Canadä



The North American Data Documentation Initiative Conference 2019

Benefits of Describing Statistical Production and Variables

April 24-26, 2019

Statistics Canada

Ottawa, Ontario



North American Data Documentation Initiative Conference

The North American Data Documentation Initiative Conference (NADDI) is an opportunity for those using metadata standards and those interested in learning more to come together and learn from each other. Modeled on the successful European DDI User Conference, NADDI is a two day conference with invited and contributed presentations, and should be of interest to both researchers and data professionals in the social sciences and other disciplines. A full day of training sessions precede the conference.

The theme for NADDI 2019, Benefits of Describing Statistical Production and Variables, emphasizes the benefits of using metadata to drive efficiencies in a research data lifecycle, as well as promote subsequent re-use of end data products.

Data Documentation Initiative (DDI)



The Data Documentation Initiative (DDI) is an open metadata standard for describing data and specifying data collection activities. DDI's principle goal is **making research metadata machine-actionable**. The specification can document and manage different stages of data lifecycles, such as conceptualization, collection, processing, analysis, distribution, discovery, repurposing, and archiving.

Host Site Location



Statistics Statistique Canada Canada

The host site of this year's conference is Statistics Canada, the national statistical office. The agency ensures Canadians have the key information on Canada's economy, society and environment that they require to function effectively as citizens and decision makers.

Wednesday, April 24, 2019 Workshops – Jean Talon Conference Room

8:00am – 9:00am	Coffee and Breakfast
9:00am — 12:00pm	Workshop 1 Workshop Chair: Stéphane Crête, Statistics Canada
	Introduction to DDI Instructor: Jane Fry, Carleton University
	Knowing what DDI is and understanding how it works and how to use it can be quite daunting – especially if you are not a programmer or coder. DDI can be a quite powerful metadata schema for the Social Sciences. But understanding the basics of it is a prerequisite to learning how to use it.
	This workshop will tell you about the background of DDI, the different tools used to interpret it and how it is integrated into a data lifecycle workflow. There will be exercises that help you to understand different facets of DDI thus allowing you to see its potential and power.
	This is not a technical workshop and is intended for those with no background in DDI. However, if you want a refresher in the 'whys' and 'wherefores' of it, then this workshop is also for you.
12:00pm – 1:00pm	Lunch Please note that lunch will not be provided by the host.
1:00pm – 4:00pm	Workshop 2
	Document Questionnaires and Datasets with DDI: a Hands-on Introduction with Colectica Instructors: Jeremy Iverson and Dan Smith, Colectica
	This workshop offers a hands-on, practical approach to creating and documenting both surveys and datasets with DDI and Colectica. Participants will build and field a DDI-driven survey using their own questions or samples provided in the workshop. They will then ingest, annotate, and publish DDI dataset descriptions using the collected survey data.
	Attendees may optionally bring their own Windows laptops to participate in the hands-on exercises.
6:00pm – 8:00pm	Informal Gathering
	Johnny Farina Restaurant 216 Elgin St, Ottawa, ON K2P 1L6



Thursday, April 25, 2019 Day 1 – Jean Talon Conference Room

8:00am - 9:00am Registration – Coffee and Breakfast 9:00am - 10:30am Welcome Statistics Canada Kathryn Stevenson, Statistics Canada DDI Alliance Welcome and Introduction of Keynote Jared Lyle, Director DDI Alliance, Inter-university Consortium for Political and Social Research Keynote - Enabling modernization of official statistics with metadata standards Anil Arora, Chief Statistician of Canada Mr. Arora has led significant transformational initiatives throughout his career, with experience and partnerships spanning all three levels of government, the private sector and international organizations, including the UN and the OECD. He has led projects on high-profile policy issues, legislative and regulatory reform, and overseen large national programs. Anil Arora was appointed Chief Statistician of Canada in September 2016. 10:30am – 11:00am Break 11:00am - 12:20pm Session 1 Session Chair: Dan Smith, Colectica **DDI Conference Evaluation Demo** Barry T Radler, University of Wisconsin-Madison Use of DDI to Document Annual Public-Use Microdata (PUMD) for the Consumer Expenditure Surveys Dan Gillman, US Bureau of Labor Statistics The US Bureau of Labor Statistics (BLS) is using DDI to document the Consumer Expenditure Surveys (CE). CE produces PUMD files on a yearly basis. The immediate goal in the development of the system is to allow users to connect PUMD variables to their questions. The current stand-alone database with the information about all the variables is incomplete, and the ability to track changes to codes, definitions, variables, and files over time and surveys is not straightforward. Our plans are to incorporate functions to view the variable / question mapping and track changes. In addition, we want to record methodological documentation such as questionnaire, weighting, and sampling designs. In the iteration of the metadata system under development now, we are mapping the variables from the 2017 PUMD files to their questions, which either wholly or in part provide the input to them. In this talk, we

describe the development effort with particular emphasis on the variable / question mapping. In many discussions about metadata models, the mapping between variables and questions is glossed over, with the implication being the mapping is straightforward. Our experience is different, and we describe the difficulties.

Use of DDI to support data curation in Dataverse

Amber Leahey, Scholars Portal, Ontario Council of University Libraries

Dataverse is an open source data repository platform developed by IQSS, Harvard University (with contributors and adopters from all around the world). Dataverse supports researchers with open and mediated data publishing, sharing, research promotion, and reuse of others data. The Dataverse platform makes it easy to deposit and publish research data, however, often deposited research data requires additional preparation and curation to be equipped for meaningful reuse. Data require accurate and rich documentation and metadata, without it data are meaningless. Standard metadata, such as DDI, are useful for describing data in Dataverse. This presentation will demonstrate how Dataverse supports the creation of metadata for improved data discovery, access, reuse, and preservation of research data in Dataverse.

Creating DDI-Lifecycle documentation for longitudinal data at ICPSR

Sanda Ionescu, Inter-university Consortium for Political and Social Research

In line with ICPSR's commitment to adopt DDI-Lifecycle for documenting some of its longitudinal studies, we have recently embarked on a new project that involves creating DDI 3.2 metadata for one of our most popular collections at the National Social Life, Health, and Aging Project (NSHAP), which is managed and distributed by our National Archive of Computerized Data on Aging (NACDA). We will briefly introduce this collection and will elaborate on the steps taken to move its metadata to DDI Lifecycle using Colectica Designer and make it publicly available on the Colectica Portal. We plan to focus on any specific challenges encountered during this process and the practical solutions applied to overcome them. We also intend to showcase some of the benefits of using DDI-L in terms of improved comparability and usability of the data.

12:20pm – 1:30pm Lunch

1:30pm – 3:00pm Session 2

Session Chair: Barry Radler, University of Wisconsin-Madison

Solve for X(ML): Transforming metadata to transform data access

John Huck, University of Alberta Libraries

At the University of Alberta, a subscription to data resources from the Linguistic Data Consortium (LDC) presented a cataloguing problem for library staff and access difficulties for users: resources were published in digital and physical formats, but the digital access was difficult to catalogue, and some datasets were only available in physical format. As a result, there was no single place to discover what datasets the



library had purchased, and accessing the data involved a different retrieval process for each format.

This presentation describes a project undertaken by staff from three library units to develop a simplified access model for LDC resources, due to be launched in 2019. LDC metadata available from an actively updated OAI-PMH host provided the basis for a transformation workflow to generate MARC records with custom XSLT stylesheets. This allowed the team to think about user goals rather than cataloguing workflows, and imagine an access model where users would not need to know anything about data format. In the new access model, each MARC record includes a link to a retrieval request form. Library staff pick the best format for retrieval and deliver the data electronically to authorized users.

metajelo+DDI: Deriving a metadata package for journals to support external linked objects from DDI

Lars Vilhuber, Labor Dynamics Institute, Cornell University

We propose a metadata package (called metajelo) that is intended to provide academic journals with a lightweight means of registering, at the time of publication, the existence and disposition of supplementary materials. Information about the supplementary materials is, in most cases, critical for the reproducibility and replicability of scholarly results. In many instances, these materials are curated by a third party, which may or may not follow developing standards for the identification and description of those materials. Researchers struggle when attempting to fully comply with data documentation and provenance documentation standards.

However, many of the required elements are present in DDI, and when properly populated by data custodians using DDI, generation of the metajelo package is straightforward. In this presentation, we describe the rationale behind metajelo, and how archives that use DDI can easily provide researchers with a compact metadata package that enhances reproducibility while reducing researcher effort.

Complex sampling design settings: bridging technical documentation and applied researcher software use

Stanislav Kolenikov, Abt Associates

This work documents our understanding of, and best practice recommendations for specifying complex sampling design settings in the statistical software used for design-based survey data analyses. We discuss complex survey data features such as stratification, clustering, unequal probabilities of selection and calibration, outline their impact on estimation procedures, *and demonstrate the dangers of incorrect analysis that ignores them*. We show how statistical software treats these features. We analyze the documentation that accompanies several surveys collected and distributed by the national statistical offices, statistical agencies or academic organizations, and rate the ease or difficulty of specifying the sampling design correctly in the software based on the dataset documentation provided. We conclude with our recommendations to survey data providers regarding how to create the transparent, accurate and efficient documentation required by survey

data users to account for sampling design features in their analyses of complex survey data. This work is co-authored with Brady T. West (University of Michigan) and Peter Lugtig (University of Utrecht).

Using DDI and metadata in research analyses, reporting and dissemination: use case for the development of risk prediction algorithms in health Doug Manuel, Statistics Canada

Despite a strong imperative for open and reproducible research, DDI metadata rarely survives to research dissemination.

We describe a workflow for DDI metadata for research analyses and dissemination using the use case of predictive algorithms using Statistics Canada's Canadian Community Health Survey linked to longitudinal mortality and disease registries (CCHS linked). The workflow starts with a with a DDI-compliant worksheet that allows investigators to collaborate across institutes to pre-specify predictive model parameters. Routine functions clean and transform the study cohort according to the worksheet specifications, adding metadata to the study cohort DDI document. DDI metadata is then re-used in a range of dissemination settings, including manuscript tables and web-based algorithm visualization tools. DDI document is published. Models are reported using Predictive Modelling Markup Language (PMML), a complementary XML schema. Algorithms are deployed using DDI, PMML and TensorFlow scoring engines.

3:00pm – 3:15pm Break

3:15pm – 4:15pm Session 3

Session Chair: James Doiron, University of Alberta Libraries

Development of a workflow for capturing, enhancing and disseminating longitudinal study metadata

Jon Johnson, UK Data Archive

CLOSER brings together eight world-leading UK longitudinal studies in order to maximise their use, value and impact. A major output has been CLOSER Discovery which allows users to search and browse questionnaire and dataset metadata. Efficient data management of complex longitudinal studies is both desirable and increasingly essential. Metadata standards are critical for maintaining this information through the data life-cycle. Many long term studies face an historical backlog which has prevented them from moving to the best metadata standards. The CLOSER project has received significant funding from the ESRC and MRC for these eight studies and provide the basis for a sustainable high quality resource for research.

CLOSER has developed a suite of tools and software using both in-house and commercially available solutions that begin to tackle come of the obstacles involved in documenting and utilising longitudinal metadata. The presentation will report on



the successes and problems faced in using the DDI-Lifecycle metadata standard to
achieve these ambitions. This will be delivered as three presentations covering:

- questionnaire metadata capture and annotation; (Charles De Barros, CLOSER)
- mapping of questions, variables, keywords and concepts and; (Charles De Barros, CLOSER)
- demonstration of the latest functionality available for the research community. (Jon Johnson, CLOSER)

Using DDI to document a complex longitudinal study

Barry T Radler, University of Wisconsin-Madison

	Comprehensive research metadata greatly clarify the methods and processes used to capture data and produce datasets In so doing, it provides data users the information needed to better analyze, interpret, preserve, and share them. Richly structured metadata are even more critical with complex longitudinal studies that contain thousands of variables, different data types, and many waves of data collection. MIDUS (Midlife in the United States) is a national longitudinal study of approximately 12,000 Americans that examines aging as an integrated bio-psycho- social process. MIDUS has a broad and unique blend of survey, experimental, and laboratory research data collected over 20 years through a variety of modes. For the last decade, MIDUS has increasingly relied on the Data Documentation Initiative (DDI) to manage these complex research data and make their metadata machine actionable in a user-friendly format. More recently, MIDUS has used Colectica tools to improve its DDI infrastructure and create a DDI-based harmonized data extraction system. Such a system allows researchers to search across datasets for variables of interest, identify and harmonize related longitudinal versions of variables, and easily create customized data exports.
4:15pm – 4:45pm	Session 4 Session Chair: James Doiron, University of Alberta Libraries
	DDI Community Engagement: Discussion & Feedback William Block, Cornell Institute for Social and Economic Research, Cornell University Barry T Radler, University of Wisconsin-Madison Jared Lyle, Director DDI Alliance, Inter-university Consortium for Political and Social Research
4:45pm – 5:00pm	Wrap-Up James Doiron, University of Alberta Libraries
7:00pm – 9:00pm	Banquet
	The Bier Markt Ottawa 156 Sparks Street, Ottawa, Ontario K1P 5C1

Friday, April 26, 2019 Day 2 – Jean Talon Conference Room

 8:00am – 9:00am
8:30am – 9:30am
Poster Session Session Chair: Flavio Rizzolo, Statistics Canada

Dataverse North Metadata Best Practices Guide

Amber Leahey, Alexandra Cooper, Martine Gagnon, et al.

One of the benefits of using Dataverse is that it provides an expansive set of metadata for describing research data at the project and dataset-levels. Using a combination of Dublin Core and the Data Documentation Initiative (DDI) standard at its core, Dataverse provides a metadata set that is well-suited to the social, behavioral, economic, and health sciences. However not all researchers may know which metadata fields to use or how best to interpret them. This poster will showcase the Dataverse North Metadata Best Practices Guide – a primer for novice to intermediate users of the Dataverse data repository platform that provides direction about how to describe research datasets using the standard metadata fields in Dataverse.

A Practical Look at the DDI 4 Methodology Pattern

Dan Gillman, US Bureau of Labor Statistics

The Methodology Pattern in DDI 4 is a general structure designed to account for descriptions of study methodology at any level of detail. The pattern accounts for two independent considerations: 1) it separates the design (the what) from the algorithm (the how - in theory) from the process (the how - in practice); and 2) it allows a recursive description into ever more detailed sections of the processing cascade. Separating design, algorithm, and process supports reuse and provides a rich descriptive range. For instance, the difference between an algorithm with an unspecified stopping criterion and a process with a specific bound provided is clearly differentiated. Also, a design is clearly differentiated from the algorithm. Many algorithms can satisfy the same basic design. Through the example of the US Consumer Expenditure Surveys (CE), we illustrate the multi-layered approach. CE accomplishes its processing through 4 discrete phases or sub-systems. Each can be described in general at a high level, at a mid-level by generally deserving each subprocess, and the low level by providing step by step details of each sub-process. We also illustrate why the lowest level descriptions might violate confidentiality requirements.

Keeping metadata alive during health research analyses and reporting: using an Open Science imperative to update the current metadata-sparse analysis approach Doug Manuel, Statistics Canada



Mapping DDI 2 to DDI 4

Larry Hoyle, Institute for Policy and Social Research, University of Kansas Joachim Wackerow, GESIS - Leibniz Institute for the Social Sciences

This poster describes the effort to add a DDI-Codebook (DDI-C) import function into the DDI 4 R package.

The DDI 4 Codebook Group did a lot of the modeling of one section of DDI 4 using a spreadsheet mapping DDI-C elements into DDI 4 properties. This started with a list of elements used by CESSDA and was refined at the May 2016 Knutholmen Sprint. Unfortunately, these mappings were not always at the leaf node level.

An R program also imported DDI-C XML from the European Social Survey and generated a list of unique XPaths of leaf elements used in that set of metadata. These elements, along with corresponding DDI 4 leaf paths, were used to update the spreadsheet.

This spreadsheet has been further refined to create an actionable table mapping DDI-C leaf values to leaf properties in DDI 4. Writing code to import the DDI-C required additional information:

- mapping from DDI-C sub-paths to DDI 4 Identifiable classes (e.g. all the information for one DDI-C "var" maps to one DDI 4 IdentifiableVariable),
- mapping abstract target classes to specific extensions,
- additional semantic property values like "typeOfMethodology"

Importing DDI-C into a lifecycle level version of DDI like DDI 4 also involves identifying repeated metadata like reused value domains (e.g. reused Likert style codelists) that are repeated for multiple variables. An R function served to do this sort of matching using the R "all.equal" function excluding differences in agency, id, and version.

The CESSDA Vocabulary Service: a new state-of-the-art tool for creating and publishing controlled terms lists

Sanda Ionescu, Inter-university Consortium for Political and Social Research Taina Jaaskelainen, Finnish Social Science Data Archive

The DDI Alliance has been creating and publishing its own controlled vocabularies since 2005. These are targeted for specific DDI classes, but are external to DDI and may be used with other metadata standards.

In its new Core Metadata Model CESSDA recommends the use of DDI Controlled Vocabularies where available, as well as the creation of new lists as needed. To support this effort, CESSDA has financed the development of a state-of-the art tool that facilitates the creation, translation, and publication of controlled vocabularies, and automates a significant part of the process. The tool is now in its final testing phase and is scheduled to go live this spring.

Our poster will introduce this new, Web-based tool in the context of our ongoing Controlled Vocabularies work, and will highlight some of its most prominent features, like cross-vocabulary searches, group sharing and editing, support for multiple language translations, automated versioning, one-click publishing, multipleformat downloads, etc.

Leveraging language codes in a style sheet transformation: OLAC (ISO 639-3) into $\ensuremath{\mathsf{MARC}}$

John Huck, University of Alberta Libraries

Being able to generate MARC records through a stylesheet transformation was a necessary element of a recently introduced solution for the discovery and access of a collection of linguistic data resources at University of Alberta Libraries. This poster shows how ISO 639-3 language codes in the source metadata were transformed into LC subject headings and MARC language encodings for approximately 750 records.

Crosswalk 4.0 adds DDI 3.2

Florio Arguillas Jr and William Block, Cornell Institute for Social and Economic Research

At last year's NADDI in Washington, DC, Crosswalk 3.0 showcased how it resurrects near dead data --those ASCII datasets with no accompanying setup files, but only hardcopy of the codebook-, back to life by simplifying the process of creating setup files and datasets just by using an Excel spreadsheet containing variable information and location, and the source ASCII dataset as inputs. Along with creating SAS, SPSS, STATA setup files and datasets, Crosswalk 3.0 also outputs DDI 2.5 Codebook. This year we are showcasing Crosswalk 4.0, the latest version of the software, in a poster session. In addition to the features present in 3.0, we have added the capabilities to create R setup files and R Workspace and DDI 3.2 Lifecycle.

United States Census Bureau Data Repository

Jared Lyle, Inter-university Consortium for Political and Social Research

We will provide an overview of the United States Census Bureau Data Repository, which preserves and disseminates survey instruments, specifications, data dictionaries, codebooks, and other materials provided by the U.S. Census Bureau. ICPSR, the host of this data repository, has also listed additional Census-related data collections from its larger holdings. We will also highlight Census-produced DDI content used in archiving collections within the repository.

9:35am – 10:45am Session 5 Session Chair: Amber Leahey, Scholars Portal



Implementing DDI 3.3 support in Colectica version 6

Dan Smith, Colectica

Colectica version 6 will include many new content areas of DDI Lifecycle 3.3, and will adopt many of the new features while still maintaining compatibility with all prior versions of the standard.

This talk will review the new content available in DDI 3.3 and how it is being added to Colectica version 6 in the following areas: Classification management based on GSIM/Neuchatel and Colectica's migration from the Copenhagen-mapping to DDI 3.3, Non-survey data collection measurements being added to Colectica Questionnaires, Sampling and Weighting, Questionnaire Design, Support for working with DDI as a Property Graph with properties on items and references via both DDI 3.3 and Colectica's RDF Services for Colectica Repository, and Quality Statement improvements. In addition, the talk will highlight the now-complete formal model that is behind DDI 3.3 and how it has been useful for software development, and how Colectica approaches DDI implementation to ensure backwards compatibility.

DDI 4 in R: New possibilities

Larry Hoyle, Institute for Policy and Social Research, University of Kansas Joachim Wackerow, GESIS, Leibniz Institute for the Social Sciences

We have been working on a representation of the DDI 4 model in R, realized as a package. We now have an R object oriented class for each DDI 4 class with associated functions to validate and print objects, manage a registry of DDI 4 objects, manage DDI 4 URNs, and import & export DDI 4 XML.

Our original goal was to enhance the ability of researchers to capture and report on metadata at the source, with the ability to embed references to metadata.

In this presentation we'll discuss the intriguing prospect of computing directly on the metadata. What could be done with these metadata objects to facilitate comparison and harmonization? In what ways could the metadata be visualized?

DDI 4 has powerful new capabilities in the Collections pattern. For classes realizing a collection, operators could be defined to return their intersection, union, and difference. Inner and outer joins could also be defined. Relationships within the collection could be visualized via network diagrams. These operators might provide efficient tools for harmonization.

Operators could also be defined on pairs of objects of the same class. Similarity measures could be computed via corresponding attributes. These could be used to create visualizations of, say, similarities among variables.

The Picasso Project

Office and Privacy Management and Information Coordination, Statistics Canada

Picasso is an enterprise solution for statistical data and metadata management. Automated business rules will ensure metadata is gathered uniformly, adhering to common architecture, governance and policy instruments.

Picasso replaces local solutions with a single hub for managing metadata for all surveys, administrative files and record linkage projects. It's a single point of access for all 'fit for use' data files, and an enterprise search and discovery engine using metadata to facilitate reuse of information. It provides a user-friendly interface with an automated workflow and integrated lifecycle management. Picasso puts the power and potential of information together in a single platform to promote reuse of data and statistical metadata to support statistical production and analysis. Its well-organized registry and repository, and its user-friendly tools, are designed to support the rapid growth in data assets and the shift to an administrative data-first approach. Information management principles are embedded into the automated workflow to ensure good stewardship and to make it easier for all employees to search, access, reuse, manage, and report on data and statistical metadata.

Picasso aligns with reference models, e.g. GSIM, GSBPM, and will support standard vocabularies, e.g. DDI, SDMX, DCAT, SKOS/XKOS, PROV-O, in a variety of formats, e.g. JSON, XML, RDF.

10:45am – 11:00am Break

11:00am – 12:15pm Session 6

Session Chair: Bill Block, Cornell University

Documenting Variable Provenance with DDI and Colectica Jeremy Iverson, Colectica

DDI supports documenting data with variable-level detail. Normally such information includes a variable's name, label, and data type, but DDI supports including more details. This can include the variable's lineage. The data behind a variable may have originated from a survey, from administrative data, or from some other dataset. The variable may have been copied directly from a source, or it may have been calculated using some formula or algorithm. For a researcher to properly understand and analyze the data, the lineage of each variable should be documented with as much detail as possible.

Variable-level provenance information can be recorded by data management staff manually. However, in some cases it is also possible to extract the transformations that were used to create a variable based on statistical source code. The C2Metadata project, sponsored by the US National Science Foundation, is building tools to accomplish this, and to record the transformations in structured ways that can be included with data documentation.



This presentation will show data documentation from real studies that are providing
variable-level provenance information, and will discuss the methods used to create
and publish the documentation.

DDI Compliance Benefit for a Generic Statistical Business Process Model

Microdata Access Division, Statistics Canada

The Microdata Access Division (MAD) is responsible for providing access to microdata to researchers outside Statistics Canada. They also provide support, expertise and standards of data provision to subject-matter areas within Statistics Canada.

The Microdata Access Division will showcase the virtual Data Lab (vDL) project and how the division would stand to benefit from a Generic Statistical Business Process Model (GSBPM) compliant to DDI, from design to dissemination.

Labour Force Survey: Improved Access for Researchers

Susan Mowers, University of Ottawa Libraries

The Canadian Labour Force Survey (LFS) is a monthly survey produced by Statistics Canada. It produces a number of data products used by researchers across the country. The monthly Public Use Microdata Files (PUMF) that are distributed via the Data Liberation Initiative (DLI) in coordination with data services provided by university libraries, requires accuracy and consistency of metadata in DDI format for access online. This presentation discusses the process for marking up the LFS across the different series of data files (1976-Current), in both English and French, and, across different metadata authorship from different institutions over the years.

- 12:15pm 1:30pm Lunch
- 1:30pm 2:30pm Session 7

Closing Plenary – Introduction: James Doiron, University of Alberta Libraries Jeff Moon, Director, Canadian Association of Research Libraries Portage Network

2:30pm – 3:00pm Wrap-Up and Evaluation

DDI Conference Evaluation Results Barry T Radley, University of Wisconsin-Madison

Wrap-Up

Jared Lyle, Director DDI Alliance, Inter-university Consortium for Political and Social Research Cory Chobanik, Statistics Canada



Sponsors and Partners

Platinum



Gold





ICPSR

Silver









CISER CORNELL INSTITUTE for Social and Economic Research





