User's Manual

UCLA-LoPucki Bankruptcy Research Database

Lynn M. LoPucki

I. Concept.

The UCLA-LoPucki Bankruptcy Research Database (BRD) is a dataset consisting of approximately 200 fields of data on each of the approximately one thousand large, public company bankruptcies filed in the United States Bankruptcy Courts since October 1, 1979. The purpose of the BRD is to promote research by providing a statistical-analysis-ready core of basic data.

The BRD includes data all large, public company bankruptcies. A company is considered "public" if it filed an Annual Report (form 10-K or form 10) with the Securities and Exchange Commission for a year ending not less than three years prior to the filing of the bankruptcy case. A company is considered "large" if that Annual Report reported assets worth $100 million or more, measured in 1980 dollars (about $314 million in current dollars). Coverage includes cases filed under Chapter 7 and Chapter 11, whether filed by the debtors or creditors.

The BRD contains several fields designed facilitate its use along with data from other sources. They include several standard company identifiers and bankruptcy court case numbers.

We keep the BRD in Microsoft Access, update it monthly, and distribute it in Microsoft Excel (.xslx) and Comma Separated Values (.csv) formats. Along with the data files, we include Stata Do Files that convert the CSV file to a Stata Dat file and run some sample regressions. Stata users can generate research results almost immediately.

The BRD's mission is to promote big case bankruptcy research that leads to publication. To accomplish that, we make an abbreviated version of the BRD available to the public on line, free of charge, along with online software that facilitates research. We license the Cases table for both commercial and academic use. To facilitate collaboration among researchers, we require that academic users provide brief (about two to ten word) descriptions of their research projects. We post the descriptions on our website. We also post citations to the growing body of research that is based on BRD data.

Financial support for the BRD has come from the UCLA School of Law, grants from the National Conference of Bankruptcy Judges Endowment, the American Bankruptcy Institute's Anthony H.N. Schnelling ABI Endowment Fund, the Turnaround Management Association, and private gifts.

II. Getting Started

A. Access to the online database.

An abbreviated version of the BRD (twenty-six fields) is available on the BRD website. http://lopucki.law.ucla.edu/. The data on all BRD cases are available in three formats: (1) by company, (2) by interactive spreadsheet, and (3) in user-customized lists and statistics (Design a Study).

1. By company. The website produces a page showing the seventeen items of data available on each BRD case. Users seeking specific facts about particular cases will find this format easiest to use. Users can reach it through the View Data and By Company tabs.

2. By interactive spreadsheet. Twenty-six fields of data are available by interactive spreadsheet. Users can choose which fields to display and order the data by any field. Users can reach the interactive spreadsheet through the View Data and BRD Spreadsheet tabs.

3. Design a Study. Up to twenty-six fields of data can be used to design customized studies. The user selects the types of cases to be included in the study and the basis for case aggregation. The program runs the study and reports case names, data for the aggregation field, and some basic statistics. The statistics are frequency for each category of the aggregation variable and means, medians, and standard deviations for each category of the duration, assets, employees, and sales variables. Users can reach Design a Study through the Query Data tab.

B. The BRD Files

We distribute the BRD in a zip file as an attachment to an email. The zip file contains these seven files:

1. BRD Import.do
2. BRD Labels.do
3. Protocols.pdf
4. Read Me First.pdf
5. UCLA-LoPucki Bankruptcy Research Database.csv
6. UCLA-LoPucki Bankruptcy Research Database.xlsx
7. User's Manual.pdf

1. Read Me First

The BRD is set up to run with Stata. If you have Stata 10 through 13 installed, these simple steps will enable you to begin running regressions:

1. Make an empty directory in which to work.
2. Click on the zip file to open it. You will see six files.
3. Choose "Extract all files" and point to the directory in which you will work.
4. Open BRD Import.do in the Stata do file editor window (on the Windows menu).
5. Change the path in the line that begins "cd" to point to the directory in which you plan to work.

6. Change the file name in the line that begins "insheet" to match the name of the csv file.
7. On the Tools menu, execute "Do" or click on BRD Import.do. Either will run both do files.
8. Click on –more – as necessary to see all the output.

If you are running Stata 13, the run will end with several regression tables; if you are running Stata 10, the run will end with one regression table and the error message iDE##c: operator invalid.

2. The Data Files

The zip file contains two copies of the UCLA-LoPucki Bankruptcy Research Database. One is an .xlsx file, the other is a .csv file. Humans tend to prefer the .xlsx file; machines tend to prefer the .csv file.

The preference is in part because numbers generally appear in the .xlsx file as integers with commas. Numbers generally appear in the .cvs file with full precision and without commas. In the .xlsx file, linked data are actually linked and appear in blue with underlining. In the .cvs file, linked data are not actually linked and appear in the format [data]#[URL].

3. The Protocols File

The BRD protocols are rules we wrote to guide our data collection and assure the comparability of data across cases. They describe in detail the standard for inclusion of cases in the BRD, the definition of the data in each field, the rules for categorization of categorical data, and, in some instances, descriptions of our data collection methods. To fully understand the data in any field, it is essential to consult the protocol for that field.

The first page of the Protocols file contains a Variables List. The list contains the names of all BRD variables in the order in which they appear in the .xlsx and .csv Excel files. With one exception, the order is alphabetical. The exception is that NameCorp – the field containing the name of the bankrupt company – appears at the beginning of the list.

We have tried to make the field names short and intuitive. To make them intuitive we concatenated words that described their contents. To make them short, we truncated the words. We used capital letters indicate the breaks between truncated words. To keep related fields together on the alphabetically organized Protocols and the Variables List we began the names of related fields with the same truncated word. Thus, for example, all liabilities fields begin with "Liab."

If you are already familiar with the language of big case bankruptcy, you can probably guess from the field names what is included in each field. If you can, use the Variables List to get an overview of the BRD. Otherwise, use the Protocols.

4. The Do Files

The BRD Import Do File imports the UCLA-LoPucki Bankruptcy Research Database from the .csv file to Stata, prepares the variables for analysis, and provides some examples of regressions based on the data. For the Do File to work, the user must point the Do File at the directory in which the csv file is located (see line 5 of the Do File).

The Do File makes these changes to the data to prepare them for regression and other statistical analysis:

1. Drop Chapter-7-at-filing cases. The Do File drops approximately 19 cases that were filed as Chapter 7s. These cases do not share the characteristics of cases filed under Chapter 11, so most data on these cases are missing.

2. Date variables. Date fields are imported from the BRD as strings. The Do File creates corresponding fields that contain the same dates in Stata date formats. The data appear the same in Stata, but users can set date ranges and subtract dates from other dates to determine the numbers of days in between.

3. Identify missing date data. In a few date fields, the BRD stores the fact that data are missing by recording values outside the relevant range. Missing dates are entered as "9/9/9999."   The Do File changes those values to "." indicating that the values are missing.

4. Identify missing numerical data. In a few numerical fields, the BRD stores the fact that data are missing by recordings values outside the relevant range. For example, missing bond prices are entered as "999999." The Do File changes those values to "." indicating that the values are missing.

5. Identify missing alphanumeric data. In many fields, the BRD stores the fact that data are missing by indicating the reason. For example, missing data on whether a company emerged from bankruptcy may be indicated as "pending" if the case has not completed, "not classified" if the case was dismissed without confirmation, 363 sale, or conversion to Chapter 7, and "No data" if we were unable to obtain data regarding emergence. The Do File creates a new variable in which all of these values to "missing."

6. Generate categorical data. The Do File converts categorical data to dummy variables if the category is sufficiently large to be useful in statistical analysis. For example, the Do File creates categorical variables for the most common SIC codes, DIP lead bankruptcy attorneys, turnaround management firms, and judges at confirmation.

7. Generate change variables.   In many instances, the BRD contains measures of a single variable at different stages of the case. For example, it contains assets before filing, at filing, and after emergence, and it contains the prime rate of interest two years before filing, one year before filing, at filing, at disposition, one year after disposition, and two years after disposition. The Do File generates "change" variables that represent the change in a variable between two dates. If a variable should be logged prior to use, the change variable is expressed as either the difference in the two logs or the log of the ratios. It is a difference in ratios. The names of change variables contain the string "change."

8. Shift zeros for logging. The log of zero is undefined. Stata will set it to "missing." In order to retain observations that are zero, the Do File adds an immaterial increment (1, .01, or .001). The increment depends on the distribution of the variable. The increment is sufficiently small that the resulting shift will not meaningfully alter the results. The names of variables adjusted in this manner end in "1."

9. Generate logged variables. Some variables are distributed lognormally. To use them in analyzes that assume normally distributed data (t-tests, regression) these data should be

converted to natural logs. We have identified several of these variables, and the Do File performs the conversion. The names of logged variables begin with the string "ln."

10. Identify observations in representative ranges. The BRD is a universe, not a sample. Some variables however, are available or sufficiently complete to be representative only within particular date ranges. For example, turnaround manager and DIP loan data are representative for the period since November 30, 2003. Refiled-within-five-years-after-confirmation data are representative only for the period before January 1, 2008. But the BRD also contains anecdotal data for some cases outside the range. F or example, some turnaround manager data are available for the period before November 30, 2003 and many refilings made after January 1, 2008 are included. Because they are not representative, the out-of-range data should be excluded from statistical analyses. To facilitate their exclusion, the Do File creates two kinds of variables. Weight variables are coded 1 for observations in the representative range and 0 for observations outside it. A weight variable can be used as a frequency weight in regression by adding tab *var1 var2* [fw=wgtvariablename] at the end. Assignment of a zero weight excludes the out-of-range observations from the analysis. The names of Weight variables begin with "wgt." Reprange variables are dummy variables that include only observations within the relevant date ranges. Some users will find them easier and more intuitive than Weight variables. The names of reprange variables contain the string "reprange."

11. Variable labels. At the end, the Do File calls another Do File (BRD Labels.do) that assigns labels to each of the variables imported or created.

C. Sample Selection

Sample selection can be done manually or automatically. Automatic sample selection occurs when a researcher specifies a set of variables for regression. The statistics program includes only observations (here BRD cases) for which data are available for all of the independent variables. The program will drop all other observations.

Automatic sample selection is problematic when BRD data are available but not representative. For example, some BRD turnaround manager field data are available for cases filed prior to December 1, 2003, but they are few and may not be representative. Researchers should consider excluding from their studies all observations with non-representative turnaround manager data. Researchers can easily do that using the TurnRepRange field.

In some numeric or date BRD fields, we indicate that data are missing by entering data outside the plausible range for the field. For example, a missing date is indicated by entering "9/9/9999" in the date field. A missing GVKEY is indicated "999999." Our Stata Do File changes these entries to missing (.). If you do not use our Stata Do File, you must make the changes in some other way before selecting your sample.

When selecting samples from the BRD based on date range, it is usually preferable to select them based on date of case disposition rather than case filing. A sample that includes all cases filed in recent years will include at least some pending cases. Pending cases are non-representative in many respects simply because they are not complete. For example in a study of

case duration, the cases filed recently and disposed of quickly will have duration data. The cases filed recently and not disposed of – the longer cases – will not have duration data. The former will be included in the study; the latter will be excluded from the study. The effect will be to bias the study toward a shorter estimate of case duration.   A similar problem will occur in a study to determine the prevalence of 363 sales, because (1) 363 sale cases are of greater average duration than non-363 sales, and (2) we do not collect 363 sale data systematically until the case is disposed of.

D. Simple Things You Can Do with the BRD

1. Find a case on Pacer.

The BRD contains court identifiers and case numbers for all large, public company bankruptcies filed in the United States since October 1, 1979. The data include case transfers. Find the case in the BRD and you have the information you need to go directly to it on Pacer.

2. Find confirmed plans on Pacer.

The BRD contains the docket number for the confirmation order. Go to that docket number, and you will usually find a link to the confirmed plan. Even if the court hasn't provided a link, the confirmation order identifies the confirmed plan unambiguously – assuring that you won't waste your time analyzing the wrong plan.

3. Make a list of cases that meet your criteria.

Want a list of all 363 sale cases? All prepacks in the Southern District of New York? The largest cases by asset size? You can generate standard lists in the online version of the database. You can also generate customized lists using the "filter" and "sort" functions on Excel.

4. Make tables and graphs showing trends or comparing cases.

Once you have lists of cases, you can count them and use Excel to make tables and graphs. Those tables and graphs can show trends or compare cases across districts, cases sizes, plan types, and many other criteria. We have posted some samples of tables and graphs made from the BRD on the BRD website.

5. Empirical studies.

For those with graphing and statistical skills, BRD data make publishable empirical studies quick and easy to do. You can do instant empirical studies on the BRD website; more sophisticated studies using the Excel version of the BRD; and even more sophisticated studies if you are willing to collect relevant data not already contained in the BRD.

III. Fields

A. Field naming protocols

The BRD contains more than 180 variables. The Do File creates hundreds more. To make the variables easier to use, we have followed naming protocols that enable users to intuit the nature of a variable from its name. These are the protocols:

1. For BRD variables.

a. Concatenations. BRD variables are concatenations of the first syllables of words that describe the variable. For example "AssetsBefore" is the amount of the debtor's assets on the last form 10-K filed by the debtor prior to bankruptcy. "BondPriceFile" is the trading price of the debtor's bonds at the time of the bankruptcy filing.

b. Capitalization. The first letter of each word is capitalized. When working in Stata, capitalization indicates that the variable was imported from the BRD, not generated in Stata. Variable generated by our Stata Do File do not contain capital letters.

d. Date fields. Regardless of subject, the names of all date fields begin with the syllable "Date."

e. Variable subjects. The names of variables related to a subject generally begin with the same syllable. Thus, variables indicating the distributions to stakeholders begin with "Distrib," gross domestic product variables begin with "Gdp," and variables indicating the numbers of employees in various categories begin with "Empl." One advantage is that variables related to a subject appear together when the fields are listed alphabetically.

2. For Stata variables.

a. Lower case. The names of variables created by our Stata Do File are entirely lower case.

b. Same name. Aside from case, the names of variables created by the Do File are the same as the names of the corresponding variables imported from the BRD. For example, bondpricefile is a Stata variable that is similar to BondPriceFile as imported from the BRD. (In this instance, the difference is that "999999" in BondPriceFile is "." (missing) in bondpricefile.)

c. Shifted variables. A shifted variable is a numerical variable the values of which have been shifted by adding a nominal amount. We shift variables to eliminate zeros and thus make it possible to log them. The name of a shifted variable ends in "1." See above.

d. Logged variables. The name of a logged variable is the name of the variable logged, preceded by "ln."

e. Change variables. If a variable measures a change over time, the variable name begins with "change." The protocol for logged variable overrides this protocol, so that the name of a changed variable that has been logged begins with "lnchange." See above.

f. Weight variables. Weight variables begin with "wgt." See above.

g. Reprange variables.   Reprange variables contain the string "reprange." See above.

h. Label definitions and label values. We have not yet created either, but eventually intend to create both.

B. BRD Field Groups

Often, no single field can capture an entire aspect of a company, case, outcome, or economic condition. For example, the size of a company or the length of a case can be measured in different ways, no single one of which is best. Nor can a single field reflect all relevant aspects of a single phenomenon, such as a DIP loan. One field is necessary to identify the lender, another to record the amount, a third to record the date on which loan was made. As a result, many fields occur as field groups. This section identifies some of the BRD field groups and indicates some of the fields that are in them. Some fields are part of more than one group.

1. Company and Security Identifiers.

Company and security identifiers enable users to link BRD data to data from other sources. The links are crucial because they enable users to check BRD data against other sources and to import data from other sources for use with BRD data.

a. Company name. Company names frequently change over time and many potential data sources contain no identifier other than company name. The BRD includes separate fields containing the companies' names at filing, upon emergence, and after emergence.

b. Coded company identifiers. Companies are commonly identified by CIK numbers (SEC), GVKEYs (Standard & Poor's), and CUSIP numbers. Although these identifiers are intended to remain the same for the duration of the company's existence, some change at the time of bankruptcy reorganization. Accordingly, the BRD contains separate identifiers for the filing and emerging companies.

c. Coded security identifiers. CUSIP extension and bond ISINs identify securities. At reorganization, securities are frequently cancelled and other securities issued. At present, the BRD contains CUSIP extension and bond ISIN fields only for the filing companies.

2. Case Identifiers.

Case number plus court or court city uniquely identifies a bankruptcy case. The BRD contains those fields of data for (1) each case that qualifies for inclusion in the BRD, (2) that same case if continued in another court after transfer, and (3) the first refiling by the emerging company.

3. Company characteristics.

The BRD contains six categories of variables that measure the financial characteristics or condition of the debtor at the time the case was filed: leverage, profits and revenues, bankruptcy cause, financial structure, company size, and industry.

a. Leverage. The BRD contains two measures of leverage at filing. AssetsBefore and LiabBefore are the debtor's total assets and total liabilities from the debtor's last 10-K prior to filing. AssetsPetition and LiabPetition are the debtor's total assets and total liabilities from Exhibit A to the debtor's bankruptcy petition.

b. Profits and revenues. The BRD contains four measures of the debtor's profitability during the last year reported prior to bankruptcy. They are net income, EBITDA, EBIT, and net income before extraordinary items (IncomeBEBefore). The BRD also contains a set of "sales" variables that measure the debtors' revenues in the last year reported prior to bankruptcy.

c. Cause. The TortCause field indicates whether the debtor filed bankruptcy principally as the result of tort liability, including fraud perpetrated against the debtor by insiders. The variable includes mutually exclusive categories for fraud, pension, patent, asbestos, other products liability, environmental liability, and other torts. Each of these categories – or combinations of these categories – can be converted to dummy variables and included as independent variables in success models. Alternatively, these categories can be used to exclude cases from, or include cases in, the regression analyses.

d. Financial structure. The BRD contains variables that measure the amounts of prepetition secured and unsecured claims in the debtor's prepetition financial structure, as reported in the disclosure statements.

e. Size. The BRD contains variables that measure company size at or before filing four ways: total assets, total liabilities, revenues, and employees.

f. Industry. The BRD contains fields for each of the four levels of Standard Industrial Classification (SIC) Codes. Researchers can convert any industrial classification – or group of classifications – that contains a sufficient number of cases into a dummy variable and include it in the regression models.

4. Case characteristics.

The BRD contains fifteen categories of variables that measure case characteristics: duration, CEO turnover, 363 sale of all or substantially all asses, case cost, case intensity, court, case type, trustees, examiners, committees, unions, attorneys, distributions, substantive consolidation, and fee reviewers.

a. Duration. The duration of a case can be measured in a variety of ways. To maximize the researcher's option, the BRD supplies the dates of several key events: (1) the filing of the case, (2) the date the court-approved sale of the company under section 363, (3) the date the court entered the plan confirmation order, (4) the date the court dismissed the case or

converted it to chapter 7, (5) the effective date of the plan, and (6) the date, if any, on which the debtor refiled bankruptcy. By subtracting the date of an earlier event from the date of a later event, researchers can quickly and easily create variables that measure the time elapsed between any two events.

b. CEO turnover. The BRD identifies the person who was CEO at filing and indicates that person's beginning and ending dates as CEO. From that data, researchers can construct several useful variables. Examples include (1) the CEO's length of service prior to filing, (2) how long the CEO survived in office after filing, and (3) the rate of CEO turnover during bankruptcy.

c. 363 Sales. The BRD uses the term "363 sale" to refer to a sale of all or substantially all of the debtor's assets in a sale that occurred prior to plan confirmation. The BRD contains (1) a field indicating whether a 363 sale occurred, (2) a field indicating the date on which the court approved the 363 sale, and (3) a field indicating whether the debtor sold all or substantially all of the debtor's assets in multiple 363 sales. .

d. Case cost. The BRD contains a field indicating the total professional fees incurred in the restructuring.

e. Case complexity or intensity. The BRD contains a variable, DktNumConf, which is the docket number of the plan confirmation order. Because docket numbers in each case begin with one and are generally assigned sequentially to documents as they are filed, the docket number of the confirmation order is a good approximation of the number of documents filed in the case by the time of plan confirmation. The number of documents filed in a case may be a useful indicator of the complexity of the case or the intensity with which the case was litigated.

f. Court. In large, public company bankruptcies, most debtors choose a court other than the one that serves their headquarters. The BRD contains variables that identify the debtor's home court, the court chosen for filing, whether the case was transferred, and the court in which the case was concluded. Additional fields indicate whether the case was filed in a court other than the home court, the distance from the debtor's headquarters to Delaware or the chosen court, and the dates of case transfers. Courts are identified by district, city, and judge.

g. Case type. Success rates may vary with whether the plan was prepackaged, prenegotiated or neither. They may also vary by whether the debtor or creditors initiated the filing. The BRD contains variables measuring both. The BRD also contains a field showing whether a voluntary case was filed under Chapter 7.

h. Trustees. The BRD contains variables indicating whether appointment of a Chapter 11 trustee was requested, whether the court ordered the appointment, and, if so, the name of the trustee and the debtor entities for which the trustee served.

i. Examiners. The BRD contains variables indicating whether appointment of an examiner was requested, whether the court appointed an examiner, and if so, the examiner's name.

j. Committees. The pattern of committee appointments may be related to success. The BRD contains fields indicating whether creditor, equity, and retiree committees were appointed.

k. Unions. The BRD contains fields indicating the numbers of employees and the number of union employees prior to filing and after emergence.

l. Attorneys. The BRD presently contains fields indicating the names of the lead law firms for (1) the debtors, (2) the creditors' committee, and (3) the equity committee. From these fields, researchers can make dummy variables to determine whether service by particular law firms or service by law firms in particular groups or categories are related to success. Relevant groups or categories might be law firm sizes, law firms in particular geographical locations, or law firms with particular levels of experience in large, pubic company cases. To facilitate the latter inquiry, we plan to construct a field that indicates the number of BRD cases in which the firm served as lead. That field will be a proxy for the level of experience of the attorneys as they served in each case.

m. Distributions. The BRD contains the amounts of money distributed to secured creditors, unsecured creditors and equity holders under plans (as specified in the disclosure statements accompanying the confirmed plans).

n. Substantive consolidation. In nearly half of large public company bankruptcy cases, the estates of corporate group members are substantively consolidated so that creditors of different entities receive the same or similar distributions. A BRD variable indicates whether substantive consolidation occurred.

o. Fee reviewers. A BRD variable indicates whether a fee auditor, a fee committee, a fee examiner, a fee review committee, or a joint fee review committee was appointed.

5. Case outcomes.

Plan confirmation is the traditional measure of case outcome. But the BRD contains eight groups of variables that plausibly measure case outcome: firm profitability, firm growth, firm survival, firm goal achievement, distribution size, leverage reduction, speed, and cost. These are the dependent variables for success modeling.

a. Firm profitability. The BRD contain four measures of the debtor's profitability, at each of two points in time. The measures are net income, EBITDA, EBIT, and income before extraordinary items (IncomeBE). A figure for each variable in each bankruptcy case is taken from the last 10-K for a period ending before the filing of the case and the first 10-K for a period ending after plan confirmation

b. Firm growth. Reorganization "success" can also be defined as an increase in company size, measured by revenues, assets, or employees. BRD users can calculate these increases from BRD data as proportions ("a 20% increase in revenues") or as normalized changes in dollar amount ("an increase in EBITDA equal to 5% of revenues").

c. Firm survival. The BRD contains several variables that measure firm survival. Emerge indicates whether the debtor's business continued to operate after plan confirmation. Refile

indicates whether the debtor refiled bankruptcy after plan confirmation. Refile 5 indicates whether the refiling occurred within five years of plan confirmation. Confirmation and refiling dates are available, so the user can easily create a variable that measures the time to refiling. We are in the process of collecting firm survival variables that indicate whether the company continues in business and, if not, the date and nature of its demise.

d. Goal achievement. Debtors may declare specific goals at the time the case is filed. For example, the BRD's SaleIntended variable indicates whether the debtor stated at the time of filing that it sought to sell the company. If the debtor sought to sell and sold, that can be regarded as success. On the same basis, plan confirmations can be regarded as successes, and dismissals or conversions can be regarded as failures. Debtors sometimes declare other goals at the time of filing, but we have not yet gathered that data.

e. Distribution size. The size of the distributions to creditors and shareholders can be regarded as a measure of success. The BRD contains variables that measure the amounts of secured claims and unsecured claims, and the distributions in dollar amount to secured creditors, unsecured creditors and equity. The dollar amounts make it possible to calculate the ratios of distributions to claims ("ten cents on the dollar"), the portions of the distributions going to particular parties, or the ratios of the distributions to other measures of firm size. An example of the latter would be the ratio of the entire distribution to the entire market or book value of the company at the time the case was filed. Lastly, if adherence to the absolute priority rule is considered a bankruptcy objective, success can be measured by how closely the distributions adhere to that rule.

f. Leverage. A commonly asserted objective of bankruptcy reorganization is to reduce debt to a manageable level. The BRD has variables for two measures of assets and liabilities at filing (10-K or Exhibit A) and one measure of assets and liabilities after emerging (10-K). From those variables, researchers can calculate the debtor's leverage at filing, the debtor's leverage after emerging, and the difference between the two. "Success" might be defined as the reduction in leverage or the reduction in excess leverage (with "excess" perhaps defined as excess over the industry norm).

g. Speed. Reorganization speed is not an end in itself. But some regard reorganization speed as success because they believe that speed serves other important objectives. The BRD provides several date fields from which case duration can be calculated in a variety of ways.

h. Cost. If reorganization outcomes are not sensitive to the identity the debtors' professionals or the actions they take, reorganization is merely a commodity. The best reorganization would be the cheapest. ProfFees10k is a variable that measures the debtor's total professional fees paid for the restructuring. Because we collect these data from the 10-Ks filed by the debtor during and after the bankruptcy case, this variable is available only for the minority of companies that survive and remain public.

6. General Economic Conditions.

The BRD contains three categories of variables that measure the condition of the economy around the time the case was pending: interest rates, gross domestic product, and the rate of inflation.

a. Interest rates. The BRD contains seven measures of the prime rate of interest around the time of the case. They are the prime rate of interest (1) two years before filing, (2) one year before filing, (3) at filing, (4) at case disposition (confirmation, dismissal or conversion), (5) one year after case disposition, (6) two years after case disposition, and (7) at refiling. From these seven variables, researchers can construct additional measures of the changing cost of capital to reorganizing firms. For example, the ratio of the prime rate at disposition to the prime rate at filing indicates the degree to which the case occurred in an environment of rising or falling interest rates.

b. Gross Domestic Product (GDP). The BRD contains seven measures of GDP around the time of the case. They are GDP for the quarter (1) two years before filing, (2) one year before filing, (3) at filing, (4) at case disposition (confirmation, dismissal or conversion), (5) one year after case disposition, (6) two years after case disposition, and (7) at refiling. From these seven variables, researchers can construct additional measures the health of the economy. For example, the ratio of GDP at disposition to GDP at filing indicates the degree to which the economy improved or deteriorated during the case. To the extent that the economy deteriorated, the economy may partly explain any failure of the reorganization process. Ratios among the prepetition GDP variables may be used to test for lagged effects of GDP changes on reorganization success.

c. Inflation. The BRD contains the consumer price index at filing and disposition. Those variables make it possible to adjust any of the amounts reported in the BRD to real dollars. They can also be used to explore whether inflation has any effect on the success of the reorganization process.

C. Representative Ranges

Bankruptcy case dockets are sporadically available for BRD cases from 1979 through October, 1995 and uniformly available since that date. Bankruptcy case files were sporadically available for BRD cases from 1997 through the mid-2000s. At present, case files are available for all cases filed after November 30, 2003 (the files are complete for all but three cases). Case files are available for cases filed prior to December 1, 2003 only if the cases remain open. If the data in a field are from bankruptcy dockets or case files, the field may be incomplete for the earlier periods, while complete for the later periods. If so, the data we collected for the earlier periods may not be representative of all data for those periods. To assist researchers in distinguishing representative from non-representative data, our policy is to add a field that indicates which data points for a particular field group are representative. At present, we have only two representative range fields, for the turnaround manager field group and the five year refiling groups. The Do File creates other representative range fields in Stata.

IV. Specific Research Types

This part provides additional information and advice regarding specific types of research.

A. Venue

Professor Lynn M. LoPucki initially designed the BRD for venue research. Although the BRD contains a field showing the district in which the case was filed, we recommend use of the city in which the case was filed instead. Forum shopping is best regarded as a choice of city rather than district because forum shoppers are seeking particular judge panels and the judge panels exist at the city, rather than the district level. For example, Boston and Worcester are both courts in the District of Massachusetts. But forum shopping is to Worcester (three of the four cases filed in Worcester) at the same time it is from Boston (24 of 29 companies headquartered in Boston). If the researcher measures forum shopping at the district level, the researcher would understand the forum shopping from Boston and miss the pattern of forum shopping to Worcester altogether.

Depending on the purpose of the research, researchers may wish to separate "voluntary" cases filed by debtors from "involuntary" cases filed by creditors. Forum shopping by debtors and creditors are to different places for different purpose and have very different effects.

Lastly, researcher will probably wish to separate Chapter 7 cases from Chapter 11 cases. By filing under Chapter 7, the debtor voluntarily surrenders control of the company and the case to the bankruptcy court.

B. Success Modeling

Most of the fields added to the BRD in recent years have been added to facilitate success modeling. Success modeling is the estimation of regression models that predict which cases will "succeed" by a variety of criteria.

Success models are most relevant and practical for the set of cases filed after November 30, 2003 (the "Privacy Cutoff"). Court files are no longer available for the set of cases filed prior to December 1, 2003. As a result, many fields have substantial missing data for the pre-Privacy Cutoff set of cases.

If a case lacks data for any of the variables in your regression model, the entire case is omitted from the calculations. If too many cases are omitted, those remaining may no longer be representative and regression may produce a faulty model. The solution is to use pre-Privacy Cutoff data only for evaluating those variables that are generally available in the pre-Privacy period. Run the models you intend to publish using only post-Privacy Cutoff data.

C. Professional Fees

The BRD contains professional fees data from Annual Reports (10-Ks) filed by the bankrupt companies with the Securities and Exchange Commission. Because only surviving companies that remain public file Annual Reports, those data are sparse and may not be representative.

The BRD website contains additional professional fees data drawn from the court files of 102 of the largest cases concluded from 1998 through 2007. Ninety-nine of the cases are BRD cases. Those data may provide a useful addition to the BRD cases for particular projects. For a complete analysis of this additional data see Lynn M. LoPucki and Joseph W. Doherty, *Professional Fees in Corporate Bankruptcies* (Oxford 2011).

V. Data Contributions to the BRD

The BRD seeks contributions of data from BRD users. One such contribution is to advise us of specific BRD data that are incorrect. (It is helpful, but not necessary that you furnish us with evidence.)

Another such contribution is to furnish us with a field or group of fields that are of sufficiently general interest that inclusion in the BRD is appropriate. If you plan to do that, please let us know in advance, so we can discuss the field protocols with you.

If we are able to agree on protocols, we will replicate your data collection, compare our data to yours, resolve the differences, and advise you of the results. Our experience with this process is that we generally collect the data more accurately than the data contributor. But the process is beneficial to us, because the data contributor is sometimes right when we are wrong.

In collecting the data, we ask that you observe these general protocols:

1. Collect the data on, or transfer the data to, an Excel spreadsheet that contains the debtor companies' BRD Primary Key numbers and names.

2. Retain copies of, or references to, the documents from which the data were drawn, and furnish them to us along with the data.