹ NCES tracked a cohort each decade through high school, through college, and for a few years of employment. Other NCES surveys tracked cohorts of college students and kindergarteners.

For education researchers interested in documenting the life cycle of human capital accumulation, these detailed longitudinal surveys were the only game in town. Their main drawbacks were their infrequency and small size. The infrequency of the surveys meant they could not be used to track short-term changes that might arise from (for example) a shift in national policy. The small size meant that they could not be used to measure variation across states, including that which might arise from (for example) a shift in state policy.

In the last 10 years large, state longitudinal data systems have emerged that track entire populations of students from kindergarten through elementary and secondary school, into college and (in some cases) into the labor force. These datasets are maintained by school districts and states. They were not built for the needs of researchers but rather as a response to federal reporting mandates. An example is the No Child Left Behind Act (NCLB), which requires that students be tested periodically and that scores be reported by race, ethnicity, eligibility for subsidized lunch, and special-education status. To comply with NCLB, states improved the student-level data systems that contained these measures.

Also spurring the development of the student longitudinal data system was the promulgation of a standardized definition of high school graduation and dropout. This new definition involves the tracking of the enrollment and attainment of students from ninth grade forward. States improved the longitudinal tracking of their students in order to generate these standardized graduation rates.

Yet another impetus to the growth of a state longitudinal data system were reporting requirements attached to federal stimulus funding during the recent recession. States

¹ The BLS longitudinal surveys include NLS Young Men, NLS Older Men, NLS Young Women, and NLS Mature Women, all initiated in the 1960s, and the National Longitudinal Surveys of Youth initiated in 1979 and 1997. See <http://www.bls.gov/nls> for more information.

that received state fiscal stabilization funds were required to report on the number of high school graduates who attended college, completed a year of credits, and took remedial courses in college. To comply with these requirements, states linked their longitudinal data on high school students to data from their postsecondary systems and/or the National Student Clearinghouse (NSC).

THE PROSPECTS FOR IMPROVING NCES SURVEYS USING ADMINISTRATIVE DATA

Combined, these various datasets track tens of millions of students, covering the entire population of elementary and secondary students in public schools and (through NSC) 93 percent of students at colleges nationwide. Unlike the NCES surveys, these administrative datasets contain a very limited number of variables. However, they hold great promise as a complement to the NCES surveys, creating the opportunity for NCES to track a subset of outcomes for their survey respondents longer, more cheaply, and more reliably.²

For example, NSC could be used to track the postsecondary attainment of respondents. While a survey was still in the field, NSC could capture enrollment of non-respondents. After the survey had closed, NSC could be used to continue to track postsecondary attainment into respondents' thirties or even forties. Since students increasingly continue postsecondary attendance into their thirties (Turner, 2004), this approach would more accurately measure postsecondary attainment than does a survey that stops at a younger age.

Administrative data could also be used to turn what are primarily *cross-sectional* datasets into *longitudinal* datasets. For example, the National Assessment of Educational Progress (NAEP) is a set of student-level cross sections. Data from state longitudinal data systems could be attached, allowing for the measurement of grade progression, high school graduation, and standardized test scores in order to examine the relationship between the nationally normed NAEP (taken in grades 4, 8, and 12) and these outcomes.

These linkages can be initiated both from the field (by researchers) and by NCES staff. Easing the process of making such linkages should be a priority for NCES. So, too, should be the development of expanded, secure channels for researchers to access NCES microdata. The NCES surveys are unusual among social science datasets in that accessing microdata of any sort requires a restricted-use data license, which in turn demands technological resources that many researchers do not have (e.g., a separate computer devoted solely to NCES data). NCES should look for inspiration to other data agencies (in particular, Census) that have made their microdata freely available to researchers.

² NCES has long supplemented its surveys with data from schools: several of the longitudinal surveys include high school transcripts, for example. What has changed is that this information is now collected, standardized, and stored at the state (and sometimes national) level, which potentially reduces the time costs of collecting and harmonizing these data.

RECOMMENDATIONS

In light of these opportunities, it is time to reconsider NCES's approach to survey design and collection. I have four recommendations:

1. Supplement NCES surveys with administrative data. Focus NCES surveys on collecting information that is not contained in administrative data sources.
2. Use administrative data to convert cross-sectional surveys into longitudinal surveys.
3. Support researcher-initiated requests to link existing NCES surveys to other data sources.
4. Explore ways to make microdata more widely available to researchers, particularly by looking to successful initiatives at Census.

I describe the rationale for these four recommendations in the next section of this paper. In the last section of the paper I use a case study to explore the rationale for these recommendations. The case study is student debt; I discuss the data needs of researchers and policy makers in this arena and prospects for filling those needs with a combination of NCES surveys and administrative data.

Recommendation 1: Supplement NCES Surveys with Administrative Data, Focusing Survey Efforts on Information Not Contained in Administrative Data

When there are administrative data available, NCES should rely on them rather than duplicate efforts with a survey. For example, NSC contains information on college attendance spells and the identity of those colleges. The Internal Revenue Service (IRS) holds similar data in the form of the information returns (Form 1098-T) that colleges file for every student. This information can also be obtained from the majority of states that now have information on college attendance in their longitudinal data systems.³

Data from these administrative sources can therefore be used to measure spells of college attendance and the identity of the college attended, rather than surveys. At the very least, these data can be used to prepopulate surveys and respondents asked to confirm their accuracy.

Similarly, every state now tracks the enrollment and grade progression of its students. These data can be used to capture the identity of the school attended, grade attainment, and high school graduation. The data also include information on special education, English learner, and subsidized lunch status. These data can be used to either replace or confirm survey responses on these topics.

IRS holds extensive data on the income and other characteristics of households. The Census Bureau has successfully negotiated access to these data for at least one of its surveys, the Survey of Income and Program Participation (SIPP). SIPP microdata are publicly available. Appended to the SIPP microdata are individual-level data from IRS and the Social Security Administration (SSA). In this dataset, the values of the

³ The disadvantage of the state data is that (unless the state has purchased NSC data or formed a data consortium with neighboring states) it will miss enrollment spells that occur outside of a state.

IRS and SSA variables have been “perturbed” to protect the privacy of the individual respondents. This is referred to as the SIPP Synthetic Beta.⁴

Researchers who have used these public data to develop their statistical models can upload their code to run their analyses on the original, unperturbed SIPP/IRS/SSA data. After being checked for compliance with privacy protocols (e.g., no cells below a certain size), the analytic results are returned to the researcher.

Similarly, perturbed versions of IRS and SSA data could be linked (for example) to the National Postsecondary Student Aid Study (NPSAS). Researchers could conduct initial analyses using these perturbed data and subsequently upload code to run on the complete, unperturbed data. Such an approach could overcome the legal, organizational, and political barriers to a linkage of NCES data with other sources.

A key drawback is that NCES does not control these data sources and has to negotiate access. This drawback must be weighed against the cost of NCES duplicating the efforts of other agencies. By obtaining key variables from administrative sources, time and money are freed up that could be used to increase sample sizes, increase the length of follow-ups, and field more frequent surveys (e.g., a high school survey every 5 years rather than every 10 years).

Furthermore, these administrative data can be used to examine the characteristics of nonrespondents. A key set of variables can continue to be collected even for those who leave the sample. These variables can be used to check the accuracy of sample weights that adjust for nonresponse and attrition. They can be used to answer these questions: Does the longitudinal behavior of nonrespondents align with that of respondents who are observationally identical at baseline (the assumption underlying the construction of weights that account for attrition)? Do the current sample weights, when used to reweight respondents, generate the same results as an administratively supplemented dataset that includes both respondents and nonrespondents?

Recommendation 2: Use Administrative Data to Convert Cross-Sectional Surveys to Longitudinal Surveys

Administrative data can be used to turn what are currently cross-sectional datasets into longitudinal datasets. For example, the periodic NAEP data are student-level cross sections. Data from states and NSC could be attached to NAEP in order to capture the relationship between the nationally normed NAEP (taken in grades 4, 8, and 12) and educational attainment. As discussed above, data from IRS or SSA on earnings could be combined in a secure setting, creating a dataset that allows us to measure whether the skills tested in NAEP translate into labor market success. These patterns could be examined separately by the variables gathered in NAEP, for example, school type and student demographics.

If the waivers signed by participants in past NAEP do not allow for merges at the individual level, group-level merges can be conducted. For example, lists of students in state-year groups could be created and sent to the relevant data holder, who could

⁴ See <http://www.census.gov/programs-surveys/sipp/methodology/sipp-synthetic-beta-data-product.html> for more information.

then return cell means and variances.⁵ This is how ED has obtained data from SSA on the earnings of college graduates in order to comply with the congressionally required “gainful earnings” reporting (U.S. Department of Education, 2012). It is also how researchers typically obtain merges of surveys such as the decennial Census with SSA earnings data (see, for example, Angrist et al., 2011). Properly defined cell means could then be used to evaluate the effect of state-specific policies on NAEP scores (as is currently done) as well as educational attainment and earnings.

This new resource would massively expand the data available to researchers to who seek to identify causal connections between state policies and educational attainment. Note that, in NCES data, such patterns can currently be tracked only in the relatively small and infrequent longitudinal surveys (every 10 years), which is too infrequent for capturing the effects of short-term shifts in policy.

Recommendation 3: Support Researcher-Initiated Requests to Link the Longitudinal Surveys to Other Data Sources

Going forward, NCES should create the conditions that allow their surveys to be regularly linked to the resources described above. However, NCES had a rich collection of existing surveys that could potentially be linked to such data. NCES should support researchers in linking the other data sources (e.g., NSC or state data) to the longitudinal surveys. This will, at low cost to NCES, improve the scope and quality of the surveys and make them more useful for researchers.

The NCES surveys are not currently amenable to such links. In older surveys, respondents did not sign waivers that would allow such administrative links; going forward this should be the standard protocol. For surveys in which these waivers have been signed, researchers need a process for requesting, making, and paying for such links. Without such an institutionalized process, the creation of such links relies on individual relationships and the entrepreneurship of individual NCES staff. These efforts will falter as NCES staff shift jobs or retire.

Below I describe one possible, institutionalized process for data requests. There are many feasible models. What matters most is that the process is publicized and understood by the research community. NCES should not be discouraged if take-up is limited in the first few years that the process is put in place; it takes years for information about new data opportunities to disseminate through and be embraced by the research community.

An online portal would be created for the initiation of proposals for data linkages. The Institute of Education Sciences (IES) would review the scientific merit of proposed data links on a rolling basis. An NCES staff member would be tasked with working with researchers to get the pertinent identifying information transferred to the organization that holds the target data. In the case of NSC, for example, the names and dates of birth would be securely transmitted from NCES to NSC, which would match on data about college attendance and securely transmit it back to NCES.

NCES would merge the new data onto the longitudinal survey and release it to the researcher for analysis. The resulting match would be available to only the researcher

⁵ As long as they remain large enough, the groups could be sliced more finely: state-year-gender, for example, or state-year-gender-race.

who paid for it for a specified period of time (one year is reasonable) and would then be available to other researchers. The resulting data could be distributed through the current channels (compact disks [CDs] sent to those who hold restricted-use licenses). In the next section I describe some alternative models for distribution of NCES data, including those matched to administrative data.

Recommendation 4: Explore New Ways to Safely Make More Detailed Microdata Widely Available to Researchers

NCES releases public-use versions of its surveys, as well as restricted-use versions. I do not know of a resource that clearly delineates exactly what additional information is available in the restricted-use versions. Broadly, they contain more detailed information about geography, as well as more finer-grained information about (for example) income.⁶ NCES sends CDs of its data to researchers who have completed restricted-use data licenses.

The intent of the restricted-use model is to prevent disclosure of respondent identities. The risk of disclosure rises as more researchers use it: it sits on more computers, for example, increasing the risk of a hack or stolen data. And the more researchers hold the restricted-use data, the more expensive it is for NCES to supervise its use. This suggests (to me, at least) that NCES should make the public-use versions as useful as possible in order to reduce demand for the restricted-use versions.

The restricted-data license model is a substantial roadblock for many researchers who want to use NCES microdata. Getting access requires completing a restricted-use data application. The speed of this application process has improved considerably, but it still discourages many who want to make use of the data. Many policy analysts and congressional staff, for example, cannot set up a standalone computer for NCES data. Yet often they need to do analyses that depend on (for example) fine measures of income or geography. How do they get this information into their hands, so that policy can be informed by the rich information NCES has gathered in its surveys?

NCES is more restrictive in what it includes in its public-use datasets than are other statistical agencies. The other surveys I use regularly as an education researcher are freely distributed as microdata, and they contain variables that NCES limits to the restricted-use versions. The American Community Survey, Current Population Survey, SIPP, PSID, and the National Longitudinal Surveys are freely available for download by any researcher. The public versions of these data meet the needs of most researchers, because they include detailed measures of variables such as geography (state, sometimes metropolitan area) and income (sometimes top coded, but always continuous rather than in brackets). Restricted-use versions contain even more information (e.g., county and school identifiers in the National Longitudinal Survey of Youth and PSID). The few researchers who need this additional information apply for restricted-

⁶ I do not know of a document that details which variables are in the public-use datasets versus the restricted-use datasets. It may well exist, but it is not easily available on the webpages that describe each data set (e.g., http://nces.ed.gov/surveys/nels88/data_products.asp), which is how researchers learn about NCES data products. Making the distinctions between variables available in each datasets would help casual users know whether they truly needed the restricted-use version or could make do with the public version.

use licenses. Everyone else can freely use the widely available microdata that do not include this information.

Census has been particularly aggressive and creative in both matching administrative data to its surveys and finding new ways to distribute those data to researchers. Public-use, microdata versions of their data products (e.g., American Community Survey, SIPP) are freely available for download by any researcher who wants access. Versions of the American Community Survey with very fine measures of geography (e.g., block) can be used in the Census Research Data Centers (RDCs) by researchers who have successfully applied for these data. A version of SIPP that is linked to data from the IRS is also available in the Research Data Centers. The Research Data Center model, however, is not one I would recommend to NCES, since they create their own barriers for researcher use.⁷

A promising route for NCES is that taken by Census for the distribution of exceptionally detailed versions of SIPP. A version of SIPP that is linked to data from the IRS is available online in the form of the SIPP Synthetic Beta.⁸ The SIPP Synthetic Beta is a promising model for NCES to handle the distribution of versions of its data that include variables sufficiently detailed that they threaten to reveal individual identities.

Census publicly releases a microdata version of its SIPP/SSA/IRS match that is “perturbed,” with some variables statistically blurred to prevent identification. An unperturbed version of the matched data sits on the Census servers. Researchers run and refine their statistical models on the publicly available data, on their own computers. They can then upload the resulting code to the Census servers, where it is run on the original, unperturbed data. Results are returned to the researcher after being checked for compliance with data standards (e.g., minimum cell sizes).

This model could be used as the standard for merges of NCES surveys with sensitive data from IRS, SSA, or the states. Agencies reluctant to allow their data to be released to researchers may well cooperate when the SIPP Synthetic model is used, with public versions being statistically perturbed. This approach appears to have worked with IRS and SSA, agencies that are notoriously protective of their data. I recommend that NCES consult closely with the Census staff who have successfully negotiated these data merges and releases.

CASE STUDY: UNDERSTANDING STUDENT DEBT USING NCES SURVEYS AND ADMINISTRATIVE DATA

In order to illuminate the promise of administrative data for improving these surveys, I focus on a particular topic in education: student loans. Student loans are the subject of heated discussion in Congress, the Department of Education, think tanks, the media, and the academic press. How have the longitudinal surveys informed this discussion? Where have they fallen short? How could administrative data be used to improve their utility? My intent is to use this particular case study to draw out les-

⁷ Many researchers do not live close to an RDC, and the process of getting a project approved is quite lengthy. There is also a fee associated with the use of the RDCs, though some universities cover this fee for researchers.

⁸ See <http://www.census.gov/programs-surveys/sipp/methodology/sipp-synthetic-beta-data-product.html> for more information.

sons and insights about the longitudinal surveys that are relevant to multiple policy domains, not just postsecondary education.

The Policy Context

Student-loan debt has mounted to \$1 trillion, now surpassing credit cards as the third-largest form of consumer debt.⁹ This has triggered a national conversation about the cost of college and the appropriate level of student borrowing. With seven million student loans in default, many ask, “Is there a student-debt crisis?”¹⁰ At the heart of this question is a concern that borrowing is out of line with the value of college.

Economists point to the high payoff to college to put student borrowing in perspective. Over a lifetime, the typical holder of a bachelor’s degree earns several hundred thousand dollars more than a high school graduate (Barrow and Rouse, 2005). Even those with only some college see lifetime gains of about \$100,000 (Greenstone et al., 2013).

How do these figures compare to student borrowing? Sixty-nine percent of undergraduates borrow less than \$10,000 and 98 percent borrow less than \$50,000. About half of Bachelor of Arts (BA) recipients borrow for college, with a debt of around \$27,000 (College Board, 2012). While those who borrow would surely prefer not to be in debt, to most these debt loads look reasonable when compared to the financial gains of college.

The recent spike in defaults on student loans is worrisome. Defaults are not driven by the small fraction of borrowers with large loans: borrowers with typical levels of student debt struggle with their payments. Those who default have borrowed less than others: the average loan in default is \$14,000 while the average loan in good standing is \$22,000.¹¹ This pattern of defaults is consistent with two scenarios with very different implications for policy.

One scenario is that defaulters have *temporarily* low earnings and their loans fall into distress during these unusual bad times. At low cost to government, an income-based repayment (IBR) program would insure borrowers against these temporary downturns by automatically reducing their payments. If *lifetime* earnings are sufficient to pay off the loans, this system can be self-funding.

An alternative scenario is that those who default have *permanently* low earnings that cannot support even moderate debt loads. An IBR plan would still help these borrowers, but the ultimate cost to government would be much higher, since many of these loans will ultimately be forgiven. The cost of making, servicing, and forgiving these loans could be so high that a grant program could be cheaper for taxpayers.

⁹ The top two forms of debt are home mortgages and car loans (Lee, 2013). Note that these widely cited statistics on loan debt are obtained not from NCES or the Department of Education but from credit records purchased by the Federal Reserve Bank of New York. The Federal Reserve Banks have been charged with improving tracking of student debt. Unable to obtain data from the Department of Education, they have resorted to buying data on federal loans from the private sector (i.e., from credit agencies such as EquiFax).

¹⁰ There are 6.5 million borrowers in default as of the third quarter of 2013. See <http://studentaid.ed.gov/sites/default/files/fsawg/datacenter/library/PortfolioByLoanStatus.xls>, accessed October 2013.

¹¹ Calculated from U.S. Department of Education, National Student Loan Data System (2013).

The Data Needed for Research

Distinguishing between these two scenarios requires individual-level, longitudinal data on student borrowing that follows former students for 25 years after college.

Why 25 years? IBR plans have students paying back their loans in 20 to 25 years, when any remaining balance is forgiven. Costing out these programs therefore requires tracking earnings for decades.

Why individual-level, longitudinal data? Individual-level data are needed to capture the shocks to income that IBR programs insure against. The payments required of borrowers with different earnings paths cannot be backed out from group averages. Any analysis that relies on averages will smooth away the within-person shocks that are needed to estimate the benefits and costs of IBR.

Understanding the relationship between earnings and borrowing is critical for designing sound aid policy. If many former students are carrying debt beyond their capacity to repay, we need to reconsider the parameters of student borrowing, such as loan limits, loan forgiveness, and repayment structures. All of these topics are currently under discussion in Washington, with little data to inform the debate.

Data Prospects: Earnings

Multiple data sources contain information on lifetime earnings: the NLSs (fielded by the Department of Labor) and PSID are well-known examples. The NCES longitudinal surveys, at present, do not contain data on lifetime earnings. The decadal cohort surveys, which follow a high school class every 10 years, do not follow respondents beyond early adulthood, typically stopping when respondents are in their twenties.¹² The postsecondary surveys do not go much later, with the last survey waves fielded 10 years after the start of college.¹³

This leaves a major gap for researchers. Multiple studies have now shown that educational interventions do not demonstrate their full effects until students are well into adulthood (e.g., Chetty et al., 2011; Dynarski et al., 2013; Kemple, 2008). The NCES surveys therefore miss many of the potentially positive effects of education policy. In the context of student loans, the NCES surveys stop before many students have finished paying off their education loans, which hampers the analysis of loan policy and the estimation of the long-term effects of college.

There have recently been advances in making administrative, longitudinal data on earnings available to researchers. As discussed earlier in the paper, the SIPP has obtained earnings data from IRS and SSA for its samples.

¹² The National Longitudinal Survey of the High School Class of 1972 surveyed students until 1986, when they were about 32. High School and Beyond, which includes the high school class of 1982, stopped surveying students when they were in their twenties (in 1986, four years after high school). So did the National Education Longitudinal Study of 1988, which stopped surveying in 2000 (when respondents were about eight years out high school). The surveys currently in the field (Education Longitudinal Study of 2002 and the High School Longitudinal Study of 2009) are not planned to survey any later in life than their predecessors. See <http://nces.ed.gov/surveys/hsb> for more information.

¹³ The Baccalaureate and Beyond has varied in how long it tracks students. Graduates who started college in 1993 were followed for 10 years, which would yield a typical exit age of late twenties. See <http://nces.ed.gov/surveys/b&b/about.asp> for more information.

Data Prospects: Borrowing and Debt

The postsecondary surveys contain data on student borrowing and debt. Some of this information is drawn from respondent surveys (e.g., private student loans), but most of it is from the administrative system of the federal loan programs, the National Student Loan Data System (NSLDS). NSLDS contains information on the universe of federal borrowers.

NCES matches NPSAS respondents to NSLDS. These matches appear to be done at baseline and when the respondents are reinterviewed (it is unclear from the documentation). They are not regularly refreshed with updated data from NSLDS. In principle, these data could be updated each academic year.

This NPSAS-NSLDS match (and the matches with NSLDS of the datasets that are drawn from NPSAS) constitutes the most comprehensive survey data on student borrowing. The data are used and cited widely. The College Board uses NPSAS data in its popular *Trends in Student Aid* series. To my knowledge, these matches constitute the only dataset in which, at the individual level, detailed data are available on both demographics and borrowing.

The drawbacks of NPSAS are its frequency and size. New samples are started every 4 years (and its students surveyed every 2 years) and then data are released with a lag of a year or two. NPSAS therefore cannot be used to track annual fluctuations in student borrowing. And while NPSAS is large (100,000 students), it is not representative at the state level, so it cannot be used to tabulate state-level estimates of student borrowing.¹⁴

Prospects for Improving Data on Borrowing and Earnings

As described above, many data sources contain individual-level, long-term earnings, while a smaller number contains information on student borrowing. The key data sources are listed in the Appendix, which serves to show that no single dataset contains the required information on both borrowing and income and to identify the datasets that could be combined. I next describe a scenario for linking these datasets.

NPSAS is the authoritative source for information about student borrowing (see Appendix). Information in NPSAS about individual students' federal borrowing is drawn from the administrative data that ED uses to run the loan programs (NSLDS). These data should be updated annually and updates should continue after the survey has "concluded." National Student Clearinghouse data can be used to update student postsecondary attendance.

A key weakness of NPSAS is its earnings data. For those who earn a BA, self-reported earnings data are collected every few years for just 10 years after graduation. For those who do not earn a BA, earnings data are collected for just 6 years after college entry. SSA and IRS hold detailed data on earnings and income. These data could be merged with the NPSAS data, just as SIPP data are currently merged with these data sources. These merges could continue well after the survey has left the field, and it can occur between surveys.

As is done with SIPP, perturbed versions of the NPSAS match with IRS and SSA

¹⁴ NCES has made substantial progress on this last point, with the recent NPSAS surveys including representative samples for some larger states.

data could be made publicly available. Researchers could conduct initial analyses using these perturbed data and subsequently upload code to run on the complete, unperturbed data.

CONCLUSION

There are no obvious conceptual, logistical, or technological barriers to linking data from the NCES surveys to administrative data on earnings and borrowing. As I see it, the barriers are lack of communication across agencies and perceived legal barriers. When agencies work together, and legal restrictions are more carefully inspected, creative solutions can emerge.

For example, there were (and are) organizational challenges in simplifying the process for applying for federal student aid (Dynarski & Scott-Clayton, 2006, 2007; Dynarski & Wiederspan, 2012). Conceptually, simplification was straightforward. A key element of aid simplification was linking IRS data to the online aid application (the Free Application for Federal Student Aid [FAFSA]). The IRS resisted this linkage, citing restrictions to sharing data with other agencies. A creative solution that emerged from dozens of meetings allowed the linkage, enabling applicants to automatically populate their aid applications with their tax information.¹⁵

A similarly creative approach would allow for more timely and complete data on borrowing and earnings. More broadly, a commitment to linking the NCES longitudinal surveys to administrative data would broaden their scope, increase their accuracy, and enhance their usefulness to researchers and policy makers.

¹⁵ The workaround is that IRS (technically) provides data not to ED but to the applicant. While completing their FAFSA, applicants are prompted to log into the IRS servers. The IRS server passes the tax data to the applicant's web browser, and the applicant approves the transfer of the data from the browser to the aid application.

APPENDIX: LIST OF KEY DATA SOURCES FOR BORROWING AND EARNINGS

1. Administrative and Survey Data on Individual Borrowing
 - a. National Student Loan Data System (NSLDS): This is a census of all federal student loans, used by the U.S. Department of Education's (ED's) Federal Student Aid office to administer the loan program. It does not contain information on private loans. There is no research version of NSLDS and only limited summary statistics are released. No information on earnings is available in these data.
 - b. National Postsecondary Student Aid Survey (NPSAS): This is a nationally representative, longitudinal survey of college students run by ED's National Center for Education Statistics (NCES). Information about students' federal borrowing is drawn from NSLDS and merged onto NPSAS. For those who earn a BA, self-reported earnings data are collected for 10 years after graduation. For those who do not earn a BA, earnings data are collected for just six years after college entry.
 - c. Federal Reserve Bank of New York (FRBNY): The FRBNY recently created a panel of credit reports that includes information about student debt (Lee, 2013). It does not contain information about earnings.
2. Administrative Data on Individual Earnings¹⁶
 - a. Social Security Administration (SSA): SSA maintains longitudinally linked records of individual earnings. These records are used to compute Social Security benefits, which are a function of lifetime earnings. Researchers have successfully linked these data to surveys, including the Census (Angrist et al., 2011). ED has used these data to calculate median earnings of graduates from career-training programs (U.S. Department of Education, 2012). Since ED has already accessed these data for regulatory purposes, it is a particularly promising prospect for linking to ED data on loans.
 - b. Internal Revenue Service (IRS): IRS maintains household-level records of income-tax returns and the informational returns that are used in the calculation of taxes. These data include information on college attendance, in the form of the 1098-T, which colleges send to the IRS to document tuition payments. In recent years, versions of these data have become available to outside researchers (e.g., Chetty et al., 2011). Treasury employees can also conduct research with these data, and outside researchers have coauthored with them on studies (e.g., Manoli & Turner, 2014).
 - c. Unemployment Insurance Earnings Records: States maintain records of earnings that occur within each state. The Bureau of Labor Statistics pools these records in order to generate local employment and earnings estimates. In theory these data are available for researchers within the Census's Research Data Centers, but a requirement that states opt into each proposed research project has hobbled their use.

¹⁶ I limit this list to earnings data that follow workers across state lines. Every state maintains longitudinally linked, individual earnings records for the purposes of administering unemployment insurance and workers' compensation. However, these records do not follow workers across state lines and so do not provide a comprehensive profile of lifetime earnings (especially for the college educated, who are the most mobile).

REFERENCES

- Angrist, J. D., Chen, S. H., & Song, J. (2011). Long-term consequences of Vietnam-era conscription: New estimates using social security data. *American Economic Review*, 101(3), 334–338.
- Barrow, L., & Rouse, C. E. (2005). Does college still pay? *The Economists' Voice*, 2(4), Article 3.
- Chetty, R., Friedman, J. N., Hilger, N., Saez, E., Whitmore-Schanzenbach, D., & Yagan, D. (2011). How does your kindergarten classroom affect your earnings? Evidence from Project Star. *The Quarterly Journal of Economics*, 126(4), 1593–1660.
- College Board. (2012). *Trends in student aid 2012*. College Board Trends in Higher Education Series. Retrieved from <http://trends.collegeboard.org/sites/default/files/student-aid-2012-full-report-130201.pdf>.
- Dynarski, S. M., & Scott-Clayton, J. (2006). The cost of complexity in federal student aid: Lessons from optimal tax theory and behavioral economics. *National Tax Journal*, 59(2), 319–356.
- Dynarski, S. M., & Scott-Clayton, J. (2007). *College grants on a postcard: A proposal for simple and predictable federal student aid* (Hamilton Project Discussion Paper No. 2007-01). Retrieved from http://www.hamiltonproject.org/papers/college_grants_on_a_postcard_a_proposal_for_simple_and_predictable_fed/.
- Dynarski, S. M., & Wiederspan, M. (2012). Student aid simplification: Looking back and looking ahead. *National Tax Journal*, 65(1), 211–234.
- Dynarski, S. M., Hyman, J., & Schanzenbach, D. (2013). Experimental evidence on the effect of childhood investments on postsecondary attainment and degree completion. *Journal of Policy Analysis and Management*, 32(4), 692–717.
- Greenstone, M., Looney, A., Patashnik, J., & Yu, M. (2013). *Thirteen economic facts about social mobility and the role of education* (The Hamilton Project Policy Memo, June). Retrieved from http://www.hamiltonproject.org/files/downloads_and_links/THP_13EconFacts_FINAL.pdf.
- Kemple, J. (2008). *Career academies: Long-term impacts on labor market outcomes, educational attainment, and transitions to adulthood*. Unpublished manuscript.
- Lee, D. (2013, February 28). *Household debt and credit: Student debt*. Federal Reserve Bank of New York. Retrieved from <http://www.newyorkfed.org/newsevents/mediaadvisory/2013/Lee022813.pdf>.
- Manoli, D. S., & Turner, N. (2014). *Cash-on-hand and college enrollment: Evidence from population tax data and policy nonlinearities* (NBER Working Paper No. 19836).
- Turner, S. (2004). Going to college and finishing college. Explaining different educational outcomes. In *College choices: The economics of where to go, when to go, and how to pay for it* (pp. 13–62). Chicago: University of Chicago Press.
- U.S. Department of Education, Federal Student Aid. (2012). *Gainful employment operations manual*. Retrieved from <http://www.ifap.ed.gov/GainfulEmploymentOperationsManual/attachments/GainfulEmploymentOperationsManualMasterFile.pdf>.
- U.S. Department of Education, National Student Loan Data System. (2013). *Direct loan portfolio by loan status and Federal Family Education Loan portfolio by loan status*. Data retrieved October 2013 from <http://studentaid.ed.gov/sites/default/files/fsawg/datacenter/library/PortfoliobyLoanStatus.xls>.