



STUDY OF THE TSUNAMI AFTERMATH AND RECOVERY (STAR): Study Design and Results

Elizabeth Frankenberg

Peter Katz

Iip U'mar Rifai

Bondan Sikoki

Cecep Sumantri

Ni Wayan Suriastini

Duncan Thomas

ACKNOWLEDGEMENTS

STAR is a collaborative project involving investigators at Duke University; SurveyMETER (Indonesia); the University of North Carolina, Chapel Hill; the University of California, Los Angeles (UCLA); the University of Pennsylvania; the University of Southern California; the World Bank; and Statistics Indonesia.

We are deeply grateful to Khusnul, Statistics Indonesia Aceh, and others at Statistics Indonesia who facilitated this survey.

In addition to the authors of this report, other key personnel at SurveyMETER include Muhammad Fajar Suminto, Amalia Rifana Widiastuti, Nasirudin, Okie Judhijanto, Muhammad Mulia, Saddam Kadafi, Alman Suriadi, and Muliani. Other key personnel at Duke include Michael Burrows, Nick Ingwersen, and Sheila Evans.

We appreciatively acknowledge financial support from the National Institute for Child Health and Human Development (HD051970, HD052762) the National Institute on Aging (AG031266), the MacArthur Foundation (05-85158-000), the National Science Foundation (CMS-0527763), the Hewlett Foundation and the World Bank.

Table of Contents

- 1. Introduction 1
- 2. Survey Instruments 2
- 3. STAR File Structure and Naming Conventions..... 4
 - Identifiers and Level of Observation 4
 - Combining Data Across Files 5
 - Concatenating Data 5
 - One-to-One Merges at the Individual or Household Level..... 6
 - One-to-Many Merges 6
 - Questions Numbers and Variable Names 7
 - Response Types 7
 - Missing Values 8
 - Special Codes and X Variables 8
 - TYPE Variables 9
 - Privacy-Protected Information 9
- 4. Special Features of the STAR Data 10
 - Symmetric Information 10
 - Duplicate Information 10
 - Family Relationships 11
 - Parents, Children, and Spouses Identified in the AR Roster 11
 - Classifying Relatives 12
 - Data Availability for Households and Individuals: HTRACK and PTRACK..... 13
- Appendix: Description of the STAR Questionnaire..... 16
 - Book K: Control Book and Household Roster 16
 - Book 1: Household Expenditures and Characteristics 16
 - Book 2: Household Economy..... 16
 - Book 3: Adult Individual Book 17
 - Book 4: Information for Ever-Married Women 18
 - Book 5: Child Information 18

I. Introduction

The Sumatran-Andaman earthquake of December 26, 2004 and the subsequent tsunami constitute one of the most powerful and deadly disasters in world history. The damage in the Indonesian province of Aceh was extreme. Some 160,000 individuals perished and an estimated 4.5 billion dollars of property was destroyed. By 2007, efforts to repair the vast damage and rebuild destroyed infrastructure constituted one of the largest projects ever undertaken in a developing country setting.

To provide evidence on the consequences of the disaster and the evolution of the recovery effort, our team fielded multiple surveys in Aceh and North Sumatra as part of the Study of the Tsunami Aftermath and Recovery (STAR). Fieldwork for the first of these surveys, STAR1, began in May, 2005, only five months after the tsunami occurred.

STAR is designed to provide data for studying many behaviors and outcomes associated with the tsunami and the subsequent recovery. The survey collects information at the individual and household levels, including multiple indicators of economic well-being (consumption, income, and assets); education, migration, and labor market outcomes; marriage, fertility, and contraceptive use; health status and use of health care; relationships among coresident and non-coresident family members; transfers among family members and inter-generational mobility; access to public services, and participation in community activities.

With these data, the analyst can address questions regarding the impact of the tsunami and post-disaster recovery on the lives of the respondents and document the effects of social, economic, and environmental change on the population.

STAR surveyed respondents living in Aceh and North Sumatra, across 13 *kabupaten*¹ with coastlines potentially vulnerable to inundation from the 2004 tsunami. The target population was located in Aceh and in the neighboring province of North Sumatra. STAR is designed to be representative at the *kabupaten* level. The original frame for the STAR survey was the 2004 SUSENAS, collected by the Indonesian Central Bureau of Statistics. *Kabupaten* were chosen to cover both the west and east coasts of Aceh as well as vulnerable areas of North Sumatra. Clusters for the survey include areas with light, medium, and heavy damage. No geographic identifiers are provided in the STAR data apart from province.


The first STAR survey (STAR1) targeted 7,157 households containing 28,376 individuals for interview. We fielded an additional four surveys over the subsequent four years and a fifth in 2014-15.

Across all survey rounds we succeeded in obtaining at least one interview for over 95% of target respondents. Because we follow individuals who were members of an original household in the first wave (panel respondents) but who have moved out of their original household, we have added "split-

¹ Indonesia is a presidential democracy with four levels of government. The provincial level is just below the national level, next are *kabupaten* which are equivalent to districts, followed by *kecamatan*, or subdistricts. The smallest unit is the *desa* or *kelurahan* which are villages and townships/wards, respectively.

off" households over time. In addition, we interview individuals who are new members of the households and we track the biological children of panel respondents as they move out of their parents' households. Outcomes for STAR1 target households are depicted in Figure 1.

FIGURE 1



It is important to note that although STAR shares a number of common features with the Indonesian Family Life Survey (IFLS), particularly with respect to the design of the questionnaire, the two projects are different on some key dimensions. First, with respect to the populations the surveys represent, IFLS was designed to represent the Indonesian population in 1993, whereas STAR is representative of *kabupatens* in 2004. Second, the IFLS is a multi-purpose survey designed to provide data on a wide array of topics, whereas STAR is designed to measure the impact of the 2004 earthquake and tsunami and on health, economic, and social well-being.

2. Survey Instruments

STAR is a comprehensive multipurpose survey that collects data at both the household and individual levels. In STAR, one or two household members provide information at the household level. In addition, we attempted to conduct an individual interview with every household member aged 11 and older, and with a parent or caretaker on behalf of children 10 and younger. Because obtaining interviews with all household members is difficult, STAR includes a proxy book to collect more limited information from a proxy respondent about individuals who were not available for an in-person interview. Table 2.1 outlines the questionnaire structure and content, which are described in more detail below and in Appendix A.

The questionnaire was divided into *books* (usually addressed to different respondents) and subdivided into topical *modules*. Three books collected information at the household level, generally from the household head or spouse²: book K, book 1, and book 2. Additional books collected individual-level data from adult respondents (book 3 and, after STAR1, book 4), and children younger than 15 (book 5). Some

² In STAR, one member of the household was designated the household head by the person who provided information on the composition of the household. Where a married couple headed the household, the husband was generally designated the head and the wife was generally designated the spouse of the head. The head of the household is permitted to change from one STAR survey to the next, according to respondent reports.

modules appear in both a household book and an individual book (for example HR, on assets), because we wanted to collect the data both for the household as a whole and from individuals.

Book K: Control Book and Household Roster. Book K recorded whether a household was found and interviewed and the location of the household. If at least one member of the target household was found, information on the activities of all current members of the target household was collected.

Book 1: Household Expenditures and Housing Attributes. This book was typically addressed to a female respondent, either the household head, the spouse of the household head, or another person knowledgeable about household affairs. The first module recorded information about household expenditures. The second module collected information about quantities and purchase prices of several staples. The third module collected information about basic housing characteristics. The fourth module recorded information about financial and non-financial transfers into and out of the household.

Book 2: Household Economy. This book was usually answered by the household head or the head's spouse. Modules asked about household businesses (farm and non-farm), business assets, household assets, borrowing, and income. Combined with individual-level data on labor and non-labor income collected in book 3, this information can be used to provide a picture of current household income from market-wage income, self-employment income, family businesses, informal-sector activities, and non-labor income.

Book 3: Individual Adult Information. This book asked all household members 15 years and older about their education and employment status, asset ownership, marriage and fertility, and migration histories. In addition, the book included questions on morbidities, health status, health diagnoses, psychological health, usual practices (such as smoking and fruit/vegetable consumption), outpatient utilization, attitudes toward time and risk, social networks, and community participation. Finally, this book included modules on individual economic shocks and exposure to the earthquake and tsunami.

Book 3P: Individual Adult Information by Proxy. The proxy book was designed to facilitate collecting data by proxy about individuals who could not be interviewed directly. The proxy book contains shortened versions of most of the sections included in Book 3, excluding sections that cannot be accurately answered by a proxy respondent. In STAR1 we have combined the data from Book 3 and Book 3P to make it easier to use the data together (b3_cov contains the variable "proxy" which indicates whether a respondent answered directly, or a proxy provided answers to the questions).

Book 3S (after STAR1): In survey rounds after STAR1 an additional book addressed to individuals 15 and older asked more detailed questions regarding interactions with non-coresident kin (parents, siblings, and children), as well as standard batteries of questions related to post-traumatic stress and depression.

Book 3B (after STAR2): In survey rounds after STAR2 a book for first-time adult respondents was created so that we could collect retrospective information for these respondents, without repeating life histories for panel respondents.

Book 4 (after STAR1): In survey rounds after STAR1 Book 4 was administered to women 15-49 to collect information on fertility preferences and, for ever-married women, full pregnancy histories and current use of contraception.

Book 5: Child Information. This book collected information about children younger than 15. For children younger than 11, the child's mother, guardian, or caretaker answered the questions. Children between


the ages of 11 and 14 were allowed to respond for themselves if they felt comfortable doing so. The five modules focused on the child’s education status, work and time use, morbidities, self-treatment, and inpatient and outpatient visits. Each paralleled a module in the adult questionnaire (book 3), with some age-appropriate modifications. For example, the list of acute health conditions specified conditions relevant to younger children.

3. STAR File Structure and Naming Conventions

This section describes the organization, naming conventions, and other distinctive features of the STAR data files to facilitate their use in analysis. Additional information about the data files is provided in the survey questionnaires and codebooks. For analysts’ convenience, each page of the questionnaire includes the names of files that contain information from that page on the bottom left-hand corner. The codebook for each questionnaire book describes the files containing the data for that book and the levels of observation represented.

STAR data files correspond to questionnaire books and modules. There are multiple data files for a single questionnaire module if the module collected data at multiple levels of observation. For example, module KP (psychosocial well-being) collected information at the individual level (on use of counselling services) and at the symptom level, so two data files are associated with that module.


File naming conventions are straightforward. The first two or three characters identify the associated questionnaire book, followed by characters identifying the specific module and a number denoting sequence if data from the module are spread across multiple data files:



Identifiers and Level of Observation

Wherever possible the data have been organized so that the level of observation within a file is either the household or the individual. Here we describe the conventions for STAR1, but these hold for later waves as well. If the level of observation is the household, variable IDHH01 uniquely identifies an observation. If the level of observation is the individual, both IDHH01 and IDP01 are required to uniquely identify a person.³

In STAR1, IDHH01 is a nine digit character variable whose digits carry the meanings described below.



³ Within STAR1 files, use IDHH01 and IDP01 to identify individuals. Use IDPLINK to identify individuals across survey rounds. Note that in the STAR AR roster in waves after STAR1, the variable IDPLINK does not uniquely identify individuals because individuals can be listed in more than one household roster (although they are a current member of only one household).

The first 3 digits correspond to a survey-defined cluster.⁴ The next five digits identify a specific household. The first three digits of the household identifier are consistent across all origin households; that is, all households arising from a STAR1 target household will begin with the same three digits. The fourth through sixth digits are the household's sequence number within a cluster. The seventh and eighth digits denote the survey round at which a household is first identified – for example, “01” for a household targeted in STAR1, or “02” for a household identified as a STAR2 split-off. The final digit designates the order in which a new split-off household is found in a particular wave. For example, a household that was intact in STAR1 may have generated two new households by STAR2 if two children moved out and each form his or her own household. These new households would have a 9th digit of 1 or 2.

The person identifier IDP01 is simply the line number of the person in the AR roster, which is two digits. In STAR, household identifiers are stored as alphanumerics, while person identifiers are stored as numerics. The alphanumeric IDPLINK identifies individuals across households. IDPLINK consists of IDHH and IDP, based on the household and roster number in which the individual was first interviewed and is constant across waves and across whatever households an individual is interviewed in. The conventions described above for STAR1 also characterize later waves of STAR.

In some data files the level of observation is something other than the household or individual. Usually this occurs because the data were collected as part of a grid, in which a set of questions was repeated for a series of items or events. For example, in the business asset data from module HU, each observation corresponds to a particular type of asset, and there are multiple observations per household. In this data file, the combination of IDHH01 and HUTYPE uniquely identifies an observation.

The variable that defines the items or events is usually named XXTYPE, where XX identifies the associated module (more is said about TYPE variables below).

In some cases, data collected as part of a grid are organized rectangularly. For example, file B2_HU contains data about 9 asset types for each household. In other cases, the number of records per household or individual varies. For example, the level of observation in file B3_MG2 is an individual migration by a respondent. Not all respondents migrated the same number of times, so some individuals appear only once, others appear twice, and some appear more than twice. Those who did not migrate do not appear at all. This file is not rectangular because the number of observations per person is not constant. To uniquely identify an observation in this file, the analyst should use IDHH01, IDP01, and MOVENUM.

Combining Data Across Files

As explained above, STAR data are stored in many different data files. To create analytic files, the analyst usually needs to combine the data from different files. How the data should be combined depends on the nature of the desired analytic file. Below we briefly describe ways to link data across files using STATA.

Concatenating Data

The analyst may wish to pool observations by concatenating two data files. For example, B3_RJ and B5_RJA both contain data on visits to outpatient providers. The data in B3_RJ pertain to adults, and the

⁴ Clusters are groups of households from ecologically similar areas with respect to degree of tsunami damage, distance to the coast, elevation, and rural-urban status.

data in B5_RJA pertain to children. The variables for adults begin with RJ, while the variables for children begin with RJA, but otherwise the information is the same. In some contexts it may be useful to combine the data for the two age groups, rather than keeping it in two separate files. The data can be combined into one file using the APPEND statement in STATA. The resulting file will contain both the observations for children and the observations for adults. Because the variable names are different, the variables in one file should be renamed so that they match the names in the other file before appending.

One-to-One Merges at the Individual or Household Level

In many cases the analyst will want to link data from one file with data about the same respondents from another file. If both files contain data at the same level of observation, the linkage will be a “one-to-one” merge. Some examples of variables to use when merging are below.

Merging Two Files at the Individual Level of Observation. Suppose the goal is to create a file that contains information on an individual’s literacy and his or her primary activity in the past week. The file B3_DL contains information on whether respondents can read or write. The file B3_TK contains information on the respondents’ primary activity in the past week. Both files contain one observation per individual. To link the desired information, sort each of the two files by IDHH01 and IDP01 and then merge on IDHH01 and IDP01.

Merging Two Files at the Household Level of Observation. For two household-level files, such as B1_KR and B2_UT0, which contain data on housing characteristics and borrowing, sort each file by IDHH01 and merge on IDHH01.

One-to-Many Merges

Often the analyst will want to merge files that are not organized at the same level of observation. Sometimes such a merge is straightforward. Other times it will require restructuring at least one of the data sets. When thinking about how to merge STAR data files, it is helpful to determine whether the identifying variables in one of the files are a subset of the identifying variables in the other file.

This is easiest to explain using an example. Suppose the analyst wishes to merge information on literacy with information on asset ownership. The identifying variables in B3_DL are IDHH01 and IDP01. The identifying variables in B3_HR1 are IDHH01, IDP01, and HR1TYPE. In B3_HR1, an individual has 7 records, one for each asset type about which we inquire. The data can be merged in two ways.

First, because the identifying variables in B3_DL are a subset of the identifying variables in B3_HR1, you could simply merge on IDHH01 and IDP01. This merge yields 7 records for each individual. Each record contains information about the individual’s literacy and information about a particular asset type.

The other option is to restructure B3_HR1 so that it is organized at the level of the individual rather than at the level of the asset, in which case the identifying variables will be IDHH01 and IDP01 in the reorganized HR data. This would involve creating a file that contained variables HR01–HR07 for asset type A (i.e., HR01A–HR07A), as well as variables HR01–HR12 for the other asset types (HR01B–HR12B, HR01C–HR12C, etc.). This file would have many more variables than B3_HR1 but many fewer observations. If the data from the B3_HR1 file are restructured so that they are at the level of the individual, merging the restructured file by IDHH01 and IDP01 with the B3_DL1 data yields one record per person that contains literacy information and all information on the different types of assets.

Restructuring data files so that they are organized at a different level of observation can be done relatively easily in STATA with the reshape commands. Some data files cannot be merged without restructuring one of the data files. For example, the identifying variables in B2_UT1 are IDHH01 and UTTYPE. The identifying variables in B2_NT1 are IDHH01 and NT1TYPE. Neither file's identifying variables is a subset of the other's identifying variables. To merge data from these two files, you must first restructure one or both of them so that IDHH01 is the identifying variable. Generally, it is not wise to merge two files that both contain data from grids and have a TYPE identifying variable without restructuring the data.

Questions Numbers and Variable Names

Most STAR variable names closely correspond to survey question numbers denoted in the questionnaire. For example, the names of variables from the DL module (education history) begin with DL and end with the specific question number.

A number of questions have two associated variables: an X variable indicating whether the respondent could answer the question and the "main" variable providing the respondent's answer. X variables are named by adding "x" to the associated question number. For example, question TK15 asked about respondent earnings in a calendar month. Variable TK15x indicates whether the respondent was able to answer the question. Variable TK15 provides the amount of earnings the respondent reported. In the questionnaire, the existence of an X variable is signaled when the interviewer is asked to circle a number indicating whether the respondent was able to answer the question (in the case of TK15x, 1 if a valid amount is provided, 3 if the amount is nothing, 7 if the respondent refused, 8 if the respondent doesn't know the amount, and 9 if the response is missing). In the codebooks, the name of the variable itself signals its X status. The label for an X variable includes an "able ans" at the end. X variables are further discussed below.

Response Types

The vast majority of STAR questions required either a number or a closed-ended categorical response; a few questions allowed an open-ended response.

The numeric questions generally specified the maximum number of digits and decimal places allowed in an answer; any response not fitting the specification was assigned a *special code* by the interviewer, and the special codes were reviewed and recoded later (explained further below). Where it was necessary to add digits or decimal places as a result of that review, we may not have updated the questionnaire. The codebook provides information on the length of each variable.

Questions requiring categorical responses usually allowed only one answer (for example, "Was the school you attended government, private, or pesantren?"). When only one answer was allowed, numeric response codes were specified. If more than four numeric response codes were possible, two digits were used so that 95–99 could serve as special codes. Some questions allowed multiple answers (for example, "Why did you leave school?"). In that case, alphabetic response codes were specified. When multiple responses were allowed, the number of possible responses set the maximum possible length for the variable.

For categorical variables, the questionnaire provides the full meanings for each response category. The codebook contains a short "format" that summarizes the response category, but analysts should check the questionnaire for the clearest explanation of response categories and not rely solely on the codebook format.

The codebook also provides information on the distribution of responses. For numeric variables, the mean, maximum, and minimum values are given. For categorical variables the frequency distribution is provided. For categorical variables where multiple responses were allowed, the codebook provides the number of respondents who gave each response. Since many combinations of responses were possible, the codebook does not provide the distribution of all responses. For example, question DL07 asked why the respondent left/missed school and allowed up to 15 reasons in response. The codebook shows how many respondents cited illness and how many respondents cited marriage but not how many respondents cited both illness and marriage.

Additional response categories were sometimes added in the process of cleaning “other” variables. Typically these categories were added below the existing “other” category. For example, question TK29 asked about why a respondent was unemployed. The questionnaire was fielded providing one substantive choice (“tsunami”) and a second “other”. When the “other” responses were reviewed, seven other additional and commonly identified categories were added.

Missing Values

Missing values are usually indicated by special codes. For numeric variables, a 9 or a period signifies missing data. For character variables, a “Z” or a blank signifies missing data.

For many variables, we can distinguish between *system missing data* (data properly absent because of skip patterns in the questionnaire) and data missing because of interviewer error. The data entry software generated some missing values automatically as a result of skip patterns. For example, question DL00 in book 3 asked whether the respondent had ever attended school, and if not, interviewers skipped to question DL12. If the interviewer recorded 3 (No) for question DL00, during data entry the software automatically skipped to DL12 and filled the book 3 DL02-DL11 variables with a period or blank. If data were missing because the interviewer neglected to ask the question or fill in the response, the data-entry editor was forced to enter 9 or Z in the data fields in order to get to the questions that the interviewer did ask.

Sometimes valid answers are missing not because of skip patterns or interviewer error but because the answer did not fit in the space provided, the question was not applicable to the respondent, the respondent refused to answer the question, or the respondent did not know the answer. In these cases special codes ending in 5, 6, 7, or 8 were used rather than 9 or Z (see below).

Special Codes and X Variables

Many STAR questions called for numeric answers. Sometimes a respondent did not know the answer or refused to answer. Sometimes the respondent said that the question was not applicable. Sometimes the answer would not fit the space provided, either because there were too many digits or decimal places were needed. Sometimes the answer was missing for an unknown reason. In all of these cases, interviewers used special codes to indicate that the question had not been answered properly. The last digit of a special code was a number between 5 and 9, indicating the reason:

- 5 = out of range, answer does not fit available space; labeled “Top Code” when this occurs in x variables
- 6 = question is not applicable
- 7 = respondent refused to answer

8 = respondent did not know the answer

9 = answer is missing

The other spaces for the answer were filled with 9's so that the special code occupied the maximum number of digits allowed.

Rather than leave special codes in the data, for most variables we created indicator (X) variables showing whether or not valid numeric data were provided. An indicator variable has the same name as the variable containing the numeric data except that it ends in X. For example, the indicator variable for KH06 (how much the respondent paid for an amount of food) is KH06X. The value of KH06X is 1 if the respondent provided a valid numeric answer and 8 if the respondent did not know what price he or she paid.

An indicator variable sometimes reveals more than whether special codes were used. For example, for KR12 (duration lived in house of interview), KR12X indicates both the units in which duration was recorded (weeks, months, or years) and the existence of valid numeric data.

TYPE Variables

As noted above, in some modules the data are arranged in grids, and the level of observation is something other than the household or individual. Examples are KS (household expenditure) data on prices, where the level of observation is a food or non-food item; HR (household assets) data, where the level of observation is a type of asset; and KP (psychological health) data, where the level of observation is a symptom. The name of the variable that identifies the particular observation level typically contains the module plus "TYPE," e.g., HRTYPE. In modules with TYPE variables, there are multiple records per household or individual, but combining HHID or HHID and PID with the TYPE variables uniquely identifies an observation. TYPE data can be either numeric or character.

Weights

HTRACK contains weight variables (HWT and HWTP). HWT is raked to Statistics Indonesia data from 2005 for the STAR districts, based on district, urban-rural status, household size, and age of the household head. HWTP scales households so that they sum to the number of households with the STAR districts in 2005.

PTRACK contains weight variables (PWT and PWTP). PWT is raked to Statistics Indonesia data from 2005 for the STAR districts, based on age, sex, district, and urban-rural status. PWTP scales individuals so that the sample as a whole equals the estimated number of individuals within the STAR districts in 2005.

Privacy-Protected Information

In compliance with regulations governing the appropriate treatment of human subjects, information that could be used to identify respondents in the STAR survey has been suppressed. The types of information include location, lengthy text responses, and very rare response codes. Interview dates, which are provided in HTRACK, have been randomized across a period within about two months of the actual interview date.

4. Special Features of the STAR Data

This section discusses the distinctive features of STAR1 data as they affect analysis files.

Symmetric Information

In some STAR modules, respondents within a household provided symmetric information. These data may permit comparisons that help validate or fill in incomplete sections in the data. In module KF, individuals provided information about the dates of their marriages. Within a household, if two individuals are married to each other, the KF data from the two individuals could be compared for consistency. Or, if one individual's data is missing, data from the spouse could be used to fill the gap. Similarly, in module MG, individuals described their migration experiences. If couples or parent-child pairs had moved together and each individual answered MG, their responses could be compared to check consistency, or the MG data of one could supply information missing from the other.

Duplicate Information

Certain pieces of information were collected in more than one place. In most cases, the respondent was one source of information and a proxy respondent (or preprinted information) was the other. For example, the household roster (module AR) contained information on a number of topics that were also in the questionnaire books addressed to individuals. Though it would be easier to use the information from the roster, data from the individual books are likely to be more accurate, since the information was self-reported rather than provided by proxy.

Age. Information on age was collected in both the AR roster (generally by proxy) and on the covers of the individual books. In addition, in certain places in the questionnaire, interviewers were required to examine the age recorded on the book cover, usually to determine whether the respondent was above a certain threshold age. We did not correct inconsistencies between the roster and book covers, but the PTRACK file (describe later) contains a "best-guess age" variable. We did not attempt to correct inconsistencies between the roster and questions within the book, since complicated skip patterns were often involved.

Birthdate. Information on birthdate was collected in individual books and recorded in the AR roster. We did not correct inconsistencies between the AR roster and the health assessment, but the PTRACK file contains a best-guess birthdate variable. If the respondent knew the year of birth but not the month or day, this variable shows month and day as 98.

Sex. After the initial round, a respondent's sex was preprinted in the AR roster and collected on the cover of books. In cases of inconsistency between the roster and book covers, we undertook extensive checks on name, STAR information, and other data to ascertain a best guess for sex. The best-guess sex values are recorded in PTRACK

Marital Status. Marital status was noted in the AR roster and on the covers of book 3. Various interviewer checks within the individual books required using marital status information from the book cover. In cleaning the data, we tried to make sure that marital status in the roster matched marital status on the book covers. We did not clean interviewer checks because that would have required complicated adjustments to skip patterns.

Education Level. The AR roster reported the highest level of schooling attained and the highest class completed within that level (AR30 and AR31), based on the answer provided by the respondent to Book K. Many respondents provide the information for themselves in response to questions in Books 3 or 5.

Generally we recommend using the data provided by the individual rather than by a proxy if there is a choice.

Earnings and Non-labor Income. Module TK asked in depth about employment and labor earnings. The proxy book also addressed these topics. As insurance in case neither module was completed for some household members, we also included a question on earnings in the AR roster (AR37 and AR38). The existence of the AR data means that a measure of total household labor income can be computed, even if not all household members provided a book 3 or proxy book. However, data from TK are preferred because they come from the respondent or a knowledgeable proxy. TK data are likely to be more accurate also because earnings were addressed in the context of related questions.

Parents' Survival Status. The AR roster recorded PID numbers for each individual's mother and father (AR20 and AR24). If the mother or father was not a member of the household, codes were used to designate whether the parent was alive and living in another household or dead, and if dead, whether that death was attributable to the tsunami. Book 5, module BAA (parental information) explicitly asked the respondent about each parent's survival status. The BAA data are preferred.

Timing of Marriage and Pregnancy. The KF module provides data on marital timing and limited information on a woman's pregnancies. KF also provides some information on contraception. In rounds after STAR1, full pregnancy histories were administered.

Family Relationships

STAR contains extensive information on family relationships, particularly between husbands and wives and between parents and children. The information is not limited to household members but also covers non-coresident kin.

Parents, Children, and Spouses Identified in the AR Roster

The AR roster provides much information on relationships among current household members, as shown in the table below:

Table 1

Variable	Information	Remarks
AR07	Which member was designated household head and how other household members were related to that person	Sometimes this information indicates how members other than the household head were related. For example, if persons 3 and 4 were both children of the head, they were either full or half-siblings. If person 4 was the mother of the head and person 3 was the child of the head, person 4 was almost certainly the grandmother of person 3. In other cases the information is not definitive. For example, if persons 5 and 6 were both grandchildren of the head, they were likely to be siblings or cousins, but we do not know which from AR07.
AR12, AR20, and AR24	PID numbers of an individual's spouse, birth mother, and birth father.	To find the education level of a child's parents, use the line numbers in AR20 and AR24 to link a child to its parents and thus to parents' education data either in the AR roster or in their individual book 3. If a person's mother or father, or spouse is not listed on the household roster (because they were not a household member) there is no PID.

Note one caution in using the AR data on family relationships. The accuracy of codes for AR21 and AR25 is not clear. The person completing the roster may not have known whether the father or mother was living or dead, especially in the displacement and confusion following the earthquake and tsunami in 2004. For the survival status of parents of children under 15 years of age in STAR, book 5 module BAA (parental information), is the preferred source of information because an explicit question was posed directly to the respondent. After STAR1, a BA module was also included in Book 3, for respondents 15 and older.

Classifying Relatives

Some relationships were not always specified with precision. In particular, the distinction between biological and through-marriage relationships was sometimes blurred. It was not always clear whether a child/parent was a biological child/parent, a step-child/-parent, or a child-/parent-in-law. Nor was it always clear whether someone classified as an aunt or cousin was related to the respondent or the respondent's spouse. We did not attempt to resolve all such inconsistencies. They were likely to arise in the contexts described below.

AR07 vs. AR20/24. Occasionally AR07 classified someone as the child of the head, but AR20 or AR24 did not list the head as the person's biological parent. The reason may be that AR20/24 asked specifically about the biological parent, whereas AR07 asked more generally about the relationship to the head. Likewise, AR07 sometimes listed an individual as the

parent of the household head, but that person's PID did not appear in the head's response for AR20/24 as a biological parent of the head.

Divorce. At the time of STAR1 some respondents' marriages had ended in divorce. During the "other" cleaning process we found responses indicating that someone was an ex-spouse or related to an ex-spouse. We created two new categories (ex-spouse and relative of ex-spouse) to account for this.

Asset Ownership. Modules HR, UT, and NT contained questions asking whether other family or household members were co-owners of various assets. In some cases it is not clear whether someone categorized as an aunt is related to the respondent or the respondent's spouse.

Data Availability for Households and Individuals: HTRACK and PTRACK

Files named HTRACK and PTRACK indicate what data are available for households and respondents, respectively, in each survey wave.

HTRACK

HTRACK also contains a record for all respondent households that were interviewed in STAR1. HTRACK provides information on whether the household was interviewed and, if so, whether data from books K, 1, and 2 are available.

HTRACK also contains household survey weights for all respondent households and an indicator for whether the cluster is rural.

Finally, HTRACK contains a measure of damage from the tsunami (the variable name is "damage"). The variable is defined at the level of the cluster. The measure takes on 3 values (no damage, some damage, and heavy damage), and is based on a combination of information collected from satellite imagery, survey supervisors, and community leaders.

PTRACK

PTRACK contains a record for every individual who was interviewed in STAR1, with information on interview status, sex, age in each wave, and variables that indicate our best effort to establish year, month, and date of birth.

PTRACK also provides individual survey weights for all STAR respondents.

Table 2

STAR1 Household Survey Questionnaires			
<i>Respondent</i>	<i>Module</i>		<i>Remarks</i>
Book K: Control Book			
Interviewer and household head, spouse, or knowledgeable other person	AR	Household roster	Preprinted with the names of all targeted household members.
	CP	<i>See Note at end of table.</i>	
Book 1: Expenditures, Prices, Housing Characteristics, Transfers			
Wife of household head, household head, or other knowledgeable person	KS	Consumption	
	KH	Prices	
	KR	Housing characteristics	
	TR	Transfers	
	CP	<i>See Note at end of table.</i>	
Book 2: Household Economy			
Household head, wife of household head, or other household member	UT	Farm business	
	NT	Non-farm business	
	HU	Business assets	
	HR	Household assets	
	HI	Income	
	BT	Borrowing	
	HE	Health expenses	
	CP	<i>See Note at end of table.</i>	
Book 3: Adult Individual Book			
Each household member age 15 and older	DL	Education	
	TK	Employment	
	HR	Individual assets	
	MA	Morbidities	
	KK	Health status	
	PB	Usual practices	
	RJ	Outpatient utilization	
	DK	Health diagnoses	

KF	Marriage and fertility
KP	Psychological health
SI	Attitudes
GEI	Individual economic shocks
MG	Migration
SN	Social and kinship networks
EX	Exposure
PM	Community participation
CP	<i>See Note at end of table.</i>

Book 5: Child Individual Book

Each child, age 0–14 (usually answered by the mother if the child was less than 11 years old)	DLA	Education
	TKA	Employment
	TAA	Time allocation
	MAA	Morbidities
	KKA	Health status
	RJA	Outpatient utilization
	EXA	Exposure
	KPA	Psychological health
	MGA	Migration history
	BAA	Parental information
	SNA	Social and kinship networks
	CP	<i>See Note at end of table.</i>

Book 3 Proxy

Someone who answered for the intended respondent to book 3.	Shortened version of other modules:	
	Book 3: DL, TK, MA, KK, PB, DK, MG, SN	
	CP	<i>See Note at end of table.</i>

Note: The CP module at the end of nearly every book asked the interviewer to record the conditions of the interview (who else was present, whether others provided assistance in answering questions), the respondent’s level of attention, and any other relevant information about the interview environment. The interviewer could also add information to explain or clarify the respondent’s answers.

Appendix:

Description of the STAR Questionnaire

This appendix expands on the summary presented in Section 2 for those interested in more detail about the survey instrument.

Book K: Control Book and Household Roster

The interviewer completed this book, or a portion of it, for all households targeted in STAR1. The cover of the book indicates (for each household, identified by HHID) whether the household was interviewed in the survey and if not, why not.

The AR Module (the household roster) lists all household members. After STAR1 the roster was preprinted with the name and characteristics of each household member from the previous wave. The roster was used to indicate household composition at the time of interview, and to enter basic information on age, sex, marital status, relationship to the head of the household, presence in the household of the individual's mother, father, and spouse, religion, whether the respondent worked or was in school, earnings in the last year, and highest level of education. Methods were intended to be sufficiently flexible to account for changes in household composition at each subsequent STAR resurvey. For individuals who had left the household between survey waves, information was collected on the reason for and date of departure and the person's current location. For individuals who joined the household since the previous survey, information was collected on the reason for and date of entry into the household.

Book 1: Household Expenditures and Characteristics

This book was answered by the household head, the spouse of the household head or by another person 18 years or older who is knowledgeable about household affairs. Module KS recorded information on expenditures for a variety of food and nonfood goods and services, including foods purchased and self-produced in the last week, personal care and household items bought during the last month, and durable goods bought in the last year. Module KH collected information about quantities and purchase prices for several staples. Module KR recorded information about ownership status and value of house and land, various household characteristics, and availability of and distance to certain public facilities. Module TR collected information about financial transfers, such as money for education or reconstruction, and in-kind transfers, such as food and building materials, into the household from different sources, including the government, NGOs, religious groups, family members, and friends and neighbors.

Book 2: Household Economy

Book 2 was answered by the household head, the spouse of the household head or by another person 18 years or older who is knowledgeable about household affairs. Modules UT and NT focused on household revenue, expenses, and value of household-owned agricultural and nonagricultural businesses. Module HU asked about the current value of household business assets (such as land, livestock, buildings, machinery, and transportation) as well as the value of any assets that were destroyed, sold, or given away. Module HR asked about the current value of household nonbusiness assets (e.g., land, livestock, jewelry), as well as asset ownership and ownership shares. Module HI asked about household-level non-labor income, by source. Module BT asked about money borrowed, received as a grant, or money that has been loaned out from the household.

Book 3: Adult Individual Book

This book elicited current and retrospective information from each household member age 15 and older.

Education history. Module DL recorded the highest level of education attended and highest grade completed, as well as basic literacy information. Information including the name and location of the most recent school attended was collected. Respondents were also asked about the highest level of schooling they expected to achieve.

Employment. Module TK asked about respondents' current employment. Employment was categorized according to industry and position within the job, and information was recorded on hours worked and earnings. Unemployment information was collected for individuals without work.

Individual non-labor income and assets. Module HR collected information on individual-level economic well-being. Respondents were asked about the current value of their nonbusiness assets (e.g., land, livestock, jewelry), as well as asset ownership and ownership shares.

Health status, physical performance, and usual practices. Module KK asked about general health status and recent health history and physical functioning. Module MA asked about morbidities in the past four weeks and about experience with conditions symptomatic of heart disease, diabetes, and high blood pressure. Module PB asked about smoking and consumption of fruits and vegetables.

Outpatient utilization and health diagnoses. Module RJ collected information on outpatient visits during the last four weeks, including name and type of provider. Module DK recorded data on injuries and accidents in the past year, and more recent obvious illnesses.

Marriage and fertility. Module KF collected information about age and date of first marriage and number of children. Unmarried respondents were asked expectations of marriage timing and number of children. Married female respondents were asked about contraceptive use.

Psychological health and attitudes. Module KP asked respondents about symptoms of depression and post-traumatic stress response. Respondents were also asked about participation in post-disaster counseling and the usefulness thereof.

Risk and Time Preferences. Module SI asked respondents about time preferences (choosing between hypothetical payments sooner or later) and risk preferences (choosing between hypothetical certain and uncertain payouts).

Individual economic shocks. Module GEI asked respondents how life changed since the tsunami, as well as expectations for how long it will take for things to return to normal. Respondents also provided their opinions on the relative quality of their lives on a 6-step ladder, before the tsunami, at the time of the interview, and five years in the future.

Migration history. Module MG collected information on the geographic mobility of individuals, as well as the causes and consequences of migratory movements, including short-stay and circulatory migration. Information was recorded about the respondent's location at birth, age 12, and each subsequent location where a move crossed a *desa* (village) boundary and lasted for 6 months or longer. For each move, data were collected on dates and locations, motivation for moving, and distance moved.

Social and kinship network. Module SN asked respondents about the number of family members and friends and neighbors that could be called upon for certain help or support. Specific questions were asked about the survival of critical family members.

Exposure. Module EX recorded data on measures of exposure to the earthquake and tsunami, such as feeling the earthquake, hearing people shouting, physically being swept away by the wave, and living afterward in destroyed areas.

Community participation. Module PM considered community development activities, which have long been a backbone of development in Indonesia. Respondents were asked about participation in, contributions of time and money to, and perceived benefits from, a slate of community development and reconstruction activities. Questions were included on participation in rotating credit schemes (*arisan*) and knowledge and use of credit sources.

Book 4: Information for Ever-Married Women

This book was administered in waves after STAR1 to collect information on children ever born, a full pregnancy history, and knowledge and use of contraceptives.

Book 5: Child Information

This book was administered to household members younger than 15. For children younger than 11, the mother, female guardian, or household caretaker answered the questions. Children between the ages of 11 and 14 were allowed to respond for themselves if they wished. Generally, each module paralleled a module in the adult questionnaire (book 3), with age-appropriate modifications. Modules that paralleled the adult questionnaire included: child's educational status and schooling disruptions (module DLA), employment status (module TKA), time allocation (module TAA), morbidities and general health characteristics (modules MAA and KKA), outpatient utilization (module RJA), exposure to the earthquake and tsunami (module EXA), child psychological health (module KPA), migration history (module MGA), social and kinship network (module SNA). Modules unique to Book 5 included module BAA, which collected information on the respondent's parents and other surviving kin, and module INA, which recorded data on type of institution in cases when the respondent lived in an institution at the time of interview.

Non-coresident family roster and transfers. Module BAA established parental survival status and, if still alive, locations of non-coresident parents. Non-coresident immediate kin (siblings, aunts/uncles, grandparents, cousins) are listed in this section, as well as any major assets inherited from parents.