A Method of Correcting for Misreporting Applied to the Food Stamp Program

Nikolas Mittag[*]

Harris School Of Public Policy Studies, University of Chicago

November 28, 2012

Abstract: Using administrative and survey data I show that survey misreporting leads to biases in common statistical analyses. Standard corrections for measurement error cannot remove these biases. I describe a method to obtain consistent estimates by combining parameter estimates from the linked data with publicly available data. This solves the problem that access to linked administrative data is usually restricted. I apply this method to study Food Stamp Program (now SNAP) receipt. I examine the degree to which this approach can be used to extrapolate across time and geography, in order to solve a common limitation of validation data. Administrative data on SNAP receipt and amounts linked to American Community Survey data from New York State show that survey data can misrepresent the program in important ways. For example, more than one billion dollars received are not reported in New York State alone. 40 percent of dollars received by households with annual income above the poverty line are not reported in the survey data, while only 13 percent are missing below the poverty line. The conditional density method I describe recovers the correct estimates using public use data only. I present evidence from within New York State that the extent of heterogeneity is small enough to make extrapolation work well across both time and geography. Extrapolation to the entire U.S. yields substantive differences to survey data and reduces deviations from official aggregates by a factor of 4 to 8 compared to survey aggregates.

Keywords: Measurement error, Survey errors, underreporting, food stamps, poverty.
JEL Classification Numbers: C15, C81, I32, I38

1 Introduction

Survey data provide the level of detail required in many economic analyses, but are known to suffer from misreporting of important variables and frequently do not include other required variables. Administrative data often contain accurate measures of these variables, but lack the detail and the number of covariates that enable researchers to answer many interesting questions. Creating validation data by linking administrative records to survey data combines the accuracy of the administrative records with the detail of the survey data and thereby overcomes both problems. I examine the patterns of food stamp receipt using administrative records from New York State linked to the American Community Survey (ACS). The linked data allow researchers to examine questions that cannot be answered with administrative data alone and to obtain more accurate answers than from survey data alone.

However, linked data are subject to strong confidentiality requirements and rarely accessible to the interested researcher. Additionally, administrative records are difficult to obtain and costly to link, so the linked data are often confined to short time periods or small geographic areas. I describe a method that allows researchers without access to validation data to obtain consistent estimates from public use data based on an estimate of the conditional density of the validated variables. The estimated density can be published without infringing upon confidentiality. This conditional density method reproduces the results from the linked data accurately from the public use data alone, so it solves the problem of access restrictions. I examine the extent to which this same method can be used to extrapolate across time and geography. This approach can mitigate the problem that validation data often only cover a short time period or small geographic area, but requires the conditional density to be the same in the data of interest and the validation data. I provide evidence that this assumption is a reasonable approximation in the case of food stamps by analyzing heterogeneity and extrapolation across time and geography within New York State. Using the conditional density method to extrapolate to the entire U.S. yields results that capture differences between

states and between metropolitan areas much better than the survey reports. The results from extrapolation are better aligned with what we know about food stamps from sources other than survey data.

In particular, I use administrative records on food stamp receipt and amounts received from New York State linked to the 2008-2010 ACS to evaluate the targeting of the program and its effect on poverty. The linked data reveal that in New York State alone, one billion dollars, a quarter of the total amount, paid to recipient households, is not reported in the survey data. With almost 800 million dollars not reported above the poverty line, underreporting leads survey data to understate the importance of the Food Stamp Program for households above the poverty line both in terms of recipients and amounts received, but overestimates recent growth rates among households with higher incomes. Survey reports also understate program participation among households in deep poverty (below 50 percent of the poverty line) by 20 percent. The linked data reveal a non-monotone relation between income and underreporting, so common methods to correct for underreporting are unlikely to solve this problem.

These findings underline that validation data, such as administrative data linked to a survey, are sufficient and often necessary to solve the problem of misreporting. However, the administrative records used to create the linked ACS data for New York State are confidential by law and can only be accessed by researchers with special sworn status at the U.S. Census Bureau. Data access requires approval by the New York Office of Temporary and Disability Assistance and the U.S. Census Bureau. The conditional density method described in this paper recovers the results from linked data accurately using only public use data and information that can be obtained from linked data and published without infringing upon confidentiality. This approach solves the problem of access restrictions and allows researchers to obtain accurate answers to detailed questions from survey data.

I also use the conditional density method to extrapolate across time within New York State and to the entire U.S. I can test how well extrapolation across time works, because the

linked data cover multiple years. Even though direct tests show that the conditional density changes moderately over time, analyses based on the conditional density method are still closer to those based on the linked data than survey reports. Extrapolation across states requires the assumption that the conditional density is the same in different states. This is not directly testable without validation data from an additional state. However, I provide evidence that most of the heterogeneity within New York State is explained by the covariates in the conditional density. This evidence is supported by the fact that extrapolation across geography within New York State works well. These tests suggest that extrapolation works very well in the case of food stamps and can additionally be improved by incorporating information from official aggregates, such as total dollars spent by region. The aggregate totals can also be used to assess the quality of the extrapolation to other states. I compare state and metropolitan area totals in the whole U.S. from extrapolation based on the conditional density from New York State to totals from survey data. I find that extrapolation reduces average error in total dollars received by a factor betwen 4 and 8 compared to survey reports.

Misreporting in survey data is by no means confined to government programs: Measurement error of a similar magnitude has been found in many other important variables, such as income and education. The errors are likely to lead to substantial biases in analyses that use these variables as well. If validation data exist or can be created, for example by linking surveys to other records such as tax or employer records, the conditional density method can be used to correct the problem of measurement error in these variables.

Section 2 reviews the evidence on measurement error in economic data, underlining that it is a substantial problem for economic analyses. Section 3 discusses consequences of misreporting for analyses of government programs. Section 4 introduces the linked data. Section 5 presents the conditional density method and shows how it could be applied to common problems. Section 6 describes the estimation of the conditional density. Section 7 uses the linked data to assess distributional effects of SNAP in New York State and shows that the conditional density method leads to the same conclusions. Finally, in section 8 I examine

3

how well the conditional density method solves the need to extrapolate across time and geography. Section 9 discusses implications and concludes.

2 The Problem of Measurement Error

Economic models frequently include variables that are known to suffer from measurement error. The main concern in this study is misreporting, i.e. measurement error due to inaccurate survey reports. This problem is known to be pervasive (see Bound, Brown and Mathiowetz, 2001, for an overview), affecting variables of key importance to economic models such as income (Bollinger, 1998; Bound et al., 1994) and even variables that seem straightforward to measure such as the highest degree obtained (Black, Sanders and Taylor, 2003). In more extreme cases, important variables are missing from the data entirely, for example information on assets in the Current Population Survey (CPS). The problem is usually transformed into a measurement error problem by replacing the missing variable with a noisy proxy. The methods in this paper also apply to missing variables, even if no proxy is available.

The economic literature often assumes "classical" measurement error, i.e. that errors are independent and identically distributed across units of observations and unrelated to any other characteristics (e.g., Fuller, 1987). This framework has been adapted from the natural sciences, where it is more appropriate because the the errors are due to imprecise measuring tools. Measurement error in economics arises primarily from human behavior, which makes the assumptions of the classical framework unlikely to hold. Several recent studies have examined the extent and nature of measurement error in economic data directly by obtaining validation data or indirectly through the use of multiple measurements or by comparing aggregate survey responses to aggregates from official statistics (e.g., Bound, Brown and Mathiowetz, 2001). Besides showing that measurement error is widespread and often large, these studies have confirmed that the assumptions of classical measurement error are oversimplified and wrong in most cases. Thus, arguments regarding the consequences of

measurement error that are often made in applied research (e.g., that it leads to attenuation or causes no bias when the dependent variable is affected) are likely to be wrong, for they only hold in the classical framework.

As an example, consider measures of inequality such as the 50/10 wage gap in Autor, Katz and Kearney (2008). They are overestimated under classical measurement error because it uniformly increases the variance of the income distribution. However, there is evidence that measurement error in earnings and wages is mean reverting (Bound, Brown and Mathiowetz, 2001) and some evidence that this is particularly driven by over-reporting of low wage earners. If this is indeed the case, rather than overstating inequality, survey measures understate the 50/10 wage gap. Similarly, mean reverting measurement error understates estimates of any other wage gap, causing bias even when income is the dependent variable. I discuss the likely consequences of measurement error for other studies in the next section and show how to correct them in section 5.

While we know that bias from non-classical measurement error is often severe, little is known about its direction and likely magnitude. The literature on parameter bounds (e.g. Horowitz and Manski, 1995; Bollinger, 2003; Gundersen and Kreider, 2008) finds that without additional restrictions on the measurement error, the bounds tend to be large. Wide bounds imply that the bias from non-classical measurement errors in these applications could be large and may well overturn conclusions drawn from biased point estimates that ignore the problem. For point estimates, the bias from non-classical measurement error has only been derived for a few special cases (e.g., Carroll et al., 2006; Meyer and Mittag, 2012). A larger literature has focuses on specific applications for which it is possible to obtain unbiased estimates (e.g., Poterba and Summers, 1986; Hausman, Abrevaya and Scott-Morton, 1998). They characterize the consequences of non-classical measurement error and emphasize its impact on substantive conclusions. However, whether their findings generalize to other cases and applications is an open question.

Common solutions to the problem of measurement error impose restrictions on the dis-

tribution of the measurement error (e.g., Hausman, Abrevaya and Scott-Morton, 1998; Hong and Tamer, 2003) or invoke an independence assumption and use instrumental variable methods (Hausman et al., 1991; Hu and Schennach, 2008) or repeated measures (Li, 2002; Schennach, 2004). Much progress has been made in estimating parameters under these assumptions, but for several cases there is evidence that the assumptions are violated (Bollinger and David, 2005; Meyer, Goerge and Mittag, 2012) and that their failure may lead to large biases (Black and Smith, 2006; Meyer and Mittag, 2012). An alternative solution is to obtain validation data. Validation data can be created by re-visiting a (sub)sample of the survey respondents (Black, Sanders and Taylor, 2003) or by linking the survey data to a more accurate data source such as administrative records (e.g., Marquis and Moore, 1990; Bound and Krueger, 1991; Meyer, Goerge and Mittag, 2012). Bound, Brown and Mathiowetz (2001) provide an excellent overview of validation studies. Estimators that use validation data to consistently estimate models in the presence of arbitrary measurement error have been proposed and seem to work well for very general classes of estimators (e.g., Chen, Hong and Tamer, 2005). The key advantage of these estimators is that they avoid any assumptions on the measurement error and its relation to the covariates. Instead, the crucial assumptions are that the validation data are accurate and that the relation between the observed and validated variables is the same in the non-validated sample. The increased awareness of misreporting and advancements in record linkage have greatly increased the number of validation studies that are conducted, but the two main problems pointed out by Bound, Brown and Mathiowetz (2001) remain: First, the validated data are usually confidential and therefore only available to the researchers involved in the original data collection. Second, validation data commonly do not exist for the exact population of interest. They are often only available for a few years or a small geographic area, such as administrative records from some states rather than the entire country (e.g., Marquis and Moore, 1990; Meyer, Goerge and Mittag, 2012).

3   Misreporting of Government Programs and its Consequences

In this section I show that survey data are insufficient to answer important questions on government programs, because of large and systematic underreporting. I focus on the Food Stamp Program (now SNAP) to illustrate the likely consequences of misreporting for several studies of this program and why previous attempts to correct them are insufficient. Section 5 shows how the method described in this paper could be used to obtain consistent estimates in these analyses. The studies I discuss employ common econometric methods, so the bias and solutions are likely to be similar in other applications.

The Food Stamp Program has grown to be the largest anti-poverty program in the U.S., spending $72 billion and serving 45 million people in fiscal year 2011. It is the only federal transfer program with almost universal eligibility based on income, so it is often the only resource for people that do not qualify for other programs. It was introduced to ensure access to a healthy diet for everyone and is still primarily thought of as a safeguard against hunger and malnutrition, even though there is evidence that it has shifted towards less poor households (e.g., Scholz, Moffitt and Cowan, 2009). Accurate information on this and other government programs is essential for understanding its impact on the target population, as well as to improve its design and avoid wasteful or misdirected spending. However, the results in this paper show that survey reports skew the evidence on the Food Stamp Program. For example, I find that the program impact is more concentrated above the poverty line, but surveys also underestimate its effect on deep poverty and distort which households benefited most from the recent expansion. Thus, the survey data may well misguide our judgment on, for example, whether or not we view the expansion of the program favorably. Understanding who benefits from the program also has implications for the most efficient way to run it with regard to eligibility criteria, efforts to encourage take-up or the question of whether a cash transfer would be preferable to in-kind benefits (Currie and Gahvari, 2008).

Analyses of survey reports will mislead the understanding and administration of government programs, because underreporting of government programs is severe and systematically related to observable characteristics. Validation studies have found high rates of false negatives (recipients that do not report food stamp receipt) in all surveys that have been examined: Between 20 and 30 percent in the 1984, 2001 and 2004 panels of the Survey of Income and Program Participation (SIPP) (Marquis and Moore, 1990; Meyer, Goerge and Mittag, 2012) and up to 50 percent in the 2002-2005 CPS Annual Social and Economic Supplement (Meyer, Goerge and Mittag, 2012). High rates of false negatives were also found in the 2001 predecessor of the ACS (between 30 and 40 percent, see Lynch et al., 2007; Meyer, Goerge and Mittag, 2012), which led to a revision of the question in the final survey (Hisnanick, Loveless and Chesnut, 2007). I find that 26 percent of recipients do not report receipt in the 2008-2010 ACS data I use below. As in other studies, false positive rates are low: only 1 percent of non-recipients report food stamp receipt. All studies find that misreporting is related to other observable characteristics such as income and household composition. The problem of misreporting is not confined to food stamps. Meyer, Mok and Sullivan (2009) demonstrate that a similar extent of misreporting plagues survey reports of other government programs. Commonly used surveys fail to capture a substantial share of program participation, so steps should be taken to improve the accuracy of the data. One way to do so is to amend survey data with administrative data.

This is particularly desirable for the ACS because it does not contain reported amounts for non-cash programs at all. There is considerable interest in program amounts in the ACS, for example, to compute poverty measures for sub-state areas (e.g., Levitan and Renwick, 2010). An advantage of the conditional density method I present here is that it can also be used when a variable is missing entirely if the variable is available in another data set. This data can be linked (administrative) data, but the approach can be applied whenever the missing variable is available in a data set that has other variables in common with the data of interest (at the cost of an assumption on heterogeneity to be made precise below).

Misreporting presents a serious problem even to simple analyses of receipt and spending of government programs. The high rate of underreporting causes analyses based on survey data to understate the take-up of the program. Even worse, the relationship to demographic variables and income affects our understanding of who benefits from the program. For example, reporting declines as income increases, which makes survey data understate the extent to which government programs assist people above the poverty line. Misreporting is known to change over time, causing analyses based on survey data to confound changes in the program with changes in reporting. Most recent papers that evaluate government programs acknowledge that misreporting influences their results (e.g., Scholz, Moffitt and Cowan, 2009; Sherman, 2009; Tiehen, Jolliffe and Gundersen, 2012).

However, only few attempt to correct their estimates for misreporting. The most common way to do so is to adjust the number of recipients or amounts to match official aggregates. This is often done by assigning program receipt to the non-recipients with the highest predicted probability of receipt and increasing amounts received proportionally for everyone until totals match official aggregates (e.g., Scholz, Moffitt and Cowan, 2009). Other approaches use program rules to determine eligibility before predicting probabilities (Sherman, 2009) or scale up aggregate numbers (Meyer, 2010). There is ample evidence that these corrections are insufficient and the results in this paper suggest that they may well increase the bias. Most importantly, all corrections assume that misreporting is unrelated to other covariates. Validation studies clearly reject this assumption, even when slightly relaxed as in Sherman (2009). Furthermore, the parameter estimates used to assign the probability of receipt are biased because they are estimated using the misreported data. This tends to overstate receipt by those likely to report while still understating receipt by groups that are unlikely to report (and is amplified if receipt is assigned to those most likely to receive rather than probabilistically). Finally, these approaches can only correct net underreporting because they cannot correct for false positives.

The effects of misreporting in more complex econometric models are complicated and

corrections for misreporting are rarely attempted. However, as I show in section 5, obtaining consistent estimates from public use data is simple if information about the relation between true and observed variables is available. Studies of program take-up (e.g., Currie and Grogger, 2001; Haider, Jacknowitz and Schoeni, 2003; Currie, 2004) use reported receipt as the dependent variable in binary choice models. Bollinger and David (1997) and Meyer, Goerge and Mittag (2012) find significant bias due to misreporting in probit models. How much of the difference in take-up between the elderly and the non-elderly is due to a decline in actual take-up and how much is due to an increase in underreporting with age is still an open question. Binary choice models are one of the few cases where the bises from misreporting have been derived (Meyer and Mittag, 2012) and some corrections have been proposed (Bollinger and David, 1997; Hausman, Abrevaya and Scott-Morton, 1998).

No such results exist for other econometric models, so it is hard to assess the impact of misreporting on other studies such as those concerned with program effects on nutritional intake (Butler and Raymond, 1996; Hoynes and Schanzenbach, 2009) or labor supply (Fraker and Moffitt, 1988; Hagstrom, 1996; Keane and Moffitt, 1998). The consequences are likely to be similar to the effect of control group substitution in the program evaluation literature: Recipients that do not report receipt were "treated" with food stamps, but appear "untreated" in the data. This tends to understate the difference between recipients and non-recipients and thereby bias the program effect towards zero, but this tendency could be overturned by other factors.

As an example, consider Keane and Moffitt (1998), who estimate a structural model of labor supply and participation in three welfare programs (food stamps, AFDC including Medicaid benefits and subsidized housing). They include stigma parameters in their participation models to allow for incomplete take-up in a model where households maximize their utility subject to the budget constraint implied by their wage rate and the program rules. One would expect bias in the parameters of the participation equations that is similar to the bias in the studies of program take-up above: Receipt is underreported, so fewer eligibles

10

participate, which makes the data suggest a stronger deterrence effect. Thus, one would expect misreporting to lead to an overstatement of the stigma of program participation. For labor supply, it is instructive to consider a much simpler case in which there is no selection and misreporting is not related to any relevant characteristics. Simple algebra yields the relation between the observed effect of (binary) participation on labor supply ($\hat{\Delta}$) and the true effect (Δ) as $\hat{\Delta} = (1 - \alpha_{01} - \alpha_{10})\Delta$, where α_{01} is the share of non-recipients among those that report receipt and α_{10} is the share of recipients among those that do not report receipt. Even though the 1984 SIPP used by Keane and Moffitt (1998) has one of the highest reporting rates[1], it would understate the true effect on labor supply by 15 percent in this simple case. While it is in line with these calculations that Keane and Moffitt find a large effect of stigma and a low impact on labor supply, there are many reasons to believe that the biases are more complex than this simple illustration suggests.

However, the examples underline that even low rates of random misreporting can lead to meaningful bias and that the size and direction of the bias is difficult to assess. I describe a simple method to obtain consistent estimates using only public use data and a set of parameters that could be estimated from linked administrative data and made publicly available in section 5 after presenting the data in the next section.

4 Linked Administrative Data

In this section, I argue that validation data have the potential to solve the problems caused by misreporting. I introduce the validation data used in the analysis below, showing that linking survey data to administrative records is a promising way to validate survey data. However, validation data are rarely available, either due to access restrictions or because they do not cover the population or time period of interest. I describe a method to solve these problems in sections 5 and 6.

---

[1]Assuming the misreporting rates in their data (wave 4) are similar to those Marquis and Moore (1990) find in wave 2.

Even though it has long been recognized that validation studies provide valuable information on measurement error and allow researchers to correct estimates of common models, few such studies have been conducted. One reason is that it is costly to re-visit a random sample of the population to obtain a response that is more accurate but may still be prone to error. Another approach is to link existing survey data to records that are believed to be more accurate, such as administrative records of program participation, wage data from companies or tax records. However, it required recent improvements in record linkage and computational power to link them to survey data at a scale and accuracy that makes them a promising solution to the problem of misreporting.

The data I use were created by linking administrative records on food stamp benefits from the New York State Office of Temporary and Disability Assistance to the New York State sample from the 2008-2010 ACS. The administrative records contain monthly actual payments made for all recipient households in New York State, so the accuracy of the numbers is very high: The overall total from these records differs from actual aggregate outlays by less than 0.001 percent. Matches are made based on a Personal Identification Key (PIK), which is obtained for 99 percent of the administrative records and more than 90 percent of the individuals in the survey using the Census Bureau's Person Identification Validation System (NORC, 2011). Following Meyer and Goerge (2011), I take the imperfect match rate into account by multiplying the weights by the inverse of the predicted probability of any household member having a PIK. Appendix table A1 reports the parameter estimates used in this correction, appendix table A2 provides summary statistics for the linked data. Meyer and Mittag (2013*b*) describe the linking process and assess the accuracy of the data in more detail.

All analyses in this paper are conducted at the household level using the ACS household definition. I consider a household to receive food stamps if anyone received food stamps in the 12 months before the interview date according to the administrative data. Given the high match rate and the quality of the administrative records, I consider the administrative data

to be "truth" even though there may be small inaccuracies due to mismatches or erroneous payments. This assumption may not be warranted with other linked records, for example if the administrative data are also based on reports, as with tax records. Nonetheless, such records may still be a better measure than survey reports and both measures can be combined to improve inference as in Abowd and Stinson (2011) or Kapteyn and Ypma (2007).

Meyer and Mittag (2013a) examine misreporting using the same data, which is substantial (26 percent false negatives and 1.19 percent false positives). It is systematically related to covariates of importance for this analysis such as income, spell length and amounts received. Consequently, analyses based on survey data are biased and administrative data will lead to more accurate conclusions about the effects of food stamps on poverty. However, the problems of access and limited geographic coverage loom large with linked administrative data. By law, administrative records are subject to strong confidentiality rules in order to protect the privacy of program participants. This makes it difficult for researchers to access the linked data and requires all results to pass a disclosure review process. Like most previous validations of program data, the administrative records are obtained from state agencies and therefore only cover one state. Previous studies cover different time periods for different states and often do not contain amounts received. Studies on the impact of food stamps at the national level or in years where validation data are not available need to either extrapolate from the existing data or rely on reports. The conditional density method I introduce in the next section can be used to solve these problems.

5   Conditional Density Method

While analyses based on validation data avoid the problem of measurement error, performing the analysis on the validated data directly is often infeasible because of confidentiality restrictions or undesirable if the validation data do not cover the entire population or time period of interest. The conditional density method introduced in this section allows re-

searchers without access to validation data to obtain consistent estimates from public use data based on an estimate of the conditional density of the validated variables, which could be published without violating confidentiality. Section 7 presents evidence that the bias from misreporting matters substantively and the conditional density method solves the problem of access, since estimates based on the linked data and the conditional density method are virtually identical. Section 8 examines how well the conditional density method can be used to extrapolate across time and geography. Obtaining consistent estimates for other time periods or other states requires more stringent assumptions, but I provide evidence that it works well for the case at hand.

The conditional density method discussed here builds on an insight from the measurement error (e.g., Hsiao, 1989; Li, 2002; Chen, Hong and Tamer, 2005) and data combination literature (e.g., Ridder and Moffitt, 2007; Ichimura and Martinez-Sanchis, 2009). The consistency of the estimators below follows from arguments analogous to theirs. However, their estimators require the validation data to be available to the researcher, which is often not the case. The conditional density method is closely related to Brownstone and Valletta (1996), who propose a more restrictive framework for a linear model and use a different method to estimate the outcome model. A more detailed description of the general approach can be found in Mittag (2012). The researcher wants to estimate an econometric model that depends on variables only available in the administrative data X^A and (potentially) other covariates Z. In the application below, X^A includes food stamp receipt and amounts received and Z contains variables such as income relative to the poverty line. The problem is that besides Z, the public use data only contain reported versions X^R of some or all of the administrative variables. For the ACS, X^R includes reported receipt only, since amounts are not reported. Even when available, the reports are subject to error, so estimating the model on the public use data would lead to biased inference. The linked data contain (X^A, X^R, Z), so they can be used to estimate the model of interest. However, there are many cases in which that is not desirable or possible, for reasons of confidentiality or coverage as discussed above. If the data

do not stem from linking records, they may also not include all of the covariates of interest. The intuition behind estimation from the public use data is that we can learn enough about the relationship between the true and the reported variables from the linked data to back out the parameters of interest from the public use data. This relationship is given by the conditional density of X^A conditional on X^R and Z, $f_{X^A|X^R,Z}(X^A|X^R,Z)$, an estimate of which can be summarized by a vector of parameters that can be published without infringing confidentiality. Once this vector of parameters has been published, consistent estimates of the model of interest are simple to compute without access to the linked data.

More formally, the objective function of any estimator that contains unobserved variables X^A can be re-expressed in terms of X^R, Z and $f_{X^A|X^R,Z}(X^A|X^R,Z)$ using the law of total probability. For example, a (non-linear) regression line $Y = g(X^A, \beta)$, implies that

$$\mathbb{E}(Y|X^R, Z) = \int g(x^A, \beta) \cdot f_{x^A|X^R,Z}(x^A|X^R, Z)dx^A \tag{1}$$

The fact that the expression on the right only depends on variables observed in the public use data has been used to derive consistent estimators for a large class of models. It covers maximum likelihood and minimum distance models (Hsiao, 1989) as well as (potentially non-linear or non-smooth) moment based models (Chen, Hong and Tamer, 2005). All of these models can also be estimated by the conditional density method. It is closely related to Chen, Hong and Tamer (2005), who propose an estimator that includes a function of the conditional density that depends on the particular model being estimated. They use both data sources simultaneously in a semi-parametric joint estimator, which requires access to the validation data and the implementation of a model-specific estimator. If both data sets are available, their estimator has the usual advantages of joint estimators. I propose to use a two-step estimator based on a parametric estimate of the conditional density that solves the problem of access and is simple to implement. The first step consists of the estimation of the conditional density and requires the linked data, the second step estimates the outcome

model and does not require the linked data. Mittag (2012) establishes consistency of this estimator and compares it to other approaches.

In order to illustrate the estimation procedure first, I assume for now that an estimate of the conditional density, $\hat{f}_{X^A|X^R,Z}(X^A|X^R,Z)$, is available to the researcher in the form a of vector of parameters. I defer estimation of the conditional density to section 6. As equation (1) suggests, the parameters of interest can be estimated by integrating X^A out of the objective function of the estimator. I do this by simulation, i.e. by generating a sample of D draws from $\hat{f}_{X^A|X^R,Z}(X^A|X^R = x_i^s, Z = z_i)$ for every observation $i = 1...N$ in the public use data[2]. This yields a data set containing $D \times N$ observations of (X_{id}^A, X_i^R, Z). Consistent estimates of the parameters of interest can be obtained from this data set in the same way as they would have been obtained if X^A were available in the original data. Mittag (2012) shows why this simulation yields consistent estimates and discusses more general cases and simplifications. Standard errors need to be adjusted for the estimation of the first stage parameters as in Newey and McFadden (1994) and simulation error as in McFadden (1989) or Li (2000). However, the correction for simulation error can be made arbitrarily small and the correction for estimation error of first stage parameters is very small with commonly used sample sizes. In summary, obtaining estimates from the public use data can usually be done in three simple steps if an estimate of the conditional density is available: First, the researcher uses the parameters of the conditional density to generate D draws for every observation. Second, (s)he runs the original program on the expanded data set. Third, (s)he adjusts the standard errors if the corrections are found to matter.

To illustrate how the conditional density method can be applied to purge bias in common analyses, consider the studies discussed at the end of section 3. As in my analysis, the statistics of interest in previous analyses of effects of food stamps on poverty are often univariate, such as poverty rates or amounts received by demographic sub-groups. For example, Sherman (2011) adds amounts received from welfare programs to family income in

---

[2]Stata programs to simulate samples from common conditional densities are available on my website.

the CPS and calculates poverty rates with and without transfers. Scholz, Moffitt and Cowan (2009) add program amounts to household income in the SIPP to analyze the poverty gap. If an estimate of the conditional density of transfer amounts received had been available for the data they use, they could have drawn D program amounts for each observation from the conditional density and calculated a value of total income y_{id} from each of these draws. A consistent estimate of the poverty rate is then given by $(D \cdot N)^{-1} \sum_{d=1}^{D} \sum_{i=1}^{N} \mathbf{1}(y_{id} < PL_i)$, where PL_i is the poverty line for household i. The (average) poverty gap can be calculated as the average of $y_{id} - PL_i$ among the observations for which $y_{id} < PL_i$. The estimators for binary choice models with misreporting (Bollinger and David, 1997; Meyer and Mittag, 2012) require predicted probabilities of misreporting for each individual. Analyses of program take-up such as Haider, Jacknowitz and Schoeni (2003) could avoid the bias in their estimates by calculating these probabilities from $\hat{f}_{X^A|X^R,Z}(X^A|X^R, Z)$ (either analytically or by simulation) and using them in one of these estimators.

It is simple to account for misreporting even in as complex a model as the one estimated by Keane and Moffitt (1998). Their structural model implies the following log-likelihood function[3]

$$\ell(\theta) = \sum_{i \epsilon E} \sum_{j=1}^{J} k_{ij} \ln \left[P(j|\theta, Z_i, w_i)\phi(w_i|\theta, Z_i) \right] + \sum_{i \epsilon U} \sum_{j=1}^{J} k_{ij} \ln P(j|\theta, Z_i) \qquad (2)$$

where E is the set of employed individuals, U the set of unemployed individuals and Z_i are observed covariates. $k_{ij} = 1$ indicates that person i chooses combination j of program participation and work hours, and is thus implied by X^A and Z. $P(j|\theta, Z_i, w_i)$ and $P(j|\theta, Z_i)$ are the probabilities of choosing participation and work hour combination j and $\phi(w_i|\theta, Z_i)$ is their model of wage determination. θ can be estimated consistently in the presence of misreporting by taking D draws of program receipt from the conditional density, calculating the implied indicators k_{ijd} and maximizing the same log-likelihood function over the simulated

---

[3]Slightly adapted to my notation.

sample, i.e., by maximizing

$$\tilde{\ell}(\theta) = \sum_{d=1}^{D} \sum_{i \epsilon E} \sum_{j=1}^{J} k_{ijd} \ln \left[P(j|\theta, Z_i, w_i) \phi(w_i|\theta, Z_i) \right] + \sum_{d=1}^{D} \sum_{i \epsilon U} \sum_{j=1}^{J} k_{ijd} \ln P(j|\theta, Z_i) \quad (3)$$

Keane and Moffitt (1998) solve the problem of maximizing (2) and the same approach can be used to maximize (3). Thus, obtaining consistent estimates from public use data is simple once the original estimation problem has been solved.

Finally, consider the estimation of wage gaps when reports of wages are subject to error as discussed above. Wage gaps, such as the gender wage gap, are usually estimated by a linear regression with wages as the dependent variable. If an estimate of the conditional density of true wages can be obtained (for example by linking tax or company wage records to survey data), the conditional density method can be applied as in the previous examples. However, a simpler and more efficient estimator in this case is to impute the conditional mean (Schafer and Schenker, 2000), which can be calculated or simulated from the conditional density for each observation. Replacing reported wages by this conditional mean in the same linear regression yields a consistent estimate of the wage gap.

These examples illustrate that the estimation procedure I discuss does not require access to the validation data and is simple to implement, which are its main advantages over the joint estimators considered in the measurement error and data combination literature (e.g., Chen, Hong and Tamer, 2005; Ichimura and Martinez-Sanchis, 2009). If the researcher has access to both data sets, the standard advantages of joint estimators may make these approaches preferable, but obtaining access to both data sets is often not an option. Section 8 shows some additional practical advantages of the approach taken here, such as easier ways to incorporate additional information. Another approach to solve the problem of access is to create and release synthetic data, i.e. repeated draws from an estimate of the joint distribution of the entire data (see e.g Rubin, 1993; Little, 1993; Abowd, Stinson and Benedetto, 2006; Drechsler and Reiter, 2010). Synthetic data methods use Bayesian methods to esti-

mate the outcome model, but the draws from the distribution could be analyzed in the same framework. The key difference between the methods is that the synthetic data approach provides a fixed set of draws from the estimated density, while the conditional density method makes the estimate of the density available to the researcher. The main disadvantage of the synthetic data approach for the purpose of this project is practical in nature: Statistical agencies appear to be more reluctant to allow the release of an entire data set than to release a vector of parameter estimates. The parameter estimates also allow the researcher to extrapolate to populations that have not been validated, which cannot be done based on the synthetic data. Finally, parameter estimates give the researcher more flexibility to choose the appropriate number of draws or to impose constraints based on additional information as I do in section 8.

6   Estimating the Conditional Density

This section discusses how to obtain an estimate of the conditional density that can be used to correct bias from non-classical measurement error as described in the previous section. After a brief discussion of how this can be done in general, I focus on estimating the conditional density of food stamp amounts and receipt in the linked administrative data from New York State.

The conditional density method can be applied whenever any consistent estimate of the conditional density can be obtained. I observe (X^A, X^R, Z) for a random sample of households from New York State, so the conditional density can be estimated by standard maximum likelihood. A more complex model may be required if the linked data come from a stratified sample or suffer from problems such as sample selection. In other cases, the assumption that the administrative records represent "truth" may be too strong, as in Abowd and Stinson (2011). Nonetheless, the assumptions they make in their analyses identify the conditional density of income conditional on survey reports, which could be used in the es-

timators discussed above. In the absence of validation data, the conditional density may be identified based on repeated measures (e.g Schennach, 2004; Bonhomme and Robin, 2010) or independence assumptions (e.g Hu and Schennach, 2008; Hu and Ridder, 2012). One could opt for a non-parametric sieve estimator (Chen, 2007), but in the application to food stamp amounts the parametric estimator I use is flexible enough. In general, maximum likelihood offers the standard advantages of efficiency and the parametric framework prevents the tails of the density from becoming unreliable. This comes at the expense of the parametric restrictions, but they can be relaxed almost arbitrarily by using mixtures of parametric families. Maximum likelihood estimation also offers a framework of well-established procedures to test the parametric restrictions and tools to guide specification search.

The estimation of the conditional density of food stamp receipt and amounts is simplified by transforming it into a univariate problem: an amount of zero is equivalent to no receipt. Consequently, I estimate the density of food stamp amounts conditional on reported receipt and other covariates, allowing for a mass point at zero that implies no receipt. The probability of receiving an amount of zero is allowed to be a function of the covariates and is modeled using the standard normal distribution function. Conditional on reports and the covariates, I find that a left truncated normal density in which the mean is a function of the covariates fits the data well:

$$
f_{X^A|X^R,Z}(X^A|X^R = x^R, Z = z) = \begin{cases} \Phi(x^R\alpha + z\beta) & \text{if } X^A = 0 \\ \left[1 - \Phi(x^R\alpha + z\beta)\right] \frac{\phi(X^A;x^R\gamma+z\delta,\sigma)}{1-\Phi(\tau;x^R\gamma+z\delta,\sigma)} & \text{else} \end{cases} \tag{4}
$$

$\alpha, \beta, \gamma, \delta, \sigma$ and τ are parameters to be estimated from the data. This model is similar to a tobit type II model where the observed varible has an additional truncation point τ. A right truncation point or a mixture of densities does not improve significantly over this model, neither does relaxing the rate of decay of the tails by using a t-distribution.[4] While the model restricts the density to be a truncated normal at every value of $x^R\gamma + z\delta$, this does

---

[4]Programs to estimate these conditional densities are available from my website.

not force the marginal density of food stamp amounts to be a truncated normal.

The choice of conditioning variables Z is a standard model specification issue, so one can use the regular tests to find the set of variables that matters. Additional guidance for similar cases is given in Steuerle-Schofield et al. (2012). Summary statistics for the variables I use are provided in appendix table A2. Estimates of selected parameter from the 2010 linked ACS are presented in table A7. The complete parameter vector and estimates for 2008 and 2009 are reported in appendix table A8. Since the model is similar to a tobit type II model, the interpretation of the parameters is similar: The parameters of the truncated normal, γ, δ, σ and τ, determine the mean and variance of amounts received if the household participates. The parameters of the mass point, α and β, are similar to the probit parameters of a selection model, with the slight difference that they determine the probability that a household does *not* receive food stamps.

Table 1: Selected Parameter Estimates of the Conditional Density, New York State 2010

	Mass Point		Amounts	
	Coef.	*SE*	*Coef.*	*SE*
SNAP receipt reported	-2.50	0.002	1,453.0	6.0
SNAP receipt imputed	-0.42	0.005	562.7	15.5
# of persons in HH	-0.21	0.001	477.3	2.9
# of children in HH	-0.02	0.002	862.1	4.1
Income relative to Poverty Line Slopes				
$\leq$ 50 %	-0.07	0.017	461.8	31.0
50-100 %	-0.90	0.018	-733.0	31.2
100-150 %	0.35	0.016	-410.3	36.2
150-200 %	0.39	0.017	-702.5	47.3
>200 %	0.04	0.000	-36.1	1.9
Sigma	2034	2.0		
Left Truncation Point	16	0.003		
# of observations	101,680			

Notes: Parameter estimates of truncated normal conditional distribution. The columns labeled "Mass Point" determine the probability that the household does *not* receive SNAP and can be interpreted like probit coefficients. Those labeled "Amounts" determine amounts conditional on receipt. Demographic characteristics refer to the household head. All analyses use household weights adjusted for PIK-probability.

The model conditions on reporting status, so the parameters are informative about the

factors that drive misreporting. However, questions about the causes of misreporting are beyond the scope of this paper and discussed in Meyer and Mittag ($2013a$). Here, I am mainly interested in how well the covariates predict true food stamp receipt and amounts. The results are not surprising: Reporting status, including non-response, is a strong predictor of receipt and amount received. Household composition and income capture most of the remaining variation. Since income relative to the poverty line plays a key role in the analysis below, I model it as a linear spline with five segments. The results in section 7 and 8 are robust to moderate changes in the specification of the conditional density. All parameter estimates and the corresponding variance matrices can be downloaded from my website to correct other analyses of food stamps.

While standard specification tests make it easy to decide between two parametrizations of the conditional density, it is harder to assess whether the conditional density fits the data well enough to provide reliable estimates from the public use data. A simple way is to simulate draws of X^A from the conditional density in the linked data. A well specified model should replicate the observed distribution of X^A accurately, which can be assessed graphically or using a Kolmogorov-Smirnov test. With the sample size here, this test rejects differences that are irrelevant in practice, but the Kolmogorov-Smirnov distance is still informative about the difference between the simulated and the actual values. However, this only tests the marginal distribution of X^A and not its relation to other variables. The conditional Kolmogorov test proposed in Andrews (1997) tests whether an observed sample comes from a specific conditional distribution, so it can be used to assess whether X^A as observed in the linked data come from the estimated conditional density. This provides a test of the full model and has been implemented in the programs mentioned above. However, it is computationally burdensome in large samples since it requires bootstrapping critical values from a Gaussian process. More work on these issues is desirable, but the results below suggest that the method can work very well despite a lack of more sophisticated specification tests.

7 Re-assessing the effect of SNAP in New York State

In this section I analyze the impact of the Food Stamp Program in New York State according to different data sources: Reports from the ACS and CPS as well as the restricted linked data. I find that survey reports significantly understate the impact of the program above the poverty line and participation by households in deep poverty. Conversely, survey reports overstate the growth rate of SNAP among households with higher incomes. In order to show that the conditional density method solves the problem of access to the confidential linked data, I replicate the entire analysis using public use data and the estimate of the conditional density from section 6. The simulated results from the public use data are virtually identical to those from the linked data, so the conditional density method works well within sample and thereby solves the problem of access.

As discussed above, I consider the linked data to be accurate, so they provide the point of reference to which I compare the results based on the survey reports and the simulations in the public use data. I provide comparisons to the ACS reports whenever possible. Since the ACS does not include program amounts, I compare the analyses that require amounts to the CPS Annual Social and Economic Supplement. While there may be some concern about the comparability due to factors such as higher non-response in the CPS, it is a sample from the same population (after excluding group quarters). It also is the main data set that is used to conduct such analyses in practice and therefore is an important point of comparison[5]. To match the reference period of the ACS, I pool the two corresponding years of the CPS. All years refer to ACS survey years, so comparisons to aggregate numbers need to take into account that they include the previous calendar year as well.

There are two versions of the ACS, a restricted internal version and the public use microdata sample (PUMS). Summary statistics for both versions and the CPS are provided in appendix tables A2-A6. The linked data were created from the restricted version, but

---

[5]The SIPP is an alternative at the national level, but cannot be used here because it is not representative for a single state.

the conditional density method would be applied to the PUMS data in practice. Some differences between the simulated estimates and the linked data may be due to differences in the data, but I have repeated all analyses using the internal ACS file. The results are virtually identical and thus not presented below.[6]

Table 2: Food Stamps in New York State by Data Source, 2010

	(1) ACS PUMS Reports	(2) CPS Reports	(3) Linked Data	(4) Simulated ACS PUMS
Total amount received (in billion $)	-	2.91	3.96	3.94
Average amount received per HH	-	388	550	549
Number of HH receiving SNAP (×1000)	996	871	1,280	1,258
...as percentage of population	13.8%	11.6%	17.8%	17.5%
Poverty rate among recipients	48.3%	53.6%	43.3%	43.6%
...including FS benefits	-	37.7%	29.3%	31.1%

Notes: All analyses conducted using household weights (adjusted for PIK probability in the linked data)

Table 2 provides a first impression of the Food Stamp Program according to different data sources. Comparing what we learn about the program from survey reports in columns (1) and (2) to the results obtained from linked data in (3) and the conditional density method in (4) immediately underlines the main points made in this paper: Reports and administrative data differ substantially due to misreporting and the conditional density method can recover the correct estimates from the public use data. The CPS underestimates the total amount of food stamps received in New York State by more than one billion dollars, which leads research based on the CPS reports to understate the impact of the program. Rows 2 and 3 show that more households benefit from the program than we are lead to believe by the reports. The estimates from the linked data and the simulated PUMS estimates increase the number of beneficiaries in the CPS by almost 50 percent. The ACS reports are more accurate, but still increase by 30 percent when taking misreporting into account. In contrast to the meaningful differences between the linked data and the reports, the linked data and the figures from the

---

[6]I focus on results from 2010, results using the 2008 and 2009 data are qualitatively similar and available upon request. Due to the large sample sizes, standard errors are too small to affect the results substantively throughout.

conditional density method are virtually identical. The last two rows of table 2 take a first look at the effect of food stamps on poverty, showing that households above the poverty line are much less likely to report food stamp receipt. This leads both surveys to overstate the fraction of recipient households that are poor and the CPS to understate the percentage of recipients moved out of poverty. The fraction of recipient households moved out of poverty is similar in all columns of table 2. However, the number of recipient households is much lower in the CPS, so these figures do not capture the difference in the number of households moved out of poverty, which is 26 percent larger in the linked data than in the CPS.

Table 3: Food Stamps by income relative to the poverty line, New York State 2010

Income in % of HH poverty line	$\leq 50\%$	50-100%	100-150%	150-200%	$\geq 200\%$
Percentage of households receiving food stamps					
ACS PUMS Reports	46.2%	54.9%	32.8%	17.4%	4.2%
CPS Reports	40.4%	51.6%	27.2%	15.5%	2.3%
Linked Data	55.5%	62.5%	39.5%	24.4%	6.7%
Simulated ACS PUMS	53.9%	61.6%	38.7%	23.7%	6.6%
Average amount received					
CPS Reports	1,740	1,710	860	462	58
Linked Data	2,305	2,043	1,174	680	170
Simulated ACS PUMS	2,240	2,048	1,178	669	166
Total amount received (in million $)					
CPS Reports	772	953	594	287	302
Linked Data	809	1,174	718	399	860
Simulated ACS PUMS	820	1,165	716	395	839
Share of SNAP budget received					
CPS Reports	26.5%	32.8%	20.4%	9.9%	10.4%
Linked Data	20.4%	29.7%	18.1%	10.1%	21.7%
Simulated ACS PUMS	20.8%	29.6%	18.2%	10.0%	21.3%
Share of households					
CPS Reports	5.9%	7.4%	9.2%	8.3%	69.1%
Linked Data	4.9%	8.0%	8.5%	8.1%	70.5%
Simulated ACS PUMS	5.1%	7.9%	8.4%	8.2%	70.3%

Notes: All analyses conducted using household weights (adjusted for PIK probability in the linked data)

Table 3 looks at the impact of the Food Stamp Program along the distribution of house-

hold income relative to the poverty line in more detail. It confirms that reports and administrative data differ substantially, that misreporting is more prevalent in the CPS than in the ACS and that the conditional density method closely reproduces the results from the linked data using public use data.

The administrative data reveal that households above the poverty line benefit from food stamps more than the reports suggest. Table 3 shows that of the roughly one billion dollars found to be missing in the reports, $792 million are underreported by households with income above the poverty line. This means that 40 percent of dollars spent above the poverty line are not reported. $258 million are unreported below the poverty line, making up 13 percent of total dollars spent below the poverty line. Taking the higher underreporting above the poverty line into account, the share of total food stamp payments received by households above the poverty line is 50 percent and not 41 percent as implied by the survey reports.

At the other end of the spectrum, reports also misrepresent the impact of the Food Stamp Program on households in deep poverty. Both surveys suggest that less than half of all households below 50 percent of the poverty line participate in the Food Stamp Program: 40.4 percent in the CPS and 46.2 percent in the ACS. The participation rate in the linked data is 55.5 percent, which is substantially higher than the reports, but still lower than the 62.5 percent recipient households between 50 and 100 percent of the poverty line. While participation increases with income below the poverty line in all data sources, the CPS data mask the decrease in average benefits received as income rises up to the poverty line: average amount received levels off at around $1700 per household below the poverty line according to survey reports. In the administrative data, the average benefit increases with need as we would expect from an anti-poverty program. This difference is mainly due to the high false negative rate in the CPS. Average reported receipt in this income category is $4303, which is slightly higher than in the linked data ($4153). As discussed above, the CPS overstates the share of the SNAP budget spent below the poverty line, but this does not hold in absolute terms - the average amount received by households in deep poverty is

understated by 25 percent in the survey data and the Food Stamp Program really fills 30 percent of the poverty gap (after all cash transfers) as opposed to 22 percent in the CPS.

Misreporting changes over time causing survey data to misrepresent dynamics of the Food Stamp Program even over the short period covered by the linked data. Between 2008 and 2010, the overall number of recipients grew by 33 percent and the total amount received grew by 97 percent according to the linked data. While the ACS understates the growth in recipients slightly, the CPS suggests only a 20 percent increase in recipient households and only a 70 percent increase of total amounts received. According to the CPS, the growth rate of the program was largest in the highest income category, both for the number of recipients and amounts received. The linked data and the simulations show that this overstates growth in this category: The program expanded most (as a percentage increase) between 100 and 150 percent of the poverty line in terms of amounts and between 150 and 200 percent of the poverty line in terms of recipients. In contrast, the survey reports understate growth below the poverty line, suggesting that the Food Stamp Program was less effective than it actually was at countervailing the increase in poverty during the recession: The number of households lifted out of poverty when including food stamp in the income definition has increased 134 percent between 2008 and 2010, while the CPS only indicates a 90 percent increase.

The differences between survey and linked data are driven by the fact that reporting is most accurate among households with income between 50 and 100 percent of the poverty line. Accuracy declines for both higher and lower incomes, so the relationship between income and misreporting seems to be U-shaped[7]. Consequently, scaling up program amounts or participation uniformly does not solve this problem: Participation and amounts above the poverty line would remain understated, while the impact below the poverty line would be overstated. Imputing receipt for those more likely to report receipt would reinforce this tendency, as the probability of receipt usually declines with income in the models used.

---

[7]For total amounts, this relationship is confounded by the fact that the number of households in the lowest category is larger in the CPS than in the ACS (see Meyer and Sullivan, 2003, for potential reasons for this). Average amounts received reveal that underreporting of amounts is U-shaped as well.

Besides making it more complicated to fix, the U-shaped relationship raises interesting questions: Common explanations for underreporting, such as stigma or recall error, are unlikely to lead to this non-monotone pattern. It indicates that misreporting is driven by multiple explanations. Underreporting both of recipiency and amounts is largest among households with annual income higher than 200 percent of the poverty line: Are these households with high income volatility that have had short spells of recipiency they are unlikely to recall? Or are they low-income food stamp units within households that include people with higher income? These questions are beyond the scope of this paper, for further evidence on them see Meyer and Mittag (2013a).

Overall, the findings suggest that if we only looked at survey data, we would underrate the impact of food stamps on households above the poverty line and in deep poverty. It would lead us to believe that a substantially larger share of program resources are received by households below the poverty line. While survey reports understate the level of receipt among households with higher incomes, they overstate the expansion of the program for these households in relative terms. It is likely that misreporting causes bias of similar magnitude at the national level, but whether this bias is similar to that observed in New York State is an open question. Section 8.2 investigates this issue further. From a methodological standpoint, the results presented here confirm that survey reports substantively skew our understanding of anti-poverty programs. Both the extent of misreporting and its U-shaped relation to income suggest that misreporting is likely to cause substantial biases that are hard to assess. Misreporting changes over time, which causes complicated biases in dynamic analyses. Consequently, survey data alone are insufficient to provide accurate information about the effects of government programs. Validation data can solve this problem and yield further insights on the causes of misreporting, but are often difficult to access due to confidentiality. The conditional density method presented in this paper can solve this problem of access because the results based on the public use data and the conditional density are virtually identical to those obtained from the confidential linked data.

8 Extrapolation

Validation data often do not exist for the exact time period or population of interest. If the data of interest (the primary data) have not been validated, the analyst needs to choose between using the survey reports or corrections based on similar validation studies (e.g., Bollinger and David, 1997). In this section, I examine how well extrapolating from existing validation data across time periods and to other states using the conditional density method works. Extrapolation is simple in practice: The implementation is the same as in section 7, but using a sample of (X^S, Z) that was not used in the estimation of the conditional density. This requires the assumption that the conditional density is the same in the validation data and the primary data. In the previous section, the validation data and the primary data were identical, so it held trivially. Whether this assumption is likely to hold or not depends on the case at hand. For food stamps, I examine this assumption directly by extrapolating across time and geography within New York State and find evidence that it works well within state. Extrapolation to the entire U.S. shows that the conditional density method is more accurate than survey reports at capturing geographic variation of the program. It reduces average deviations from total dollars received according to administrative aggregates by a factor between 4 and 8 compared to survey aggregates. However, validation data from other states will be required to evaluate how much error remains. The conditional density method makes it simple to incorporate additional information in the extrapolation and thereby relax the crucial assumption of (conditional) geographic homogeneity.

8.1 Extrapolation Across Time

In order to evaluate extrapolation for the same population over time, I use estimates of the conditional density from a previous survey year, but confine the analysis to New York State. For many of the large annual surveys, validation data have been created for a few years in the past. This raises the question of whether past validation studies can be used to improve

estimates in current surveys or at least assess their sensitivity to misreporting. From a survey design perspective, it is interesting to know how long the conditional distribution remains stable in order to find out how frequently the data should be validated. I use the parameter estimates of the conditional density from 2009 (see appendix table A8) to adjust estimates in the 2010 ACS[8]. In practice, one would obviously prefer to use the current parameters, since validation data are available for both years, but I want to evaluate how well extrapolation would work in the absence of current validation data. The availability of multiple years of validation data allows me to test the crucial assumption that the relationship between X^A and (X^S, Z) remains the same directly. I conduct a test whether the conditional density $f_{X^A|X^R,Z}(X^A|X^R, Z)$ is the same in the linked data from 2009 and 2010 by estimating a fully interacted model on data for both years and testing for parameter equality. The test rejects the assumption, which is not surprising given the sample size of more than 200,000 households. While most parameters are qualitatively similar, there are some differences in the age and income profile. This reflects changes in economic conditions and the program and underlines that the assumption of a stable conditional density is at best an approximation.

A more important question is whether the extrapolation improves upon the results from survey data. Usually, the primary data have not been validated, making it hard to assess whether it actually improves the results. The extrapolation can be assessed by comparing it to information from other sources, such as information on aggregate spending. An advantage of the conditional density method is that it is simple to include such information as restrictions in the extrapolation. This can be used to relax the assumption that the conditional density does not change. For example, total dollars received from the Food Stamp Program are published by the Bureau of Economic Analysis (BEA) and the number of recipient households are available from the U.S. Department of Agriculture (USDA)[9]. One can

---

[8]Future versions of this paper will (also) use the density of 2008 to extrapolate to 2010, but the results have not been disclosed yet.

[9]The USDA uses a different household definition and publishes average monthly participation instead of anual participation, so the numbers are not directly comparable to the linked data. For New York State, the number of participating households is 8-10 percent lower in the linked data and I use this factor to adjust the official aggregates for other states. I pool weighted fiscal years to match the time period of the ACS as

adjust the intercepts for the mass point and amounts of the conditional density to make the expected value of food stamp recipients and amounts match the official numbers. Aggregate total dollars received are also available at lower geographic levels, so one could make total amounts match at the county level by including county specific intercepts.

I use the aggregates from the BEA and USDA in the extrapolation to other states below, but use aggregates calculated from the 2010 linked data in this application to isolate the difference in the results that is due to extrapolation[10]. I repeat the analysis above twice, using the parameter vector from 2009 and an adjusted version that includes an intercept for each of the 143 public use microdata areas (PUMA) in New York State in β and δ. The PUMA-specific intercepts in β adjust the probability of receipt to make the expected number of participating households match participation in the linked data. Those in δ adjust expected amounts received such that total amounts match the amount actually received in each PUMA. I calculate them by first adjusting each intercept in β iteratively until the expected number of recipient households match. The intercepts in δ are calculated the same way, but take the adjustment of β into account.

Columns (5) and (6)[11] of table 4 apply the conditional density method to the 2010 data using the parameters from 2009. Column (3) contains the ideal case of having current parameter estimates and column (4) the results that would have been obtained from the CPS in the absence of validation data. Columns (1) and (2) use the linked data for 2009 and 2010 to place the results in context. The results from the extrapolation in coulmn (5) all fall between the the numbers in columns (1) and (2), indicating that the extrapolation captures some of the expansion of the program, but not all of it. However, they are much closer to the true results in column (2) than the survey reports, so if validation data are available for recent years, extrapolation may be preferable to reports. The lower part of the table shows mean absolute deviations of PUMA aggregates based on the conditional density to the 2010

in Meyer, Mok and Sullivan (2009).

[10]For the same reason, I use the internal ACS file instead of the public use data here.

[11]Column (6) of table 4 is based on a preliminary ajustment that adjusted the probability of receipt to make total dollars spent match. They will be replaced as soon as I have access to the data again.

	(1) Linked Data 2009	(2) Linked Data 2010	(3) CPS Reports 2010	(4) Using Cond. Density from 2010	(5) 2009	(6) 2009, adjusted
Total amount received (in million \$)	2,800	3,959	2,909	3,954	3,088	3,956
Average amount received per HH	390	550	388	549	429	550
# of HH receiving SNAP (×1000)	1,114	1,280	871	1,268	1,198	1,731
...as percentage of population	15.5%	17.8%	11.6%	17.6%	16.7%	24.1%
Poverty rate among recipients	46.5%	43.3%	53.6%	43.4%	45.1%	37.1%
...including FS benefits	34.9%	29.3%	37.7%	30.8%	34.7%	28.8%
Mean absolute deviations to PUMA level...						
# of HH receiving SNAP (×1000)	-	-	-	0.95	1.00	3.27
total amount received (in million \$)	-	-	-	3.57	6.70	0.30

Notes: Column 1 contains the results for 2009 for comparison. Columns 2, 3 and 4 replicate columns 2,3 and 4 of table 2. Column 5 used the same method as column 4, but parameter estimates from the 2009 linked data. Column 6 is preliminary and will be replaced. It adds PUMA-specific intercepts to this parameter vector that make expected total dollars spent match total dollars spent in the linked data in each PUMA. Mean absolute deviations are the average absolute value of the difference between PUMA totals in the linked data and the method in the respective column. All analyses conducted using household weights (adjusted for PIK probability in the ACS)

linked data. They are small relative to the level of these aggregates. The extrapolation does better at predicting participation than at reproducing total amounts received for small geographic areas. Column (6) uses a preliminary version of the adjustment described above: I adjusted the probability of receipt to match total dollars spent by PUMA.[12] This leads to an overestimate of the number of participating households, showing that it is important to adjust both interecepts. However, it improves analyses of amounts received, suggesting that adjusting both margins may significantly improve extrapolation.

In conclusion, the results show that extrapolation to other years is not trivial because the conditional density changes over time. It is likely to be more stable in less turbulent times, but probably changes more over longer time periods. However, a rich set of conditioning variables and incorporating additional information can provide quite accurate results. While some bias in the results is likely to remain, extrapolation is more accurate than survey reports in this application.

---

[12]I currently do not have access to the restricted ACS from Chicago, but will replace this column by an improved version that also adjusts the number of households as soon as possible.

8.2 Extrapolation to the entire U.S.

Most validation studies use a non-random subsample from a specific geographic area because of data availability. Administrative records are often kept by state agencies, so validating a sample from the entire U.S. would require obtaining linkable data from all 51 such agencies, which is infeasible. Therefore, it is important to assess how much can be learned about the entire population from a geographically confined subsample. I provide evidence that most of the geographic heterogeneity of the Food Stamp Program within New York State is captured by observable covariates. The assumption required to obtain consistent estimates for the entire U.S. using the conditional density method is that the conditional density is identical across states. One may expect differences for food stamps, e.g., because the program is administered by the states, so there may be differences in the way the program is run. However, SNAP is a federal program and its rules are uniform across states with few exceptions. The assumption that the conditional density is the same in all states is more likely to be an approximation, but there are reasons to believe that it is sufficiently similar to work well in practice. I find that extrapolation reproduces administrative aggregates well and substantially improves over the alternative, which usually is to use survey reports that are subject to severe misreporting.

In contrast to the previous section, the underlying assumption is not directly testable. It could be tested if validation data from another state were available for the same survey, which is likely to happen in the future. In lieu of a formal test, evidence on the extent of geographic heterogeneity can be obtained from within New York State. There is no variation in state institutions in these tests, but New York State is a large state with much demographic heterogeneity and variation in the administration of the program. I first examine how much of the between-PUMA variation in receipt and amounts is explained by the variables included in the conditional density. To do so, I estimate a probit model for receipt and a linear regression for amounts conditional on receipt at the household level using the administrative variables in the linked data and the covariates used in the conditional density. I use the

parameter estimates to calculate average predicted receipt and amounts for each PUMA and regress the actual PUMA averages on the predicted averages. The R^2 from these regressions are measures of how much of the difference between geographic areas is explained by the covariates I use and are 0.98 for receipt and 0.62 for amounts. These results are encouraging, indicating that a substantial share of the geographic variation is captured by the covariates.

Table 5: Extrapolation Across Geography Within New York State, 2010

	(1) Linked Data	(2) unadjusted Conditional Density	(3) adjusted Conditional Density
Total amount received (in million $)	1,853	1,990	1,849
Average amount received per HH	529	568	528
# of HH receiving SNAP ($\times$1000)	607	628	571
. . . as percentage of population	17.3%	17.9%	16.3%
Poverty rate among recipients	41.2%	41.2%	43.7%
. . . including FS benefits	28.3%	29.4%	31.3%
Mean absolute deviations to PUMA level. . .			
number of HH receiving SNAP ($\times$1000)	-	0.82	0.82
total amount received (in million $)	-	3.29	0.29

Notes: Column 1 uses the actual records from the linked data for the western PUMAs. Column 2 uses the parameter estimates of the conditional density from the eastern PUMAs and the internal ACS data from the western PUMAs. Column 3 is preliminary and will be replaced. It adds PUMA-specific intercepts to this density that make expected total dollars spent match total dollars spent in the linked data in each PUMA. Mean absolute deviations are the average absolute value of the difference between PUMA totals in the linked data and the method in the respective column. All analyses conducted using household weights adjusted for PIK probability

To investigate this issue in more detail, I split the linked data into two samples based on geography. To avoid extrapolating from New York City to Upstate New York or vice versa, I create a subsample that includes the western PUMAs of New York City and the western PUMAs of Upstate New York. Summary statistics are provided in table A2. Differences between regions may be more pronounced when extrapolating to other states, but this setup is reasonably similar because of the size and diversity of the population of New York State. Parameter estimates of the conditional density for both subsamples are in appendix table A9. The coefficients on the main predictors are similar, but there are some significant differences for other variables. Consequently, a test of parameter equality rejects that the densities are

the same.

However, the results in table 5 indicate that the extrapolation works well. Column (1) was calculated from the linked data for the western subsample, while columns (2) and (3) use the conditional density from the eastern PUMAs to predict food stamp receipt and amounts in the western subsample. Column (3) uses the same preliminary adjustment of the parameter vector as the previous section.[13] Both methods yield results that are similar to the true results obtained from the linked data. As before, the preliminary adjustment improves amounts, but fares slightly worse on the number of recipients. This indicates that extrapolation can be improved by incorporating additional information, but that a better adjustment is desirable. However, adjusting the parameters does not change the results much, underlining that most of the geographic variation is captured by the covariates.

Table 6: Extrapolation across Geography: to entire U.S., 2010

	(1) ACS Reports	(2) CPS Reports	(3) Conditional Density New York	(4) adjusted Conditional Density
Total amount received (in million $)	-	35,282	57,646	60,689
Average amount received per HH	-	299	503	530
# of HH receiving SNAP (×1000)	13,586	11,827	17,910	17,544
. . . as percentage of population	11.9%	10.0%	15.6%	15.3%
Poverty rate among recipients	49.4%	52.7%	43.7%	43.9%
. . . including FS benefits	-	41.7%	32.3%	31.6%
Mean absolute deviations of total $ received (in million $). . .				
to admin. state totals	-	497.2	102.0	5.1
to admin. totals for large MSAs	-	210.0	54.5	22.3
to admin. totals for county groups	-	-	10.2	8.7

Notes: Column 3 uses the 2010 U.S. ACS PUMS data excluding group quarters and the conditional density from New York State, column 4 adds state specific intercepts to this density that make the number of recipient households and total $ spent in each state match administrative totals in expectation. Administrative totals on total $ spent are from the BEA regional economic information system, the # of households is from the USDA and adjusted for the different household definition and time period. Mean absolute deviations are the average absolute value of the difference between administrative totals from the BEA and the method in the respective column. County groups are the smallest geographic area that does not split counties or PUMAs. All analyses conducted using household weights.

Whether this is also true when extrapolating to other states is a difficult question, but

---

[13]I currently do not have access to the restricted ACS from Chicago, but will replace this column by an improved version that also adjusts the number of households as soon as possible.

the results from extrapolation to the entire U.S. in table 6 are promising. Column (3) uses the unadjusted parameter estimates from New York State to predict food stamp receipt and amounts in the entire U.S. sample of the 2010 ACS public use data. The results in column (4) adjust the parameters such that the expected number of recipient households matches USDA state aggregates and total dollars spent by state match the state totals published by the BEA as described in section 8.1. The BEA also publishes total dollars spent at the county level, so researchers interested in small areas could match expectations at the county group[14] level to obtain better results. Making the adjustment at the state level allows me to use the sub-state totals to assess the accuracy of the extrapolation. The CPS is not representative at the sub-state level except for large metropolitan statistical areas (MSAs), so in order to provide a fair comparison for the adjusted results, I also compare them to total dollars spent according to the BEA for MSAs with more than half a million inhabitants.[15] The lower part of table 6 computes mean absolute deviations of total dollars spent at the state, MSA and county group level. The improvements over survey reports are drastic, even with the unadjusted parameters: Mean absolute deviation in the survey data is almost 5 times as large at the state level and nearly 4 times as large at the MSA level. After adjustment, MSA totals are off by less than 11 percent of the average deviation according to survey reports. The difference at the state level due to the fact that these numbers were forced to match in expectation. Nonetheless, they suggest that inference about sub-state areas may greatly benefit from adjustments to county group totals. The number of participating households is only available at the state level. The extrapolation based on the unadjusted density deviates from these numbers by only 33,291 households on average, which is a third of the average deviation in the CPS and less than 10 percent of actual state-level participation. While these comparisons only evaluate one aspect of the extrapolation, they certainly favor the extrapolation over survey reports.

---

[14]I define county groups as the smallest combination of counties that does not cut PUMA boundaries.

[15]I exclude MSAs with more than 1 percent of the population in rural PUMAs to ensure comparability to the ACS.

The upper part of table 6 shows that there are substantive differences between the results based on reports and the conditional density method. As in table 2, the ACS is closer to the extrapolation than the CPS, but both suggest a much smaller program and a larger share of food stamp recipients in poverty than the extrapolated administrative data. The results using the adjusted and unadjusted conditional density are similar, suggesting again that the covariates capture most of the heterogeneity. Besides the evidence from aggregates above, it is hard to assess how close these numbers are to "truth." The extrapolation is certainly better aligned with what we know about food stamps from other sources, for example that 18.4 million households received food stamps in 2010 (U.S. Department of Agriculture, 2011). Neither the time period nor the household definition is directly comparable to the numbers in table 6, but they suggest that the conditional density method depicts the program better than the survey reports.

Overall, the analyses in this part provide evidence that extrapolations based on the conditional density method perform better than survey reports of food stamp participation and amounts, but several unsettled issues remain. Validation data from a second state would be useful to examine how closely extrapolation reproduces true recipiency and amounts and could also be used to examine which adjustments work well. Whether extrapolation works for other variables is an open question, but the findings here suggest that approaches like the conditional density method have the potential to mitigate the problem that linked data often cover only a limited time period or geographic area.

9   Conclusion

The linked administrative data from New York State show that survey data misrepresent the impact of the Food Stamp Program on poverty. Most importantly, they conceal a large part of receipt above the poverty line and a substantial share of participation by households in deep poverty. They also distort dynamic aspects of the program, such as the expansion in

the recent recession. This distortion is not uniform across the distribution of recipients, but varies by income relative to the poverty line. Even though the results are only for one state, the analysis in section 8.2 points towards similar differences in the entire country. SNAP is the largest anti-poverty program and expanding rapidly, so an accurate understanding of its impact and beneficiaries is essential. Survey data skew our beliefs about which segments of the population receive a substantial amount of the resources spent to reduce poverty and how this spending has changed in recent years. Whether increased spending above the poverty line or among households in deep poverty is desirable is a normative question and beyond the scope of this paper. However, to answer such questions, we certainly need a more accurate description of the program than that provided by survey reports. The differences between survey reports and administrative data also raise many questions about how the program should be run. Efforts to increase participation by eligibles may be misguided if the populations we reach out to already receive food stamps but do not report it. Or we may want to re-consider eligibility criteria if a lot of money turns out to go towards households that we do not want the program to support. Finally, the administrative data show that food stamps serve a more heterogeneous population than we thought, since underreporting is more prevalent at both extremes of the income range. This raises the question of whether a uniform program is the optimal way to serve the entire current target population. For example, households in higher income categories are likely to spend more money on food than the food stamp benefit they are issued and are thus not constrained to increase their food consumption. Providing a cash transfer to these households may be cheaper and reduce stigma (Whitmore, 2002). Average spell lengths among these households suggest that longer re-certification periods for these households would have little effect (unless it increased fraud). Yet, longer periods may be desirable for households in deep poverty in order to increase take-up. In conclusion, accurate information about government programs is necessary to target them well and run them efficiently, but the inaccuracy of survey data conceal important details.

From a methodological perspective, this study confirms that the extent of misreporting is large and that it is systematically related to important variables. The patterns of misreporting in this study raise interesting questions about the reasons for misreporting and how it can be avoided. The U-shaped relationship to income suggests that it is driven by multiple causes. The extent of underreporting among households with annual income above two times the poverty line is striking. Whether this is driven by income volatility or differences between the food stamp unit and the ACS household is an interesting topic for future research. Misreporting creates a form of non-classical measurement error that affects the conclusions we draw about food stamps and that cannot be corrected by common methods for measurement error. Therefore, research that examines inference from contaminated data is important to improve our understanding of the Food Stamp Program. Similar problems are likely to be present in studies that use other variables known to suffer from non-classical measurement error, such as income or education. Consequently, improving corrections for measurement error is desirable. Validation studies can solve the problem and could be conducted for other variables in a similar way, for example by linking tax records to survey data. However, the usefulness of validation data is reduced by the difficulty of access to the data and the limited time periods and areas for which they are available. The conditional density method in this paper can solve the problem of access by offering a simple way to obtain estimates from public use data in common applications. The results from extrapolation to other states and across time have shown its potential to mitigate or solve the problem that validation data are often not available for the population of interest. Thereby, it makes existing validation studies more useful and additional validation data more valuable. As more validation data become available, we will be in a better position to evaluate and improve the performance of such methods. In the application to food stamp receipt and amounts received, my simulations based on the conditional density clearly improve upon survey reports even in the extrapolations. This underlines the importance of additional validation studies and methods to reduce or correct bias arising from measurement error in survey data.

Appendix

Table A1: The Determinants of a Household having a PIK, Probit Coefficients

	(1)	(2)	(3)	(4)	(5)	(6)
	2008		2009		2010	
	Coef.	*SE*	*Coef.*	*SE*	*Coef.*	*SE*
Not Married, no children	-0.047	0.037	-0.137	0.035	-0.107	0.036
Not Married, children	0.203	0.036	0.193	0.036	0.040	0.036
Married, no children	-0.037	0.039	-0.048	0.037	-0.080	0.039
Number of Persons in HH	0.085	0.011	0.051	0.011	0.071	0.011
Number of Persons under 18	-0.034	0.020	0.017	0.019	-0.020	0.020
Age 18-29	-0.010	0.030	-0.129	0.030	-0.024	0.032
Age 30-39	-0.024	0.026	-0.097	0.026	-0.096	0.028
Age 50-59	0.169	0.026	0.145	0.026	0.089	0.027
Age 60-69	0.230	0.030	0.274	0.031	0.159	0.031
Age $\geq$ 70	0.278	0.032	0.362	0.033	0.221	0.034
Less than High School	-0.060	0.028	-0.090	0.028	-0.083	0.030
High School Degree	-0.129	0.021	-0.165	0.022	-0.125	0.022
College or more	0.062	0.022	0.087	0.023	0.061	0.023
Hispanic	0.023	0.024	0.039	0.025	0.060	0.026
White	0.087	0.019	0.103	0.019	0.072	0.020
HH head unemployed	0.011	0.049	0.030	0.039	0.083	0.040
HH head not in labor force	-0.109	0.024	-0.105	0.024	-0.075	0.025
Income/poverty line	0.015	0.002	0.015	0.002	0.011	0.002
Disability	-0.003	0.042	0.070	0.046	0.074	0.048
Disabled, not working	0.117	0.049	0.062	0.053	0.057	0.055
Speaks poor English	-0.346	0.033	-0.274	0.034	-0.180	0.036
Speaks no English	-0.473	0.051	-0.388	0.053	-0.333	0.056
Not a US citizen	-0.484	0.025	-0.512	0.025	-0.324	0.027
Public Assistance (reported)	0.076	0.050	0.079	0.049	0.137	0.053
Constant	1.209	0.049	1.293	0.047	1.369	0.050
Observations	107,349		107,250		106,656	

Notes: All analyses conducted using household weights. Demographic characteristics refer to the household head. The omitted family type is married with children, the omitted age category is 40-49 and the omitted education category is some college.

Table A2: Summary Statistics Linked Data

| Variable | Linked Data Piked New York State Sample | | | | | | NYS Subsample | |
| | 2008 | | 2009 | | 2010 | | 2010 | |
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
HH head had a PIK	0.975	0.157	0.973	0.161	0.985	0.121	0.985	0.121
# of members piked	2.355	1.399	2.333	1.384	2.436	1.456	2.551	1.484
SNAP receipt reported	0.109	0.311	0.128	0.334	0.141	0.348	0.139	0.346
Admin. SNAP receipt	0.135	0.342	0.155	0.362	0.178	0.382	0.173	0.378
Admin. SNAP amount	281.4	941.8	389.7	1182.9	550.2	1511.2	529.4	1465.8
NY subsample	0.490	0.500	0.487	0.500	0.487	0.500		
Below poverty line	0.125	0.331	0.124	0.329	0.129	0.335	0.121	0.326
...incl. SNAP	0.115	0.318	0.106	0.307	0.104	0.305	0.098	0.298
Out of poverty by SNAP	0.011	0.103	0.018	0.133	0.025	0.156	0.022	0.148
HH income	80,209	109,104	80,010	106,088	78,142	102,134	75,280	83,395
Income/Poverty Line	5.090	6.975	5.078	6.780	4.843	6.293	4.527	4.978
SNAP receipt imputed	0.009	0.095	0.009	0.093	0.013	0.113	0.014	0.116
Age 18-29	0.095	0.293	0.095	0.293	0.102	0.302	0.087	0.281
Age 30-39	0.175	0.380	0.170	0.376	0.168	0.374	0.163	0.369
Age 50-59	0.206	0.404	0.209	0.407	0.210	0.407	0.217	0.412
Age 60-69	0.144	0.351	0.150	0.357	0.153	0.360	0.154	0.361
Age 70 or older	0.156	0.363	0.157	0.364	0.158	0.364	0.160	0.366
Any income from capital	0.264	0.441	0.246	0.431	0.222	0.416	0.215	0.410
# of persons in HH	2.503	1.491	2.512	1.507	2.544	1.538	2.662	1.563
# of children in HH	0.601	1.039	0.597	1.036	0.596	1.037	0.642	1.047
Not married, no children	0.450	0.498	0.455	0.498	0.455	0.498	0.417	0.493
Not married, children	0.099	0.299	0.098	0.297	0.099	0.298	0.105	0.306
Married, no children	0.255	0.436	0.254	0.435	0.255	0.436	0.269	0.444
Linguistic Isolation	0.080	0.271	0.078	0.268	0.078	0.269	0.080	0.271
# of persons employed	1.753	1.485	1.715	1.479	1.710	1.486	1.792	1.523
Anyone in HH employed	0.781	0.414	0.770	0.421	0.766	0.423	0.772	0.420
Elderly or disabled in HH	0.438	0.496	0.443	0.497	0.442	0.497	0.459	0.498
Single household	0.295	0.456	0.294	0.456	0.288	0.453	0.258	0.438
HH head unemployed	0.029	0.169	0.045	0.207	0.051	0.219	0.052	0.223
HH head not in labor force	0.312	0.463	0.317	0.465	0.322	0.467	0.323	0.468
HH head female	0.498	0.500	0.501	0.500	0.505	0.500	0.496	0.500
HH head not a U.S. citizen	0.095	0.293	0.092	0.289	0.096	0.294	0.099	0.298
White (HH head)	0.722	0.448	0.726	0.446	0.713	0.452	0.706	0.455
Less than High School	0.140	0.347	0.136	0.342	0.133	0.340	0.138	0.344
High School Degree	0.257	0.437	0.256	0.437	0.254	0.435	0.270	0.444
College or more	0.344	0.475	0.347	0.476	0.349	0.477	0.313	0.464
HH head disabled	0.156	0.363	0.154	0.361	0.148	0.355	0.151	0.358
Disabled, not working	0.119	0.324	0.120	0.325	0.118	0.323	0.120	0.325
HH head speaks no English	0.016	0.125	0.015	0.122	0.015	0.123	0.015	0.121
Speaks English poorly	0.051	0.220	0.050	0.219	0.051	0.220	0.054	0.226
Number of Observations	101,378		101,346		101,680		49,576	
	SNAP Recipients Only							
Admin. SNAP amount	2,082.8	1,677.6	2,514.3	1,920.2	3,093.9	2,229.5	3,055.1	2,164.2
Below poverty line	0.498	0.500	0.465	0.499	0.433	0.495	0.412	0.492
. . . incl. SNAP	0.418	0.493	0.349	0.477	0.293	0.455	0.283	0.451
Out of poverty by SNAP	0.080	0.271	0.116	0.320	0.140	0.347	0.129	0.335
Number of Observations	11,287		13,117		15,464		7,272	

Notes: All statistics at the household level using household weights adjusted for the probability of having a PIK.

Table A3: Summary Statistics ACS Public Use Data New York State

Variable	2008 Mean	2008 SD	2009 Mean	2009 SD	2010 Mean	2010 SD
SNAP receipt reported	0.107	0.310	0.124	0.329	0.138	0.345
NY subsample	0.487	0.500	0.487	0.500	0.489	0.500
Below poverty line	0.125	0.330	0.123	0.328	0.130	0.336
HH income	80,176	98,852	80,335	97,636	78,306	92,777
Income/Poverty Line	5.086	6.304	5.097	6.198	4.858	5.703
SNAP receipt imputed	0.011	0.105	0.011	0.103	0.016	0.125
Age 18-29	0.096	0.294	0.095	0.293	0.102	0.302
Age 30-39	0.176	0.381	0.171	0.376	0.169	0.375
Age 50-59	0.205	0.404	0.208	0.406	0.210	0.407
Age 60-69	0.143	0.351	0.150	0.357	0.152	0.359
Age 70 or older	0.157	0.364	0.157	0.364	0.157	0.364
Any income from capital	0.263	0.440	0.246	0.431	0.221	0.415
# of persons in HH	2.500	1.489	2.510	1.503	2.539	1.538
# of children in HH	0.597	1.034	0.596	1.033	0.595	1.037
Not married, no children	0.452	0.498	0.455	0.498	0.455	0.498
Not married, children	0.099	0.299	0.098	0.297	0.099	0.298
Married, no children	0.254	0.435	0.253	0.435	0.255	0.436
Linguistic Isolation	0.085	0.279	0.082	0.274	0.083	0.275
# of persons employed	1.751	1.480	1.713	1.476	1.704	1.485
Anyone in HH employed	0.780	0.414	0.770	0.421	0.765	0.424
Elderly or disabled in HH	0.435	0.496	0.441	0.496	0.441	0.497
Single household	0.296	0.457	0.297	0.457	0.291	0.454
HH head unemployed	0.028	0.165	0.046	0.208	0.051	0.219
HH head not in labor force	0.312	0.463	0.316	0.465	0.323	0.468
HH head female	0.497	0.500	0.501	0.500	0.505	0.500
HH head not a US citizen	0.095	0.293	0.092	0.290	0.096	0.295
White (HH head)	0.719	0.450	0.720	0.449	0.707	0.455
Less than High School	0.140	0.347	0.136	0.343	0.135	0.342
High School Degree	0.257	0.437	0.256	0.436	0.252	0.434
College or more	0.343	0.475	0.348	0.476	0.348	0.476
HH head disabled	0.155	0.362	0.154	0.361	0.149	0.356
Disabled, not working	0.119	0.324	0.120	0.325	0.120	0.325
HH head speaks no English	0.016	0.124	0.015	0.122	0.015	0.123
Speaks English poorly	0.051	0.220	0.050	0.217	0.052	0.222
Number of Observations	73,106		73,685		74,105	

Notes: All statistics at the household level using household weights.

Table A4: Summary Statistics ACS Public Use Data

	NY Subsample 2010		U.S. 2010	
Variable	*Mean*	*SD*	*Mean*	*SD*
SNAP receipt reported	0.135	0.342	0.119	0.323
Below poverty line	0.121	0.326	0.128	0.334
HH income	75,968	79,911	67,801	70,314
Income/Poverty Line	4.572	4.688	4.251	4.380
SNAP receipt imputed	0.016	0.127	0.013	0.114
Age 18-29	0.088	0.283	0.115	0.319
Age 30-39	0.164	0.371	0.172	0.377
Age 50-59	0.217	0.412	0.205	0.404
Age 60-69	0.153	0.360	0.152	0.359
Age 70 or older	0.159	0.366	0.152	0.359
Any income from capital	0.213	0.409	0.218	0.413
# of persons in HH	2.655	1.561	2.520	1.483
# of children in HH	0.640	1.048	0.624	1.060
Not married, no children	0.419	0.493	0.417	0.493
Not married, children	0.106	0.308	0.097	0.296
Married, no children	0.268	0.443	0.286	0.452
Linguistic Isolation	0.083	0.276	0.046	0.210
# of persons employed	1.787	1.516	1.714	1.482
Anyone in HH employed	0.773	0.419	0.767	0.423
Elderly or disabled in HH	0.457	0.498	0.444	0.497
Single household	0.260	0.439	0.274	0.446
HH head unemployed	0.052	0.222	0.055	0.228
HH head not in labor force	0.322	0.467	0.314	0.464
HH head female	0.494	0.500	0.470	0.499
HH head not a US citizen	0.100	0.300	0.065	0.246
White (HH head)	0.700	0.458	0.782	0.413
Less than High School	0.140	0.347	0.125	0.331
High School Degree	0.265	0.441	0.262	0.440
College or more	0.314	0.464	0.303	0.460
HH head disabled	0.152	0.359	0.164	0.370
Disabled, not working	0.122	0.327	0.126	0.332
HH head speaks no English	0.014	0.119	0.010	0.101
Speaks English poorly	0.055	0.229	0.030	0.171
Number of Observations	36,192		1,203,777	

Notes: All statistics at the household level using household weights.

Table A5: Summary Statistics CPS New York State, 2008-2009

Variable	2008 Mean	SD	2009 Mean	SD
SNAP receipt reported	0.095	0.294	0.106	0.308
Reported SNAP amount	224.3	904.5	333.5	1295.8
Household income	70,211	75,984	70,818	76,675
Income/Poverty Line	4.405	4.662	4.439	4.692
Below poverty line	0.143	0.350	0.135	0.341
HH income incl. SNAP below poverty line	0.133	0.340	0.120	0.325
Lifted above poverty line by SNAP	0.010	0.098	0.015	0.121
Number of Observations	6,687		6,735	
SNAP Recipients Only				
Reported SNAP amount	2,351.3	1,891.9	3,136.6	2,646.7
Below poverty line	0.622	0.485	0.565	0.496
HH income incl. SNAP below poverty line	0.520	0.500	0.426	0.495
Lifted above poverty line by SNAP	0.102	0.302	0.139	0.346
Number of Observations	696		798	

Notes: Years refer to the ACS survey years, i.e. 2009 pools the 2009 and 2010 CPS ASEC. All statistics at the household level using household weights adjusted for pooling years.

Table A6: Summary Statistics CPS 2010

Variable	New York State Mean	SD	U.S. Mean	SD
SNAP receipt reported	0.116	0.320	0.100	0.300
Reported SNAP amount	388.1	1470.6	298.7	1225.5
Household income	70,764	81,664	67,741	73,077
Income/Poverty Line	4.351	4.763	4.146	4.403
Below poverty line	0.141	0.348	0.128	0.334
HH income incl. SNAP below poverty line	0.122	0.327	0.117	0.322
Lifted above poverty line by SNAP	0.019	0.136	0.011	0.103
Number of Observations	6,689		151,368	
SNAP Recipients Only				
Reported SNAP amount	3,340.7	2,959.5	2,983.2	2,644.0
Below poverty line	0.560	0.497	0.542	0.498
HH income incl. SNAP below poverty line	0.396	0.489	0.436	0.496
Lifted above poverty line by SNAP	0.163	0.370	0.106	0.308
Number of Observations	869		16,311	

Notes: Years refer to the ACS survey years, i.e. 2010 pools the 2010 and 2011 CPS ASEC. All statistics at the household level using household weights adjusted for pooling years.

Table A7: Parameter Estimates of the Conditional Density, New York State 2010

	Mass Point		Amounts	
	Coef.	SE	Coef.	SE
SNAP receipt reported	-2.50	0.002	1,453.0	6.0
SNAP receipt imputed	-0.42	0.005	562.7	15.5
Income relative to Poverty Line Intercepts				
50-100 %	0.62	0.015	376.8	25.8
100-150 %	-0.26	0.020	-244.9	45.7
150-200 %	-0.41	0.029	113.6	82.5
>200 %	0.48	0.005	-1,139.3	13.2
Income relative to Poverty Line Slopes				
≤ 50 %	-0.07	0.017	461.8	31.0
50-100 %	-0.90	0.018	-733.0	31.2
100-150 %	0.35	0.016	-410.3	36.2
150-200 %	0.39	0.017	-702.5	47.3
>200 %	0.04	0.000	-36.1	1.9
Age 18-29	0.00	0.003	53.5	8.0
Age 30-39	-0.13	0.003	23.5	7.0
Age 50-59	0.04	0.003	-57.3	7.6
Age 60-69	0.16	0.004	-17.7	9.7
Age 70 or older	0.33	0.004	-116.0	10.3
Any income from capital	0.35	0.003	-597.2	11.9
# of persons in HH	-0.21	0.001	477.3	2.9
# of children in HH	-0.02	0.002	862.1	4.1
Not married, no children	-0.54	0.004	253.7	9.1
Not married, children	-0.54	0.003	65.2	7.4
Married, no children	-0.21	0.004	221.7	10.8
Linguistic Isolation	-0.05	0.003	22.1	7.8
# of persons employed	0.08	0.001	-295.7	3.7
Anyone in HH employed	0.02	0.003	-51.7	8.9
Elderly or disabled in HH	-0.14	0.003	-187.2	6.6
Single household	0.23	0.003	-1,387.7	9.6
Unemployed	-0.12	0.004	-358.5	8.7
Not in labor force	-0.10	0.003	-63.6	7.2
Female	-0.06	0.002	141.1	5.3
Not a US citizen	0.27	0.003	-453.6	7.1
White	0.41	0.002	-208.7	4.7
Less than High School	-0.11	0.003	140.3	6.4
High School Degree	-0.09	0.002	68.6	6.1
College or more	0.20	0.003	66.7	8.7
Disabled	-0.07	0.005	173.3	13.8
Disabled, not working	-0.07	0.006	-138.1	14.6
Speaks English poorly	-0.16	0.004	289.2	8.0
Speaks no English	0.10	0.006	266.0	12.2
Constant	1.60	0.007	486.1	17.7
Sigma	2,034	2.0		
Left Truncation Point	16	0.003		
Number of observations	101,680			

Notes: Parameter estimates of truncated normal conditional distribution. The columns labeled "Mass Point" determine the probability that the household does *not* receive SNAP and can be interpreted like probit coefficients. The columns labeled "Amounts" determine amounts conditional on receipt. Demographic characteristics refer to the household head. All analyses use household weights adjusted for PIK-probability.

Table A8: Parameter Estimates of the Conditional Density, New York State 2008-2009

	2008				2009			
	Mass Point		*Amounts*		*Mass Point*		*Amounts*	
	Coef.	*SE*	*Coef.*	*SE*	*Coef.*	*SE*	*Coef.*	*SE*
SNAP receipt reported	-2.61	0.003	1,487.4	6.5	-2.50	0.002	1,429.1	6.7
SNAP receipt imputed	-0.54	0.006	514.1	18.4	-0.50	0.006	535.8	19.6
Income relative to Poverty Line Intercepts								
50-100 %	-0.26	0.015	522.6	25.0	-0.12	0.015	671.9	26.7
100-150 %	-0.52	0.022	-1,127.1	51.1	-0.22	0.021	-362.5	49.0
150-200 %	-0.26	0.035	-564.8	89.4	-0.39	0.032	1,386.6	95.8
>200 %	0.32	0.006	-1,188.7	14.4	0.53	0.006	-1,451.2	14.3
Income relative to Poverty Line Slopes								
≤ 50 %	-0.47	0.019	-72.5	32.0	-0.48	0.018	255.2	31.9
50-100 %	0.11	0.018	-1,245.2	30.6	0.01	0.018	-1,295.7	32.4
100-150 %	0.48	0.017	89.2	41.0	0.25	0.016	-483.3	39.1
150-200 %	0.31	0.020	-223.8	51.5	0.41	0.018	-1,541.7	55.0
>200 %	0.05	0.001	-10.2	2.2	0.02	0.000	-18.4	2.0
Age 18-29	-0.07	0.004	57.0	8.1	-0.06	0.003	168.8	8.4
Age 30-39	-0.08	0.003	78.4	7.1	-0.05	0.003	191.8	7.5
Age 50-59	0.04	0.003	-80.5	8.2	0.01	0.003	32.1	8.1
Age 60-69	0.22	0.004	64.3	10.3	0.10	0.004	21.4	10.3
Age 70 or older	0.42	0.004	0.1	11.0	0.35	0.004	-96.8	10.9
Any income from capital	0.28	0.003	-135.9	11.8	0.33	0.003	-211.9	12.4
# of persons in HH	-0.17	0.002	278.7	3.4	-0.20	0.001	416.9	3.3
# of children in HH	-0.07	0.002	604.3	4.3	-0.01	0.002	691.5	4.5
Not married, no children	-0.56	0.004	-162.2	9.7	-0.43	0.004	30.9	10.0
Not married, children	-0.58	0.004	-77.3	7.8	-0.50	0.003	-96.8	8.1
Married, no children	-0.13	0.004	-45.7	11.9	-0.04	0.004	-88.8	11.9
Linguistic Isolation	-0.04	0.004	-49.0	8.0	-0.03	0.003	33.4	8.2
# of persons employed	0.09	0.002	-73.4	4.0	0.10	0.002	-255.9	4.2
Anyone in HH employed	0.02	0.004	-25.3	9.8	0.04	0.004	-1.0	9.5
Elderly or disabled in HH	-0.26	0.003	-271.7	6.8	-0.20	0.003	-87.6	7.0
Single household	0.27	0.004	-1,307.8	10.5	0.19	0.003	-1,306.9	10.1
Unemployed	-0.17	0.005	-17.0	10.0	-0.06	0.004	-277.4	9.7
Not in labor force	-0.04	0.004	-123.1	7.6	-0.04	0.003	-81.3	7.7
Female	-0.05	0.002	16.5	5.7	-0.06	0.002	40.0	5.7
Not a US citizen	0.21	0.003	-318.3	7.5	0.12	0.003	-290.9	7.6
White	0.39	0.002	-87.8	4.9	0.42	0.002	-168.4	5.0
Less than High School	-0.20	0.003	34.8	6.5	-0.17	0.003	102.8	6.7
High School Degree	-0.06	0.003	30.5	6.5	-0.12	0.002	55.2	6.6
College or more	0.17	0.003	245.3	9.2	0.18	0.003	30.5	9.5
Disabled	-0.07	0.005	228.1	12.6	0.05	0.005	26.5	14.0
Disabled, not working	-0.13	0.006	-62.8	13.5	-0.18	0.006	-60.1	14.8
Speaks English poorly	-0.17	0.004	217.0	8.4	-0.06	0.004	296.7	8.5
Speaks no English	-0.12	0.006	447.6	12.0	-0.12	0.006	439.4	12.1
Constant	1.84	0.008	11.9	19.1	1.60	0.008	261.9	19.4
Sigma	1,680	2.1			1,903	2.2		
Left Truncation Point	4	0.002			10	0.002		
Number of observations	101,378				101,346			

Notes: Parameter estimates of truncated normal conditional distribution. The columns labeled "Mass Point" determine the probability that the household does *not* receive SNAP and can be interpreted like probit coefficients. The columns labeled "Amounts" determine amounts conditional on receipt. Demographic characteristics refer to the household head. All analyses use household weights adjusted for PIK-probability.

Table A9: Parameter Estimates of the Conditional Density, New York State Subsamples 2010

| | Eastern Counties | | | | Western Counties | | | |
| | Mass Point | | Amounts | | Mass Point | | Amounts | |
	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
SNAP receipt reported	-2.49	0.003	1,523.7	8.2	-2.53	0.003	1,372.1	8.6
SNAP receipt imputed	-0.31	0.007	568.1	21.6	-0.54	0.007	579.0	21.8
Income relative to Poverty Line Intercepts								
50-100 %	0.47	0.021	381.1	35.5	0.72	0.022	407.1	37.3
100-150 %	-0.32	0.028	688.5	63.9	-0.19	0.030	-1,134.5	64.5
150-200 %	-0.16	0.041	-152.5	112.5	-0.63	0.042	423.6	119.5
>200 %	0.31	0.007	-1,099.9	18.1	0.59	0.008	-1,155.2	19.2
Income relative to Poverty Line Slopes								
≤ 50 %	-0.20	0.023	113.5	41.7	0.00	0.027	724.3	46.0
50-100 %	-0.83	0.025	-705.2	42.7	-0.94	0.026	-795.9	45.2
100-150 %	0.31	0.022	-1,055.8	50.6	0.36	0.023	216.0	51.1
150-200 %	0.21	0.023	-581.5	64.5	0.55	0.024	-834.3	68.6
>200 %	0.05	0.000	-28.0	2.5	0.03	0.001	-36.7	2.8
Age 18-29	-0.05	0.005	-94.7	11.2	0.05	0.005	166.0	11.3
Age 30-39	-0.15	0.004	15.8	10.0	-0.12	0.004	36.0	9.6
Age 50-59	0.01	0.004	-141.4	10.8	0.07	0.004	31.3	10.5
Age 60-69	0.13	0.005	-201.3	13.6	0.19	0.005	141.6	13.8
Age 70 or older	0.31	0.006	-340.3	14.4	0.37	0.006	82.2	14.7
Any income from capital	0.31	0.004	-600.3	16.7	0.38	0.004	-580.9	16.6
# of persons in HH	-0.22	0.002	373.2	4.1	-0.20	0.002	567.5	4.2
# of children in HH	0.00	0.003	874.4	5.8	-0.04	0.002	825.2	5.6
Not married, no children	-0.50	0.005	171.2	12.8	-0.58	0.005	348.5	12.7
Not married, children	-0.54	0.005	84.3	10.5	-0.54	0.004	111.4	10.3
Married, no children	-0.19	0.006	160.9	15.1	-0.24	0.005	295.1	15.2
Linguistic Isolation	-0.02	0.005	-160.0	11.0	-0.07	0.005	159.6	10.8
# of persons employed	0.08	0.002	-273.3	5.2	0.08	0.002	-297.6	5.2
Anyone in HH employed	0.04	0.005	-3.7	12.2	-0.02	0.005	-94.9	12.8
Elderly or disabled in HH	-0.14	0.004	30.1	9.5	-0.14	0.004	-363.8	9.0
Single household	0.16	0.004	-1,351.8	13.0	0.29	0.005	-1,388.7	14.0
Unemployed	-0.15	0.006	-243.5	12.7	-0.10	0.005	-445.3	11.8
Not in labor force	-0.09	0.004	-117.9	9.9	-0.11	0.004	7.4	10.3
Female	0.00	0.003	23.1	7.4	-0.13	0.003	229.7	7.4
Not a US citizen	0.17	0.004	-372.8	10.2	0.35	0.004	-503.8	9.7
White	0.34	0.003	-292.7	6.5	0.47	0.003	-138.6	6.7
Less than High School	-0.04	0.004	321.9	8.9	-0.17	0.004	-16.7	9.0
High School Degree	-0.04	0.003	86.4	8.5	-0.14	0.003	60.7	8.6
College or more	0.20	0.004	110.3	11.9	0.20	0.004	9.8	12.5
Disabled	-0.12	0.008	16.4	19.7	0.00	0.007	326.0	19.1
Disabled, not working	0.04	0.008	-12.5	20.7	-0.20	0.008	-270.6	20.4
Speaks English poorly	-0.14	0.005	269.1	11.6	-0.17	0.005	316.2	11.0
Speaks no English	0.23	0.009	233.0	16.8	-0.01	0.009	259.7	17.7
Constant	1.73	0.010	689.9	24.3	1.53	0.010	266.8	25.5
Conditional Variance	1,983	2.7			2,054	2.7		
Left Truncation Point	16	0.004			16	0.004		
# of observations	52,104				49,576			

Notes: Parameter estimates of truncated normal conditional distribution. The columns labeled "Mass Point" determine the probability that the household does *not* receive SNAP and can be interpreted like probit coefficients. The columns labeled "Amounts" determine amounts conditional on receipt. Demographic Characteristics refer to the household head. All analyses use household weights adjusted for PIK-probability.

References

Abowd, J.M., M. Stinson, and G. Benedetto. 2006. "Final report to the social security administration on the SIPP/SSA/IRS public use file project." U.S. Census Bureau, Longitudinal Employer-Household Dynamics Program.

Abowd, John, and Martha Stinson. 2011. "Estimating Measurement Error in SIPP Annual Job Earnings: A Comparison of Census Bureau Survey and SSA Administrative Data." Center for Economic Studies, U.S. Census Bureau Working Papers 11-20.

Andrews, Donald W. K. 1997. "A Conditional Kolmogorov Test." *Econometrica*, 65(5): 1097–1128.

Autor, David H., Lawrecne F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics*, 90: 300–323.

Black, Dan A., and Jeffrey A. Smith. 2006. "Estimating the returns to college quality with multiple proxies for quality." *Journal of Labor Economics*, 24(3): 701–728.

Black, Dan A., Seth Sanders, and Lowell Taylor. 2003. "Measurement of Higher Education in the Census and Current Population Survey." *Journal of the American Statistical Association*, 98: 545–554.

Bollinger, Christopher R. 1998. "Measurement Error in the Current Population Survey: A Nonparametric Look." *Journal of Labor Economics*, 16(3): 576–594.

Bollinger, Christopher R. 2003. "Measurement error in human capital and the black-white wage gap." *Review of Economics and Statistics*, 85(3): 578–585.

Bollinger, Christopher R., and Martin H. David. 1997. "Modeling Discrete Choice With Response Error: Food Stamp Participation." *Journal of the American Statistical Association*, 92(439): 827–835.

Bollinger, Christopher R., and Martin H. David. 2005. "I didn't tell, and I won't tell: dynamic response error in the SIPP." *Journal of Applied Econometrics*, 20(4): 563–569.

Bonhomme, Stephane, and Jean-Marc Robin. 2010. "Generalized Non-parametric Deconvolution with an Application to Earnings Dynamics." *Review of Economic Studies*, 77(2): 491–533.

Bound, John, and Alan B. Krueger. 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics*, 9(1): 1–24.

Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Measurement error in survey data." In *Handbook of Econometrics*. Vol. 5, , ed. James J. Heckman and Edward Leamer, Chapter 59, 3705–3843. Amsterdam:Elsevier.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers. 1994. "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics*, 12(3): 345–368.

Brownstone, David, and Robert G. Valletta. 1996. "Modeling Earnings Measurement Error: A Multiple Imputation Approach." *The Review of Economics and Statistics*, 78(4): 705–717.

Butler, J. S., and Jennie E. Raymond. 1996. "The Effect of the Food Stamp Program on Nutrient Intake." *Economic Inquiry*, 34(4): 781–798.

Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. 2006. *Measurement Error in Nonlinear Models: A Modern Perspective. Monographs on Statistics and Applied Probability*. 2nd ed., Boca Raton:Chapman & Hall/CRC.

Chen, Xiaohong. 2007. "Large Sample Sieve Estimation of Semi-Nonparametric Models." In *Handbook of Econometrics.* Vol. 6, , ed. James J. Heckman and Edward Leamer, Chapter 76, 5549–5632. Amsterdam:Elsevier.

Chen, Xiaohong, Han Hong, and Elie Tamer. 2005. "Measurement Error Models with Auxiliary Data." *The Review of Economic Studies*, 72(2): 343–366.

Currie, Janet. 2004. "The Take Up of Social Benefits." National Bureau of Economic Research Working Paper 10488.

Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature*, 46(2): 333–83.

Currie, Janet, and Jeffrey Grogger. 2001. "Explaining Recent Declines in Food Stamp Program Participation." *Brookings-Wharton Papers on Urban Affairs*, 203–244.

Drechsler, Jrg, and Jerome P. Reiter. 2010. "Sampling With Synthesis: A New Approach for Releasing Public Use Census Microdata." *Journal of the American Statistical Association*, 105(492): 1347–1357.

Fraker, Thomas, and Robert Moffitt. 1988. "The effect of food stamps on labor supply: A bivariate selection model." *Journal of Public Economics*, 35(1): 25–56.

Fuller, Wayne A. 1987. *Measurement Error Models. Probability and mathematical statistics*, New York:Wiley.

Gundersen, Craig, and Brent Kreider. 2008. "Food Stamps and Food Insecurity What Can Be Learned in the Presence of Nonclassical Measurement Error?" *Journal of Human Resources*, 43(2): 352–382.

Hagstrom, Paul A. 1996. "The Food Stamp Participation and Labor Supply of Married Couples: An Empirical Analysis of Joint Decisions." *The Journal of Human Resources*, 31(2): 383–403.

Haider, Steven J., Alison Jacknowitz, and Robert F. Schoeni. 2003. "Food Stamps and the Elderly: Why Is Participation so Low?" *The Journal of Human Resources*, 38: 1080–1111.

Hausman, Jerry A., Jason Abrevaya, and Fiona M. Scott-Morton. 1998. "Misclassification of the dependent variable in a discrete-response setting." *Journal of Econometrics*, 87(2): 239–269.

Hausman, Jerry A., Whitney K. Newey, Hidehiko Ichimura, and James L. Powell. 1991. "Identification and estimation of polynomial errors-in-variables models." *Journal of Econometrics*, 50(3): 273–295.

Hisnanick, John J., Tracy Loveless, and John Chesnut. 2007. "Evaluation Report Covering Receipt of Food Stamps." U.S. Census Bureau American Community Survey Content Test Report H.6, Washington, D.C.

Hong, Han, and Elie Tamer. 2003. "A simple estimator for nonlinear error in variable models." *Journal of Econometrics*, 117(1): 1–19.

Horowitz, Joel L., and Charles F. Manski. 1995. "Identification and Robustness with Contaminated and Corrupted Data." *Econometrica*, 63(2): 281–302.

Hoynes, Hilary W., and Diane W. Schanzenbach. 2009. "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program." *American Economic Journal: Applied Economics*, 1(4): 109–39.

Hsiao, Cheng. 1989. "Consistent estimation for some nonlinear errors-in-variables models." *Journal of Econometrics*, 41(1): 159–185.

Hu, Yingyao, and Geert Ridder. 2012. "Estimation of nonlinear models with mismeasured regressors using marginal information." *Journal of Applied Econometrics*, 27(3): 347–385.

Hu, Yingyao, and Susanne M. Schennach. 2008. "Instrumental Variable Treatment of Nonclassical Measurement Error Models." *Econometrica*, 76(1): 195–216.

Ichimura, Hidehiko, and Elena Martinez-Sanchis. 2009. "Estimation and Inference of Models with Incomplete Data by Combining Two Data Sets." Unpublished Manuscript.

Kapteyn, Arie, and Jelmer Y. Ypma. 2007. "Measurement Error and Misclassification: A Comparison of Survey and Administrative Data." *Journal of Labor Economics*, 25(3): 513–551.

Keane, Michael, and Robert Moffitt. 1998. "A Structural Model of Multiple Welfare Program Participation and Labor Supply." *International Economic Review*, 39(3): 553–589.

Levitan, Mark, and Trudi Renwick. 2010. "Using the American Community Survey to Implement a National Academy of Sciences-Style Poverty Measure: A Comparison of Imputation Strategies." American Statistical Association, Social Statistics Section Unpublished Manuscript.

Li, Tong. 2000. "Estimation of nonlinear errors-in-variables models: a simulated minimum distance estimator." *Statistics & probability letters*, 47(3): 243–248.

Li, Tong. 2002. "Robust and consistent estimation of nonlinear errors-in-variables models." *Journal of Econometrics*, 110(1): 1–26.

Little, Roderick J.A. 1993. "Statistical analysis of masked data." *Journal of Official Statistics*, 9: 407–407.

Lynch, Victoria, Dean M. Resnick, Jane Stavely, and Cynthia M. Taeuber. 2007. "Differences in Estimates of Public Assistance Recipiency Between Surveys and Administrative Records." U.S. Census Bureau, Washington, D.C.

Marquis, Kent H., and Jeffrey C. Moore. 1990. "Measurement Errors in SIPP Program Reports." In *Proceedings of the 1990 Annual Research Conference.* 721–745. Washington, D.C.:U.S. Bureau of the Census.

McFadden, Daniel. 1989. "A method of simulated moments for estimation of discrete response models without numerical integration." *Econometrica*, 995–1026.

Meyer, Bruce, and Nikolas Mittag. 2013*a*. "Evidence on Misreporting of Government Programs in Household Surveys From Linked Administrative Data." Unpublished Manuscript.

Meyer, Bruce, and Nikolas Mittag. 2013*b*. "Report on Linking Administrative Data on Government Programs From New York State to Survey Data." Unpublished Manuscript.

Meyer, Bruce, and Robert Goerge. 2011. "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation." Center for Economic Studies, U.S. Census Bureau Working Papers 11-14.

Meyer, Bruce D. 2010. "The Effects of the Earned Income Tax Credit and Recent Reforms." *Tax Policy and the Economy, Volume 24*, 153–180. The University of Chicago Press.

Meyer, Bruce D., and James X. Sullivan. 2003. "Measuring the Well-Being of the Poor Using Income and Consumption." *The Journal of Human Resources*, 38: 1180–1220.

Meyer, Bruce D., and Nikolas Mittag. 2012. "Misreporting in Binary Choice Models." Unpublished Manuscript.

Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan. 2009. "The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences." Harris School of Public Policy Studies, University of Chicago Working Papers 0903.

Meyer, Bruce, Robert Goerge, and Nikolas Mittag. 2012. "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation." Unpublished Manuscript.

Mittag, Nikolas. 2012. "Imputations: Benefits, Risks and a Method to Handle Missing Data." Unpublished Manuscript.

Newey, Whitney K., and Daniel L. McFadden. 1994. "Large Sample Estimation and Hypothesis Testing." In *Handbook of Econometrics.* Vol. 4, , ed. Robert F. Engle and Daniel L. McFadden, Chapter 36, 2111–2245. Amsterdam:Elsevier.

NORC. 2011. "Assessment of the US Census Bureaus Person Identification Validation System." NORC at the University of Chicago Final Report presented to the US Census Bureau.

Poterba, James M., and Lawrence H. Summers. 1986. "Reporting Errors and Labor Market Dynamics." *Econometrica*, 54(6): 1319–1338.

Ridder, Geert, and Robert Moffitt. 2007. "The Econometrics of Data Combination." In *Handbook of Econometrics.* Vol. 6B, , ed. James J. Heckman and Edward Leamer, Chapter 75, 5469–5547. Amsterdam:Elsevier.

Rubin, Donald B. 1993. "Discussion: Statistical Disclosure Limitation." *Journal of Official Statistics*, 9(2): 461–468.

Schafer, Joseph L., and Nathaniel Schenker. 2000. "Inference with Imputed Conditional Means." *Journal of the American Statistical Association*, 95(449): 144–154.

Schennach, Susanne M. 2004. "Estimation of Nonlinear Models with Measurement Error." *Econometrica*, 72(1): 33–75.

Scholz, John .K., Robert Moffitt, and Benjamin Cowan. 2009. "Trends in income support." *Changing poverty, changing policies*, 203–41.

Sherman, Arloc. 2009. "Safety Net Effective at Fighting Poverty But Has Weakened for the Very Poorest." Center on Budget and Policy Priorities Policy Report, Washington, D.C.

Sherman, Arloc. 2011. "Poverty and Financial Distress Would Have Been Substantially Worse in 2010 Without Government Action, New Census Data Show." Center on Budget and Policy Priorities Policy Report, Washington, D.C.

Steuerle-Schofield, Lynne, Brian Junker, Lowell J. Taylor, and Dan A. Black. 2012. "Limitations of Institutional Plausible Values." Unpublished Manuscript.

Tiehen, Laura, Dean Jolliffe, and Craig Gundersen. 2012. "Alleviating Poverty in the United States: The Critical Role of SNAP Benefits." U.S. Department of Agriculture, Economic Research Service ERR-132.

U.S. Department of Agriculture. 2011. "Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2010." The Food and Nutrition Service, Office of Research and Analysis, Alexandria, VA.

Whitmore, Diane. 2002. "What Are Food Stamps Worth?" Princeton University, Department of Economics, Industrial Relations Section. Working Papers 847.