

Survey Measurement of Income Tax Rates*

Michael S. Gideon[†]

November 11, 2014

Abstract

This paper uses survey techniques to characterize income tax rate perceptions. Data come from self-reported average tax rates (ATRs) and marginal tax rates (MTRs) in two subsequent waves of the Cognitive Economics Study. Comparing survey measures of tax rates with those computed from survey measures of income, respondents systematically overestimate ATRs while perceptions of MTRs are accurate at the mean. A statistical model is used to analyze heterogeneity in perceptions and to impute true rates that account for measurement error in income. Perceptions of MTR and ATR (conditional on the true rates) vary depending on cognitive ability and use of tax preparers. Combining information from self-reported income and tax rates reduces measurement error in the imputed true tax rates.

Average tax rate; Marginal tax rate; Measurement error; Perceptions

1 Introduction

Income tax rates influence household choices about labor supply, consumption and wealth accumulation. Of all policy parameters, perhaps none is more germane to such a wide array of important household decisions. Tax rate changes are often used to estimate structural parameters, such as labor supply elasticities and intertemporal substitution elasticities, that are subsequently used to evaluate the efficiency costs of labor income taxation, as well as forecast the impact of economic policy and macroeconomic shocks over time.

Despite the extensive literature on identification and estimation of responses to tax changes, the issue of tax rate measurement has been largely ignored. Tax returns provide a “gold standard” measure of true

*This paper is based on the first chapter of my Ph.D. thesis at the University of Michigan, Ann Arbor. I thank Matthew Shapiro, Robert Willis and Pamela Giustinelli for continued feedback and support, and David Agrawal, Charlie Brown, Andrew Goodman-Bacon, Italo Gutierrez, James Hines, Joanne Hsu, Peter Hudomiet, Joshua Hyman, Gretchen Lay, Brooke McFall, Eric Ohrn, Daniel Reck, Tyler Shumway, Dan Silverman, Joel Slemrod, Jeffrey Smith, Bryan Stuart, Caroline Weber and numerous seminar participants for helpful discussions and comments at various stages of the project. Financial support from the F. Thomas Juster Economic Behavior Research Award and National Institute on Aging program project grant P01-AG026571 is gratefully acknowledged.

[†]Present address: NORC at the University of Chicago, 1155 E. 60th Street, 2nd Floor, Chicago, IL 60637. Email: gideonmi@gmail.com. phone: 734-756-1896

income tax rates. When unavailable, tax rates are typically computed from measures of self-reported income using NBER's TAXSIM tax simulator. Income measurement error and unobserved filing behavior (such as claiming deductions, exclusions and credits) make computed measures inherently noisy. Moreover, if perceptions are inaccurate, then it is not obvious whether to focus on perceptions or true rates. If people act upon their beliefs rather than true rates, then estimated parameters using tax returns are still difficult to interpret without additional information about perceptions.

While the survey approach introduces other problems—namely, whether respondents give accurate answers—it can provide a potentially important source of information about tax beliefs and private information about filing behavior. If perceptions matter, we might as well just ask people about their tax rates. And if perceptions are accurate, then survey questions provide an important source of information about actual tax rates.

My central contribution is the measurement, identification and characterization of heterogeneous tax perceptions using survey measures of average tax rates (ATR) and marginal tax rates (MTR). My methodological contribution is the statistical model used to combine information from survey and computed MTRs and ATRs, which provides a way to use survey measures to separate measurement error when imputing the true tax rates.

Data from the Cognitive Economics (CogEcon) Study provides a unique opportunity to analyze income tax perceptions. First, there are two waves of data containing self-reported marginal tax rates and average federal income tax rates. While previous surveys have asked about marginal tax rates (MTR) or average tax rates (ATR), this is the first to ask about both. Second, it includes data on whether people use tax preparers to file tax returns, information that is rarely (if ever) available on household surveys. Finally, the data are linked to measures of numerical and verbal ability from in-depth cognitive assessments, as well as a measure of financial sophistication constructed using a long battery of questions on previous waves of the survey. Data on labor supply, savings decisions and wealth are also available, along with a rich set of demographic variables.

Survey tax rates and survey income (which is used to compute MTR and ATR) provide three indicators of latent taxable income. I exploit the known relationship between taxable income and tax rates to identify jointly the distribution of income measurement error and tax rate biases. By making parameteric assumptions on the distribution of income and the tax rate heterogeneity, these measures of tax perceptions help identify income measurement error.

I find evidence that respondents systematically overestimate average tax rates and exhibit heterogeneity in perceived rates. Perceptions of marginal tax rates are accurate, at the mean, but exhibit mean reversion and substantial heterogeneity. Respondents with tax preparers report systematically larger MTR and ATR. Number series and financial sophistication scores are associated with ATR errors, even conditioning on education, wealth and having a financial occupation. This means measures of ability help explain variation in tax perceptions, above and beyond what is explained by education. At the same time, there appears to

be a uniform cluelessness about marginal tax rates.

I then impute true rates using the survey measures alongside computed rates even when I allow both the survey and computed ones to be subject to error. Self-reported measures help improve upon the inherently noisy computed rates. In particular, in both 2010 and 2012, 8.3% of observations have imputed MTR that are not equal to the MTR computed from their reported income. And 15% had a revised MTR in at least one of the waves of the study. Asking about both marginal and average tax rates provides leverage when distinguishing between income measurement error and heterogeneity in beliefs.

Relationship to existing literature It is surprising that little attempt has been made to directly measure and systematically analyze tax perceptions. The CogEcon data is the first, to my knowledge, to have perceptions of both marginal and average income tax rates and the first with perceptions across multiple years. Nevertheless, it is useful to compare results about ATR and MTR, individually, to several older studies of tax perceptions that use survey or interview data.

Older studies of tax perceptions find mixed results across time and different countries. Some studies find taxpayers slightly overestimate their tax liability (Wagstaff, 1965; Kapteyan and Ypma, 2007), while others find they slightly underestimate it (Enrick, 1963; 1964).¹ Some studies find that taxpayers overestimate their marginal tax rate (Brown, 1968) while others find taxpayers underestimate it (Lewis, 1978; Fujii and Hawley, 1988). There is also evidence of mean reversion, with lower income households overestimating and higher-income households underestimating tax liabilities (Wagstaff, 1965; Auld, 1979).

I find evidence suggesting that taxpayers overestimate how much they pay in federal income taxes and report approximately unbiased perceptions of marginal tax rates, at the mean. The results affirm and contrast with previous evidence about perceptions of tax rates.

2 Survey methodology

2.1 Cognitive Economics Study

The CogEcon sample consists of households over 50 years old who were chosen using a random sample selection design. Partners were also included in the sample, regardless of their age. These respondents also participated in the Cognition and Aging in the USA Study (CogUSA), which included a detailed 3-hour face-to-face psychometric evaluation, consisting of detailed cognitive and personality assessments.² While

¹Kapteyn and Ypma (2007) do not focus on tax perceptions, yet have results that address the question of self-reported versus true tax liability. They use data from Sweden and find that people overreport their income tax liability relative to the tax liability in administrative records. The Swedish administrative data comes from LINDA (Longitudinal Individual Data for Sweden) and the survey data is from SHARE (Survey of Health, Ageing, and Retirement in Europe).

²The Cognitive Economics Study is supported by National Institute on Aging program project P01-AG026571, "Behavior on Surveys and in the Economy Using HRS," Robert J. Willis, PI. More information about CogEcon is available online (<http://cogecon.isr.umich.edu/survey/>) and in the data documentation (Fisher et al., 2011; Gideon et al., 2013). CogUSA is part of the Unified Studies of Cognition (CogUSC) led by cognitive psychologist Jack McArdle at the University of Southern California. All respondents who completed the first waves of CogUSA, and were not otherwise involved with the Health and Retirement Study (HRS), were invited to complete the first CogEcon survey. More information about CogUSA is available at cogusc.usc.edu.

the CogEcon/CogUSA sample was initially selected to be nationally representative of the older population, internal work by the CogUSA team suggests that this population has higher cognitive ability than the population at large.

Data come from CogEcon 2011 and CogEcon 2013, which are the third and fourth waves of the study. New questions asking about federal income tax rates were added in 2011 and repeated in 2013. Other new questions ask about tax return filing behavior and the use of paid tax preparers and tax software (such as TurboTax). It also repeats questions from prior waves (fielded in 2008 and 2009) about income, employment, assets and debts, among other topics.

CogEcon 2011 was fielded between October 2011 and January 2012, while CogEcon 2013 was fielded between October 2013 and January 2014. There was an internet mode and a mail mode. Households with internet access were invited to the web version. Questions are typically the same across mode. In the main estimation sample, 68.9% completed the web version in both 2011 and 2013, 26.8% completed the mail version both waves, and 4.3% completed the web version one wave and the mail version in the other.

My estimation sample includes the 348 respondents who gave valid responses to all four tax rate questions and reported total income above \$5,000 in both waves. This includes 302 households, overall, and 46 households with two respondents. I include both respondents, except where specified otherwise.

There were 748 respondents who completed CogEcon 2011, among whom 564 answered both tax rate questions. There were 694 respondents who completed CogEcon 2013, among whom 488 answered both tax rate questions. Among the 629 respondents who submitted both CogEcon 2011 and 2013, 382 respondents answered the tax rates questions in both waves. Given the importance of the income measurement in this paper, I restrict my estimation sample to the 348 respondents who answered both tax rate questions and reported total income above \$5,000 in both waves.

Table 1 provides summary statistics for the entire sample of respondents who completed both CogEcon 2011 and CogEcon 2013, as well as for the sample with complete income and tax rates data who are used in the analyses. The typical respondent is 67 years old, has higher income than the overall population (mean \$84,000; median \$66,750), and roughly two-thirds are married. Among the sample of respondents, those who answered the tax questions are on average younger, higher income, less likely to use professional tax preparers and score higher on the number series cognitive assessments.

2.2 The survey instrument

Questions were designed to elicit perceptions of federal individual income tax rates. The set of questions in CogEcon 2011 opens with the following.

These questions focus on current and future federal income tax rates, both in general and for you personally. The marginal tax rate is the tax rate on the last dollars earned. For example, if a household's income tax bracket has a marginal tax rate of 15%, then a household owes an extra \$15 of taxes when it earns an extra \$100. Answer each question with a percentage between 0 and 100. Please provide your best estimate of the marginal tax rate even if you are not sure.

Table 1: Sample summary statistics

Percent	All respondents	Sample
Male (percent)	44.5	52.0
Married (percent)	68.4	71.6
Non-white (percent)	8.9	7.2
Education (years)	14.5	14.7
Age (years, in 2011)	67.4	66.2
Employed during 2010 (percent)	41.0	58.9
Employed during 2012 (percent)	43.1	51.7
Paid tax preparer (percent)	51.5	48.5
Mail response mode (percent, in 2011)	31.2	21.0
Number Series Score (standardized)	0.04	0.24
Log of Income (Individual Average)	11.00	11.10
	(0.84)	(0.78)
Log of Wealth (Individual Average, Positive)	12.76	12.95
	(1.53)	(1.33)
Responses	629	348

Notes: This table shows summary statistics for the entire sample and the estimation sample. The entire sample includes respondents who completed CogEcon 2011 and CogEcon 2013, whereas the estimation sample is restricted to respondents with complete income and tax rates data, as described in the text. Number series score is standardized around mean 0. The demographics are similar, although, respondents who answered the tax questions had slightly higher number series scores, were more likely to be male, and less likely to use paid tax preparers.

These questions are about federal income taxes only; please do not include state or local taxes, or payroll taxes for Social Security and Medicare.

After this introduction, there were two questions about marginal tax rates imposed on households in the top income tax bracket. The first was about marginal tax rates in 2010 and the second asked about expected rates in 2014. These were followed by three questions about the respondents' own tax rates. The first was about their average tax rate in 2010, the second about their marginal tax rate in 2010 and the third about their expected marginal tax rate in 2014. I focus on the following questions eliciting perceptions of average and marginal income tax rates in 2010.

We now want to ask you about your household's federal taxes. Please use the same definitions of federal income tax and marginal tax rate as on the previous page.

[ATR 2010]. Please think about your household's income in 2010 and the amount of federal income tax you paid, if any. Approximately what percentage of your household income did you pay in federal income taxes in 2010? _____%

[MTR 2010]. Now we want to ask about your household's marginal income tax rate. Please think about your household's federal income tax bracket and the tax rate on your last dollars of earnings. In 2010, my household's marginal tax rate was _____%.

Appendix A.1 presents the exact wording and ordering of all five.

The questions were written in precise yet simple language to elicit perceptions of average and marginal tax rates. The question refers to an "income tax bracket" as a way to elicit perceptions of one's statutory marginal income tax rate without explicitly distinguishing between statutory and effective rates. While respondents might be confused about whether we want statutory or effective rates, the people who understand the

difference are expected to lean toward giving the statutory rate. For the average tax rate question we intentionally used a clearly specified tax concept but vague definition of income.³

Pilot testing revealed that respondents better understood questions explicitly asking for the marginal tax rate than questions asking for the tax rate on the "last \$100 of income."⁴ After testing both versions we learned that people would sometimes say zero because they were thinking of the paper tax table and whether or not the additional money would actually move them from one cell in the tax table to another cell. Other testers mentioned that \$100 is small and that the tax can't be much.

Questions eliciting marginal and average tax rates were asked again in CogEcon 2013. This set of questions opens with the following description of the section and definition of marginal tax rates.

These next two questions focus on your federal income tax rates. These questions are about federal income taxes only; please do not include state or local taxes, or payroll taxes for Social Security and Medicare.

The first question asked for an approximate average federal income tax rate in 2012, the next for the marginal tax rate in 2012. The exact wording was as follows:

[ATR 2012]. Please think about your household's income in 2012 and the amount of federal income tax you paid, if any. Approximately what percentage of your household income did you pay in federal income taxes in 2012? _____%

[MTR 2012]. Now we want to ask about your household's marginal income tax rate. The marginal income tax rate is the tax rate on the last dollars earned. Please think about your household's federal income tax bracket and the tax rate on your last dollars of earnings. In 2012, my household's marginal tax rate was _____%.

There are a few important differences in question wording across waves. First, the definition of marginal tax rate was given in the question about marginal tax rate rather than at the beginning of the section. The definition seems most appropriate at the beginning of the question for which it is applicable. Second, the definition of marginal tax rate did not include an example involving 15% that might have anchored respondents to that number. This would show up as more people reporting fifteen in CogEcon 2011 than in CogEcon 2013 and would likely be among those who are not confident about their rates. Third, in CogEcon 2011 the questions were immediately following questions about marginal tax rates for households in the top income tax bracket. This could have anchored respondents to report higher rates in 2011 than in 2013, when the questions were not following questions about the top rates. There is no noticeable difference in the level of the rates across the two waves. Last, the questions in 2013 did not require that respondents give a number between 0 and 100.

³This is because we measure household income using a broad question earlier in the survey. If we specified adjusted gross income (AGI) then knowledgeable respondents would use this income concept while others would not. We would have no way to disentangle knowledge of the tax concept from knowledge of the income concept.

⁴Preliminary versions of the survey instruments were tested on two groups. First, they were tested by people who were in the demographic of the Cognitive Economics Study respondents— Americans age 50 and older, who are slightly better educated and with higher income than the age demographic at large. Second, they were tested as part of the standard Survey Research Operations procedure for testing survey instruments.

Table 2: Taxable Income Thresholds by Tax Bracket and Filing Status, 2010 and 2012

MTR	Single Filers		Married Filing Jointly	
	2010	2012	2010	2012
10	0	0	0	0
15	8,375	8,700	16,750	17,400
25	34,000	35,350	68,000	70,700
28	82,400	85,650	137,300	142,700
33	171,850	178,650	209,250	217,450
35	373,650	388,350	373,650	388,350

Notes: This table shows the lower end of the taxable income thresholds for each tax bracket. The tax brackets are associated with marginal tax rates (MTR), and the columns distinguish between single and married tax filing status, and between tax years 2010 and 2012.

2.3 Measuring tax rates using income

Self-reported income data are used to compute Adjusted Gross Income (AGI) and taxable income (TI). Filing status and TI determine statutory MTR and tax liability, and the computed ATR equals tax liability divided by AGI. I assume the true marginal tax rate is the statutory rate, which is consistent with how the question was worded.

Statutory marginal tax rates for wage and salary income were 10%, 15%, 25%, 28%, 33% and 35%, the levels set in the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA). Table 2 presents the taxable income thresholds associated with these marginal tax rates for tax years 2010 and 2012, broken down by filing status.

Unfortunately, I do not have data on respondents' filing status and I must make assumptions based on marital status. Single respondents are assumed to file as single (rather than head of household), and married respondents are assumed to file jointly. Nevertheless, if married respondents file separately then I will underestimate their true MTR and, therefore, overestimate the difference between perceived and true MTR; the opposite will be true for singles, for whom I will overestimate their MTR and underestimate their bias. The income variables come from self-reported information about different sources of income. I use the NBER's TAXSIM tax rate calculator to transform this vector of income variables into AGI.⁵ Taxable income equals AGI minus exemptions and deductions. Information about dependent exemptions was collected in the survey and I assume all taxpayers claim the standard deduction. Details about constructing AGI and TI are in Appendix B.1.

This procedure ignores a couple important institutional details. First, income from some investments gets taxed according to a different rate schedule. Prior to 2003, dividends were taxed as ordinary income, but JGTRRA reduced the top tax rate on qualified dividends to the long-term capital gains rate of 15%. If I incorrectly assume income is taxed as ordinary income rather than investment income then computed rates will be larger than true rates. Second, certain earners are subject to the Alternative Minimum Tax (AMT). While there is no definitive way to know whether respondents were among this group, the fraction of the

⁵See Feenberg and Coutts (1993) for a discussion of TAXSIM.

Table 3: Tax rate responses in 2013 conditional on responses in 2011

Tax rate responses in 2011	Total	Tax rate responses in 2013							
		skipped both	MTR, skip ATR	ATR, skip MTR	ATR > MTR	ATR = MTR	ATR < MTR	both=0	
	N=645	14.6	0.8	15.4	12.3	22.5	23.4	11.2	100
		Tax rate responses in 2013, conditional on 2011 response							
skipped MTR & ATR	16.4	43.4	1.9	17.0	11.3	9.4	7.6	9.4	100
MTR, skipped ATR	0.2	100	0	0.0	0	0.0	0.0	0	100
ATR, skipped MTR	10.4	10.5	1.5	34.3	11.9	11.9	23.9	6.0	100
ATR > MTR	14.7	9.5	0	12.6	13.7	33.7	24.2	6.3	100
ATR = MTR	28.1	7.7	0.6	11.1	16.6	37.6	21.6	5.0	100
ATR < MTR	21.6	7.2	0.7	15.1	11.5	16.6	42.5	6.5	100
both=0	8.7	12.5	0	8.9	0	7.1	10.7	60.7	100

Notes: This tables shows categories of tax rate responses in CogEcon 2011, again in CogEcon 2013, and then for the responses in 2013 conditional on the response in 2011.

sample is likely small. In Tax Year 2010, 2.8% of tax returns were required to pay additional tax because of the AMT. Around 4.3% of households between ages 55 and 65 were subject to the AMT, whereas 2.6% of those 65 and older (SOI, Publication 1304, Table 1.1).

3 Descriptive analysis

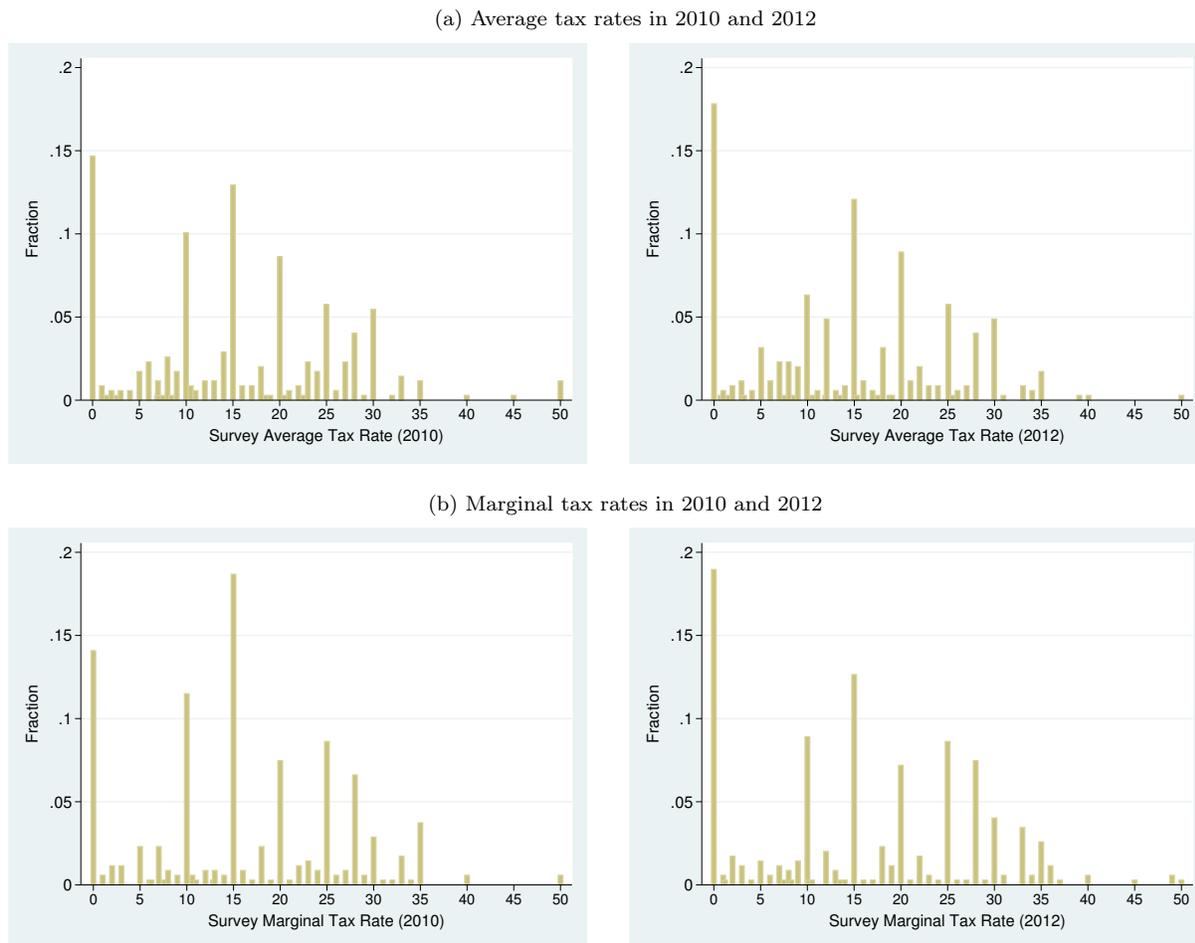
Table 3 presents the distribution of respondents' answers across categories of average and marginal tax rates. The first three columns are the percent of respondents (of the subsample) who skipped both questions, the percent who only skipped the MTR question but answered the ATR question, and the percent who answered zero for both. The next three columns show the percent of respondents with survey MTR larger than, smaller than, and equal to ATR, conditional on answering both with at least one non-zero numbers. Around sixteen percent of respondents skipped both questions and another ten percent answered the question about ATR but not about MTR. Only one respondent skipped the question about the average tax rate and answered the question about their marginal tax rate.

Comparing average and marginal rates provides an indicator of whether respondents understand that in the progressive U.S tax system the average tax rate is always (weakly) less than the marginal tax rate. The relationship between MTR and ATR is due to the fact that (i.) statutory MTRs increase with the level of taxable income, and (ii.) deductions and exemptions are always strictly positive. Whenever taxable income is positive, (ii.) implies that MTR is strictly greater than ATR. Only a third of responses are consistent with the progressive tax schedule ($MTR > ATR$). The largest percent of respondents gave the same number for both ATR and MTR, which was 37 percent of the sample, or 29 percent if I exclude those who reported zero for both.

Figure 1 displays a histogram of the survey measures of MTR and ATR from CogEcon 2011 and CogEcon

2013. Reported marginal tax rates are in the top panel, while average tax rates are in the bottom panel. In both waves there are concentrated responses at 0, 15 and other statutory MTRs or round numbers. In CogEcon 2011, for example, roughly a quarter of responses were multiples of ten, around half were multiples of five. Overall, there is substantial heterogeneity in reported tax rates.

Figure 1: Histograms of tax rate responses with labeled mass points and statutory rates



Notes: Reported tax rates greater than 50 are binned at 50.

Differences between the survey and computed rates is perhaps more interesting. Table 4 shows summary statistics across the various measures of ATR and MTR, as well as for the difference between survey and computed rates. There is substantial heterogeneity across respondents in the size of the difference between survey and computed tax rates (for both MTR and ATR). The survey minus computed difference for ATR is positive and significant (albeit heterogeneous).

Figure 2 plots the mean survey and computed rates across values of total income. The top panel (Figure 2a) shows survey ATR larger than computed ATR across the income distribution. This pattern holds across both waves. The mean of the survey ATR tracks the mean of the computed ATR across the income distribution. The bottom panel (Figure 2b) compares survey and computed marginal tax rates across values

Table 4: Summary statistics for survey and computed MTR and ATR

	Mean	25th %-tile	50th %-tile	75th %-tile	Std. Dev.	Obs.
Survey MTR (2010)	16.0	9.0	15.0	25.0	11.9	565
Computed MTR (2010)	15.0	0.0	15.0	25.0	10.8	565
[Survey MTR] - [Computed MTR]	0.98	-6.0	0.0	7.0	12.4	565
Survey ATR (2010)	15.3	7.0	15.0	23.0	12.2	641
Computed ATR (2010)	8.9	0.0	9.5	14.4	7.2	641
[Survey ATR] - [Computed ATR]	6.4	0.0	4.4	12.0	11.5	641
Survey MTR (2012)	15.8	5.0	15.0	25.0	12.2	487
Computed MTR (2012)	13.4	0.0	15.0	25.0	10.6	487
[Survey MTR] - [Computed MTR]	2.4	-4.0	0.0	9.0	11.6	487
Survey ATR (2012)	14.3	5.0	15.0	22.0	11.1	585
Computed ATR (2012)	8.0	0.0	8.2	13.0	7.1	585
[Survey ATR] - [Computed ATR]	6.3	0.0	4.2	12.0	10.3	585

Notes: This table presents descriptive statistics for the survey tax rates, computed tax rates, and the difference between the rates. The sample includes all respondents with strictly positive adjusted gross income (AGI).

of total income. The mean survey MTR tracking the mean computed MTR, with slight mean reversion, or the tendency to report rates closer to the mean tax rate in the population. People at low levels of total income overestimate their MTR relative to the computed rates, whereas people at higher incomes underestimate MTR. Again, this pattern holds across both waves.

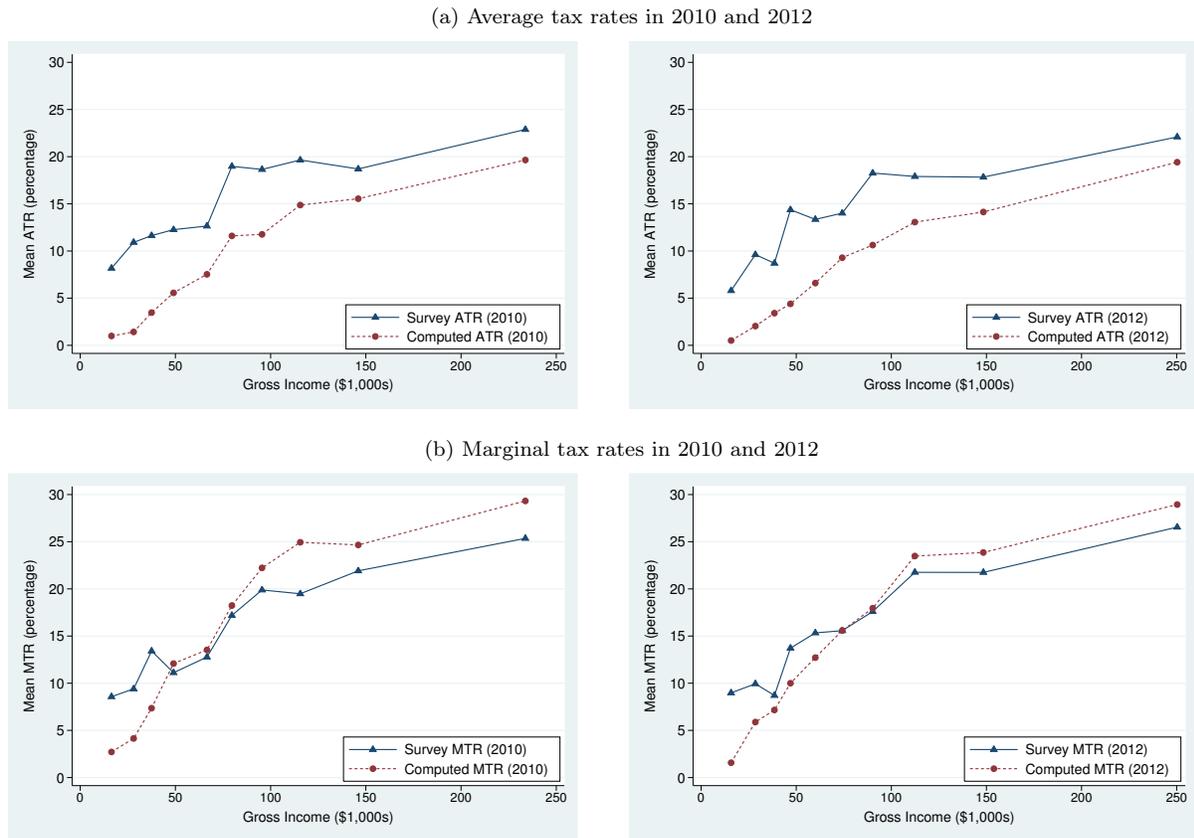
Survey tax rates rise with income. This pattern suggests that survey measures of tax rates may provide information about income. To gain intuition, consider respondents answering completely randomly. The mean would be flat across the distribution of income because the survey measures would contain no information about income. After accounting for random noise, the distribution of survey measures might inform estimation of the distribution of income.

It is also interesting to compare survey and computed rates at an individual level by looking at survey rates minus computed rates. Figure 3 summarizes these differences. Each observation in Figure 3a consists of a pair of ATR errors, defined as survey ATR minus computed ATR. Figure 3b does the same but for survey MTR minus computed MTR. The mass of observations in the first quadrant of Figure 3a means that respondents systematically report ATR larger than what is computed from their income. In contrast, Figure 3b shows respondents scattered throughout all four quadrants, providing evidence of both heterogeneity in systematic errors, along with substantial randomness in the survey measures of marginal tax rates.

4 Structural measurement model

The descriptive statistics provide compelling evidence of systematic differences between survey and computed tax rates, and substantial heterogeneity in these differences across individuals. For a given individual, these differences could reflect systematic misperceptions, unobserved filing behavior, or survey noise. Survey

Figure 2: Binned means of survey and computed measures of marginal and average tax rates, across gross income

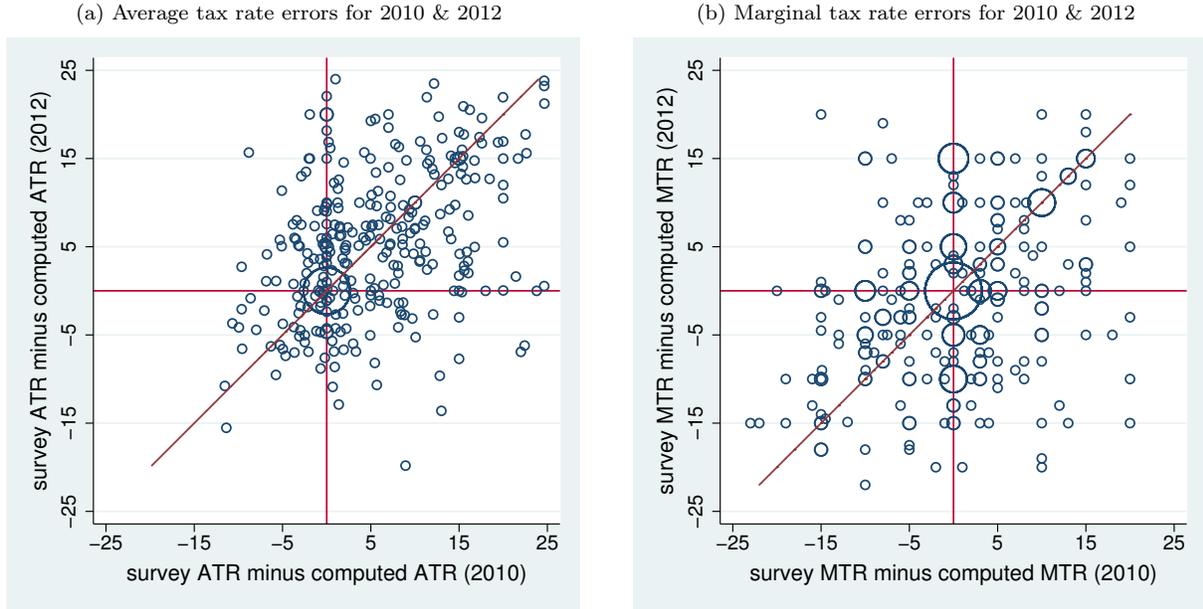


Notes: Observations were grouped into ten equally sized bins of income. The mean tax rate is graphed at the mean income within the bin. The sample used for these figures was restricted to people who provided both ATR and MTR, such that the same simple is in both figures. Respondents with income above \$500,000 are removed to limit the mean of the binned income for the top bin.

responses were interpreted at face value, ignoring the fact that survey data is subject to misreporting or other errors with data handling, and that respondents do not always have information available when completing the survey. Income measurement error causes misclassification of marginal tax rates, and variation in the survey measures comes from systematic misperceptions and random noise. While it is unlikely that random noise in either variable could generate the systematic errors in the previous section, it could lead to erroneous inference about individual observations. Both would lead to inconsistent estimates of behavioral parameters.

I develop a statistical measurement model to combine information from survey and computed average and marginal tax rates. Survey responses about MTR, ATR and income (which is used to compute MTR and ATR) provide three indicators of latent taxable income. I estimate the error distribution for both survey tax rates and income. Separately identifying income and tax rate errors hinges on the known non-linear relationship between income and tax rates, and parametric assumptions about the structure of the three errors. I then impute true rates while correcting for measurement error and systematic differences between

Figure 3: Scatterplot of ATR and MTR errors across waves (weighted)



Notes: Each observation in 3a consists of a pair of ATR errors, which is the difference between survey and computed ATR. Each observation in 3b consists of a pair of ATR errors, which is the difference between survey and computed ATR. In both figures, the x-axis is the error in 2010 and the y-axis is the error in wave 2012. The area of the symbol is proportional to the frequency of observation. Observations in the first quadrant are systemically biased upward whereas observations in the third quadrant report rates smaller than what is computed from income. The diagonal line is at 45%, which corresponds to the difference between survey and computed rates being the same in both waves. Observations with errors greater than 25 or less than -25 are not displayed.

the measures, and analyze heterogeneity in tax rate misperceptions.

My approach is similar to recent models in which the error structure of survey responses is compared to that from administrative data (Kapteyan and Ypma, 2007; Abowd and Stinson, 2013). Unlike these papers, my two measures are fundamentally different variables although they have a known functional relationship. My model uses information from MTR to detect income errors, which is conceptually similar to the approach of Sullivan (2009), who uses information from wages to detect misclassification of occupations. Both wages and occupations come from self-reported survey data.

This estimation procedure is an extension of Kimball, Sahm and Shapiro (2008), who show how to use discrete survey responses to estimate the distribution of a latent continuous variable when responses are subject to error. In KSS, the true underlying latent variable is continuous but the survey responses are discrete. In this application, the true marginal tax rate is a discrete variable, while responses are continuous.

4.1 Variables

Each individual i has survey measures of income, MTR and ATR across two waves of data ($w = 1, 2$). Wave 1 refers to data collected in CogEcon 2011, with reference to tax year 2010. Wave 2 refers to data collected

in CogEcon 2013, with reference to tax year 2012. The income, MTR and ATR variables with superscript $*$ are the unobserved true values. There is no superscript when the variable is observed in the data.

Income

- $\mathbf{y}_i \equiv (y_{i,1}, y_{i,2})'$ is the vector of the log of reported income for individual i .
- $\mathbf{y}_i^* \equiv (y_{i,1}^*, y_{i,2}^*)'$ is the vector of the log of true income for individual i .

Marginal tax rates (MTR)

- $\mathbf{m}_i \equiv (m_{i,1}, m_{i,2})'$ is the vector of reported marginal tax rates for individual i .
- $\mathbf{m}_i^* \equiv (m_{i,1}^*, m_{i,2}^*)'$ is the vector of true statutory marginal tax rates for individual i .

Average tax rates (ATR)

- $\mathbf{a}_i \equiv (a_{i,1}, a_{i,2})'$ is the vector of reported average tax rates for individual i .
- $\mathbf{a}_i^* \equiv (a_{i,1}^*, a_{i,2}^*)'$ is the vector of true average tax rates for individual i .

4.2 Known income tax functions

4.2.1 Statutory relationship between income and MTR

Individual i 's true statutory marginal tax rate ($m_{i,w}^*$) in wave w is a function of true taxable income ($I_{i,w}^*$), written as $m_{i,w}^* = M_w(I_{i,w}^*)$. Taxable income ($I_{i,w}^*$) is the greater of zero and adjusted gross income ($AGI_{i,w}^*$) minus deductions ($D_{i,w}^*$), which is either the standard deduction or itemized deductions, and exemptions ($E_{i,w}^*$)

$$I_{i,w}^* = \max\{AGI_{i,w}^* - D_{i,w}^* - E_{i,w}^*; 0\}. \quad (1)$$

Adjusted gross income can also be written as the AGI share of total income ($s_{i,w}^*$) times total income: $AGI_{i,w}^* = s_{i,w}^* \cdot Y_{i,w}^*$. The marginal tax rate function $M_w(\cdot)$ is defined by tax law

$$M_w(I_{i,w}^*) = \tau_j \iff \overline{I_w^{j-1}} \leq I_{i,w}^* < \overline{I_w^j} \quad (2)$$

and maps true taxable income into one of seven statutory tax rates, $\tau_j \in \{0; 10; 15; 25; 28; 33; 35\}$. The statutory marginal tax rate is τ_j when taxable income in wave w is between thresholds $\overline{I_w^{j-1}}$ and $\overline{I_w^j}$. The top threshold for bracket $j - 1$ is the bottom threshold for bracket j . These income thresholds depend on filing status (married or single), but subscripts for filing status are suppressed to simplify notation. While the statutory rates are the same in 2010 and 2012, the function $M_w(\cdot)$ changes because of adjustments to the taxable income thresholds.

4.2.2 Statutory relationship between income and tax liability

Taxpayer i 's true tax liability in wave w , $T_{i,w}^* = T_w(I_{i,w}^*)$, is a continuous, piece-wise linear deterministic function of taxable income

$$T_w(I_{i,w}^*) = \tau_j \cdot (I_{i,w}^* - \overline{I_w^{j-1}}) + C_w^j \text{ when } \overline{I_w^{j-1}} \leq I_{i,w}^* < \overline{I_w^j}. \quad (3)$$

That is, someone with taxable income between $\overline{I_w^{j-1}}$ and $\overline{I_w^j}$ is taxed at rate τ_j on all taxable income above $\overline{I_w^{j-1}}$. Everyone with marginal tax rate τ_j (and the same filing status) pays tax C_w^j on all income less than threshold $\overline{I_w^{j-1}}$, where $C_w^j \equiv \tau_1 \cdot \overline{I_w^1} + \tau_2 \cdot (\overline{I_w^2} - \overline{I_w^1}) + \dots + \tau_{j-1} \cdot (\overline{I_w^{j-1}} - \overline{I_w^{j-2}})$.

By ignoring tax credits and assuming all income is taxed as wage or salary income, the tax liability computation is an upper bound on true tax liability at a given taxable income. This is because tax credits reduce tax liability and tax-preferred investments are taxed at lower rates than wages and salary income. Therefore, my estimate of the average tax rate errors, if anything, results in a lower bound on the ATR error.

4.3 Measurement equations

There are two definitions of income relevant to this analysis. Total income (Y) refers to money received from almost any source during the specified tax year. This corresponds to the definition of income that is observed in the data. Taxable income (I) is defined as total income minus all deductions, exemptions and adjustments.⁶

Survey measures of income are inherently noisy, and there has been an extensive literature on the causes and consequences of such errors.⁷ To account for measurement error in observed income, I follow the literature (e.g., Bound and Krueger (1991); Kapteyan and Ypma (2007)) and model observed log income as $y_{i,w} = \log Y_{i,w}$, equal to true log income ($y_{i,w}^*$) plus random noise ($e_{i,w}$)

$$y_{i,w} = y_{i,w}^* + e_{i,w} \quad (4)$$

I assume true log income ($y_{i,w}^*$) across the two waves 1 and 2 has a bivariate normal distribution with wave-specific mean $\mu_{y_w^*}$ and standard deviation $\sigma_{y_w^*}$, and correlation ρ_{y^*} across the two waves: $\mathbf{y}_i^* \equiv (y_{i,1}^*, y_{i,2}^*)' \sim BVN(\mu_{y_1^*}, \mu_{y_2^*}, \sigma_{y_1^*}^2, \sigma_{y_2^*}^2, \rho_{y^*})$.

Measurement error is normally distributed random noise with mean zero, $e_{i,w} \sim N(0, \sigma_e^2)$. Assuming the income error is unbiased is consistent with existing evidence from validation studies (Bound et. al (2001)). Assuming errors are mean zero with the same standard deviation (σ_e) in the two waves, uncorrelated with

⁶More specifically, what I refer to as total income might typically be called gross income (GI). This includes money received from almost any source during the specified tax year and is the broadest income concept used for tax purposes. Adjusted Gross Income (AGI) equals gross income minus "above the line" deductions such as IRA contributions, student loan interest, and self-employed health insurance contributions. Taxable income equals AGI minus exemptions and deductions.

⁷See Bound et. al (2001) for a detailed account of measurement error in survey data.

true income and independent across waves, yields $Cov(y_1, y_2) = \rho_{y^*} \cdot \sigma_{y_1^*} \cdot \sigma_{y_2^*}$.⁸ Hence, the joint distribution of observed income in waves 1 and 2 is

$$\begin{bmatrix} y_{i,1} \\ y_{i,2} \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_{y_1^*} \\ \mu_{y_2^*} \end{bmatrix}, \begin{bmatrix} \sigma_{y_1^*}^2 + \sigma_e^2 & \rho_{y^*} \cdot \sigma_{y_1^*} \cdot \sigma_{y_2^*} \\ \rho_{y^*} \cdot \sigma_{y_1^*} \cdot \sigma_{y_2^*} & \sigma_{y_2^*}^2 + \sigma_e^2 \end{bmatrix} \right)$$

The three main assumptions are (i.) the functional form of the income distribution; (ii.) measurement error is “classical”; (iii.) measurement error is independent across waves. In my main specification, I also assume the parameters are the same in the two waves, such that $\mu_{y_w^*} = \mu_{y^*}$ and $\sigma_{y_w^*} = \sigma_{y^*}$, and assume that the cross-wave correlation of true income (ρ_{y^*}) is known (from another source).

The functional form seems reasonable given the likely sources of error in my data. Measurement error that is additive in logs is equivalent to having multiplicative error that follows a log-normal distribution. Respondents who answer the question with an exact amount often report rounded numbers (e.g., someone with income of 41,348 might say 40,000). Those who do not provide an exact amount are then prompted to select a range of income from a list of ranges. For these respondents we assign the midpoint but in fact all we know is that their income is within an interval of values; the intervals increase at higher levels of income, consistent with a lognormally distributed error term.

I also assume measurement error is “classical,” such that $Cov(e_{i,w}, y_{i,w}^*) = 0$. This assumption warrants discussion, as there is mixed evidence in previous studies about the reasonableness of this assumption. Bound and Krueger (1991), Bollinger (1998), Kapteyan and Ypma (2007), and Hudomiet (2013) all find evidence that reported income exhibits mean-reverting measurement error. People with low income overreport and those with higher income underreport. The implications of this assumption will be addressed when I interpret the results. However, Hudomiet (2013) finds that this is almost entirely at very low levels of income, and that there is no evidence of mean reversion for those with income above \$10,000.

Finally, I assume that the income measurement error is uncorrelated across waves. By making this assumption I am attributing all systematic differences between the survey and computed tax rates to errors in tax rates. This assumption would be violated if respondents, for example, have private information about itemized deductions that is not captured in the CogEcon measures of income but impacts their tax rate responses.

4.3.1 Marginal and average tax rates

Taxpayer i 's reported marginal tax rate in wave w equals the true MTR plus an error term

$$m_{i,w} = m_{i,w}^* + \varepsilon_{i,w}^m \tag{5}$$

⁸This comes from

$$Cov(y_1, y_2) = Cov(y_1^* + e_1, y_2^* + e_2) = Cov(y_1^*, y_2^*) = \rho_{y^*} \cdot \sigma_{y_1^*} \cdot \sigma_{y_2^*}$$

where true MTR ($m_{i,w}^*$) is defined in the previous section. Taxpayer i 's reported average tax rate in wave w equals the true ATR plus systematic error and stochastic noise

$$a_{i,w} = a_{i,w}^* + \varepsilon_{i,w}^a \quad (6)$$

where true ATR is true tax liability (from previous section) divided by true adjusted gross income, $a_{i,w}^* = \frac{T_{i,w}^*}{AGI_{i,w}^*}$. The errors $\varepsilon_{i,w}^a$ and $\varepsilon_{i,w}^m$ capture heterogeneity in survey reports, conditional on the true tax rates. If respondents had precise beliefs about their tax rates, then the errors $\varepsilon_{i,w}^a$ and $\varepsilon_{i,w}^m$ capture heterogeneity in misperceptions and random survey noise. If respondents do not have precise beliefs, then the reported values can be interpreted as an unbiased estimate of the mean subjective rates. Heterogeneity in survey measures includes systematic misperception, variation in reporting beliefs due to resolveable (but unresolved) uncertainty (on behalf of the respondent), and random survey noise.

There are a couple reasons why I define true ATR with adjusted gross income ($AGI_{i,w}^*$) in the denominator. First, this definition is often used by the IRS when analyzing average tax rates and by TurboTax when computing the average tax rate after filing one's return. Respondents who recalled estimates of ATR from their tax software or tax professionals are likely reporting a noisy measure of average tax rate defined using AGI in the denominator. A more intuitive definition, and one more consistent with the wording of our survey question, uses a broader definition of income. As a robustness check I present estimates using true ATR defined as tax liability divided by total income, $a_{i,w}^* = \frac{T_{i,w}^*}{Y_{i,w}^*}$. This definition of true ATR is the preferred specification if respondents answer the question by dividing their perceived tax liability by perceived total income. As is expected, the mean bias in ATR is larger when using this broader measure of total income ($Y_{i,w}^*$). The mean bias in perceived average tax rates is larger when using gross income or another measure with a broader tax base.

Modeling the marginal tax rate with additive error in the true rate could be contrasted with modeling the subjective rate in terms of the discrete statutory MTR. If people knew the tax table then I could model the tax response in terms of an underlying latent variable determining the categorical tax rate. However, only about half of respondents answer using one of the statutory marginal tax rates. Binning the tax responses ignores information contained in responses that are in between statutory tax rates. On the other hand, the presumed data generating process will not give the spikes in the data at the statutory MTRs. I address this issue in a companion paper (Gideon, 2014), in which respondents come from mixture of two data generating processes. The first type report MTR with additive error, as in the model above. The second type report statutory MTR associated with their perceived level of taxable income.

While respondents were asked for values between 0 and 100 in CogEcon 2011, they were not restricted like this in CogEcon 2013. Nevertheless, the distributions look similar in the two waves and the restricted domain in 2011 does not appear to have impacted responses. The functional form assumptions are used to tame the data and are not instrumental to interpreting the parameter estimates. Therefore, the baseline

specification ignores censoring at 0 and 100 in CogEcon 2011.

To close out the structural model I need to make distributional assumptions about MTR errors $\varepsilon_{i,w}^m$ and ATR errors $\varepsilon_{i,w}^a$. Define the vector of tax rates responses as $\mathbf{r}_i = (a_{i,1}, a_{i,2}, m_{i,1}, m_{i,2})'$. I assume that the marginal and average tax rates, conditional on true income, have a multivariate normal distribution

$$\mathbf{r}_i \mid \mathbf{y}_i^* \sim N(\mathbf{r}_i^* + b_{\mathbf{r}}, \Sigma_{\mathbf{r}})$$

over the average and marginal tax rate responses in both waves (four dimensional distribution). The vector $b_{\mathbf{r}} = (b_{a_1}, b_{a_2}, b_{m_1}, b_{m_2})'$ captures mean biases in tax rate perceptions, where $b_{a_w} = E(a_{i,w} - a_{i,w}^*)$ and $b_{m_w} = E(m_{i,w} - m_{i,w}^*)$ in each wave w , and

$$\Sigma_{\mathbf{r}} = \begin{pmatrix} \Sigma_{aa} & \Sigma_{am} \\ \Sigma_{ma} & \Sigma_{mm} \end{pmatrix} = \begin{pmatrix} \sigma_{a_1}^2 & \rho_a \sigma_{a_1} \sigma_{a_2} & \rho_{am} \sigma_{a_1} \sigma_{m_1} & 0 \\ \rho_a \sigma_{a_1} \sigma_{a_2} & \sigma_{a_2}^2 & 0 & \rho_{am} \sigma_{a_2} \sigma_{m_2} \\ \rho_{am} \sigma_{a_1} \sigma_{m_1} & 0 & \sigma_{m_1}^2 & \rho_m \sigma_{m_1} \sigma_{m_2} \\ 0 & \rho_{am} \sigma_{a_2} \sigma_{m_2} & \rho_m \sigma_{m_1} \sigma_{m_2} & \sigma_{m_2}^2 \end{pmatrix}$$

is the variance-covariance matrix. I let σ_{a_w} represent the standard deviation of the ATR error (ε_w^a) in wave w , and σ_{m_w} is the standard deviation of the MTR error (ε_w^m) in wave w . The parameter $\rho_a = Corr(\varepsilon_1^a, \varepsilon_2^a)$ is the correlation of ATR errors across waves, $\rho_m = Corr(\varepsilon_1^m, \varepsilon_2^m)$ is the correlation of MTR errors across waves, and $\rho_{am} = Corr(\varepsilon_1^m, \varepsilon_1^a) = Corr(\varepsilon_2^m, \varepsilon_2^a)$ is the correlation of ATR errors and MTR errors within the same wave. The values of ρ_a , ρ_m , and ρ_{am} are between -1 and 1. Correlation of ATR and MTR errors across different waves is set to zero: $Corr(\varepsilon_2^m, \varepsilon_1^a) = 0$. In the baseline specification I assume the mean and variance are the same across waves, such that $b_a = b_{a_w}$ and $\sigma_a = \sigma_{a_w}$ for ATR errors, and $b_m = b_{m_w}$ and $\sigma_m = \sigma_{m_w}$ for MTR errors. To summarize, the ATR and MTR errors are potentially correlated within the same wave, that ATR errors could be correlated across waves, MTR errors could be correlated across waves, but the ATR errors in one wave are uncorrelated with the MTR errors from the other wave.

4.4 Likelihood function

Let Ω denote the set of observed data, which is a vector $\Omega_i = (\mathbf{y}_i, \mathbf{r}_i)'$ of reported income and tax rates for individuals $i = 1, \dots, N$. The likelihood $L(\theta \mid \Omega)$ that the distributional parameters for the specified model are θ given data Ω is proportional to the probability $\Pr(\Omega \mid \theta)$ of observing Ω given the specified model and parameters θ :

$$L(\theta \mid \Omega) \propto \Pr(\Omega \mid \theta) = \prod_{i=1}^N \Pr(\Omega_i \mid \theta) \quad (7)$$

The log likelihood function for the model is then given by

$$\text{LL}(\theta) = \sum_{i=1}^N \ln \{\text{Pr}(\Omega_i | \theta)\} \quad (8)$$

Using Bayes law, $\text{Pr}(\Omega_i | \theta)$ equals the probability of reporting income \mathbf{y}_i times the probability of reporting tax rates \mathbf{r}_i conditional on having reported income \mathbf{y}_i

$$\text{Pr}(\Omega_i | \theta) = \text{Pr}(\mathbf{r}_i | \mathbf{y}_i, \theta) \cdot \text{Pr}(\mathbf{y}_i | \theta) \quad (9)$$

The term $\text{Pr}(\mathbf{y}_i | \theta)$ is simply the density $f(\mathbf{y}_i | \theta)$. To calculate $\text{Pr}(\mathbf{r}_i | \mathbf{y}_i)$, I need to integrate over the two-dimensional income measurement error distribution because latent true income is not observed in the data. Each realization of income measurement error vector \mathbf{e} (conditional on observed income \mathbf{y}_i) pins down true income \mathbf{y}_i^* , since $\mathbf{e} = \mathbf{y}_i - \mathbf{y}_i^*$, which in turn determines true average tax rates \mathbf{a}_i^* and marginal tax rates \mathbf{m}_i^* .

$$\begin{aligned} \text{Pr}(\mathbf{r}_i | \mathbf{y}_i) &= \int \int f(\mathbf{r}_i | \mathbf{e}_i, \mathbf{y}_i) \cdot f(\mathbf{e}_i | \mathbf{y}_i) \, d\mathbf{e} \\ &= \int \int f(\mathbf{r}_i | \mathbf{y}_i^* = \mathbf{y}_i - \mathbf{e}_i) \cdot f(\mathbf{e}_i | \mathbf{y}_i) \, d\mathbf{e} \end{aligned}$$

This integral for $\text{Pr}(\mathbf{r}_i | \mathbf{y}_i)$ does not have a simple closed-form representation due to the complicated non-linear functions that are used to map the true income into tax rates, which are needed to integrate over the reported tax rates. Therefore, the likelihood function does not have a tractable closed-form analytical expression.

I use Maximum Simulated Likelihood (MSL) to estimate the parameters of the model. The major complication in evaluating the likelihood function arises from the fact that true income is not observed. Estimating the parameters of the model by maximum likelihood involves integrating over the distribution of these unobserved income errors. This problem is solved by simulating the likelihood function. The simulated log-likelihood function is then given by

$$\text{SLL}(\theta) = \sum_{i=1}^N \ln \check{L}_i(\theta) \quad (10)$$

The contribution of each individual i is $\check{L}_i(\theta)$, which is a simulated approximation to $L_i(\theta)$, derived as

$$\check{L}_i(\theta) = \frac{1}{K} \sum_{k=1}^K L_i^k(\theta) \quad (11)$$

where the average is over the likelihood evaluated at each simulation draw

$$L_i^k(\theta) = \text{Pr}(\mathbf{r}_i | \mathbf{y}_i, e_{i(k)}) \cdot f(\mathbf{y}_i) \quad (12)$$

and K is the number of pseudorandom draws of the vector of errors $e_{i(k)}$. The algorithm involves simulating a distribution of income errors for each respondent. The individual's likelihood contribution is computed for each set of income errors, and density of the implied tax errors are averaged over the K values to obtain the simulated likelihood contribution.

4.5 Identification

Given my statistical model for observable data Ω given a parameter vector θ expressed via the probability function $\Pr(\Omega | \theta)$, the identification problem boils down to the following: do differences between survey and computed tax rates arise because of misperceptions or because measurement error in income induces errors in the computed rate?

The model is said to be theoretically point identifiable if there is a unique set of parameter estimates $\theta = (\mu_{y^*}, \sigma_{y^*}, \rho_{y^*}, \sigma_e, b_a, b_m, \sigma_a, \sigma_m, \rho_a, \rho_m, \rho_{am})$ given the sample moments of the data. The parameter vector includes three means, four standard deviations and four correlations. The sample moments include 3 means (3 in each wave), 3 standard deviations (3 in each wave). As for the correlations, within-wave there are 3 correlations.

Identification comes from having multiple indicators and repeated measures. Reported income, marginal tax rate and average tax rate provide three indicators of latent taxable income, and these variables are observed across two waves of data. Having a known functional relationship between true taxable income and true tax rates separately identifies the distribution of income measurement error and tax rate biases. Intuitively, identification can be broken into two steps, corresponding to parameters governing the income data generating process and parameters characterizing the tax rate errors.

The income parameters include those governing true income (μ_{y^*} , σ_{y^*} and ρ_{y^*}) and the standard deviation of income measurement error (σ_e). When income measurement error is mean zero, the mean of true log income (μ_{y^*}) is identified by the mean of observed log income (μ_y), since $\mu_{y^*} = E(y^*) = \mu_y$. Because measurement error in income is assumed to be uncorrelated with true income and is uncorrelated across waves, two relationships hold:

1. $\sigma_y^2 = \sigma_{y^*}^2 + \sigma_e^2$: variance of observed income is equal to the variance of true income plus the variance of income measurement error.
2. $\rho_{y^*} = \rho_y \cdot \frac{\sigma_{y^*}^2 + \sigma_e^2}{\sigma_{y^*}^2}$, which says that correlation of true income equals the correlation of observed income times one over the reliability ratio $\frac{\sigma_{y^*}^2}{\sigma_{y^*}^2 + \sigma_e^2}$.

Hence, once I identify the standard deviation of true income (σ_{y^*}), the standard deviation of income measurement error (σ_e) is identified by the standard deviation of observed income (σ_y). The correlation of true income across waves (ρ_{y^*}) is then pinned down by the correlation of observed income (ρ_y). Alternatively, identifying the correlation of true income (ρ_{y^*}) is enough to identify the standard deviation of income mea-

surement error (σ_e). Therefore, identification of σ_e requires knowledge about the correlation of true income (ρ_{y^*}) or the standard deviation of true income (σ_{y^*}).

The tax rate error parameters $b_a, b_m, \sigma_a, \sigma_m, \rho_a, \rho_m$ and ρ_{am} are identified once the income parameters are known. The bias terms b_a and b_m are identified from the observed tax rates and income, and mean zero income measurement error. This is a normalization to make errors have mean zero. The assumption that the tax rate errors are conditionally independent of income implies that $Var(m_{i,w}) = \sigma_{m^*}^2 + \sigma_m^2$, or that the variance of observed MTR is equal to the variance of true MTR plus the variance of the MTR errors. The variance of true MTR ($\sigma_{m^*}^2$) is a deterministic function of the true income distribution (characterized by μ_{y^*} and σ_{y^*}) and the known functional relationship between true income and true tax rates. The same holds for identifying the ATR errors.

Finally, identification of the correlation of ATR and MTR errors within wave (ρ_{am}), the correlation of the ATR error across waves (ρ_a), and the correlation of MTR errors across waves (ρ_m) exploits the known correlation of true income (ρ_{y^*}) and the non-linear relationship between true MTR and true ATR across the income distribution.

It is known that non-linear functions can greatly aid identification in measurement error models (Carroll et al., 2006, pp. 184). In this context, there are two non-linear relationships that facilitate identification. First, as true income increases, both MTR and ATR increase, but at different rates. This creates variation in the difference between true MTR and true ATR. Another important mechanical source of identification comes from variability in how reliable reported income is depending on its location within tax bracket. When income falls toward the middle of a tax bracket, the true rate is known with greater likelihood than when income is close to a income cut-off of the tax bracket. I know the true MTR whenever adjusted gross income is measured without error and deductions perfectly observed. This helps separately identify the correlation of ATR and MTR errors within wave (ρ_{am}), the correlation of the ATR error across waves (ρ_a), and the correlation of MTR errors across waves (ρ_m).

Therefore, identifying σ_e requires knowledge about the correlation of true income (ρ_{y^*}) or the standard deviation of true income (σ_{y^*}). And knowing these three parameters is sufficient to identify the tax rate error parameters $\sigma_a, \sigma_m, \rho_a, \rho_m$ and ρ_{am} .

The standard deviation of true income is also identified from the correlation of observed income and tax rates within waves. By assumption, the shared component of this correlation is due to true income. Combining income and tax rates helps identify the distribution of true income.

As described above, identification depends on a few important assumptions. First, tax errors are uncorrelated with true income and, therefore, uncorrelated with the true tax rate. Second, tax and income errors are uncorrelated. This assumption is critical, and relies on the fact that errors reflect random survey noise and systematic differences between perceived and true tax rates. If people who overreport income also overreport taxes then the pair of responses would appear consistent with one another and having a second measure is no longer helpful.

4.6 Estimation

Estimation of $\theta = (\mu_{y^*}, b_a, b_m, \sigma_{y^*}, \sigma_e, \sigma_a, \sigma_m, \rho_a, \rho_m, \rho_{am})$ is done using the method of maximum simulated likelihood. In practice, I use between $K=50$ and 200 draws per individual to simulate the likelihood in the baseline specification. Quasi-random Halton sequence draws, rather than random draws, are used to simulate the likelihood because of the documented superior performance of quasi-random Halton draws in the simulation of integrals (e.g., Train 1999, Bhat 2001). Halton sequences are used to construct draws over the two-dimensional income measurement error. The individual's likelihood contribution is computed for each set of income errors, and density of the implied tax errors are averaged over the K values to obtain the simulated likelihood contribution. A detailed description of the simulation algorithm and likelihood evaluation are in Appendix C.1.

In the main specification, I calibrate the correlation of true log income (ρ_{y^*}) across years 2010 and 2012 to reduce computational issues associated with simulating the likelihood function. I set $\rho_{y^*} = 0.9$, which coincides with estimates of the second-year autocorrelation of log earnings from other studies. Haider and Solon (2006) estimate an average first-order (one year) autocorrelation of 0.89 and average second-order autocorrelation of 0.82 using Social Security earnings data for a sample from the cohort born in 1931-1933 and studied between the ages of about 43 and 52. They show that the correlation for an alternative time frame, 1980-1991 rather than 1975-1984, results in slightly higher larger estimates, with first-order autocorrelation of 0.91 and second-order autocorrelation of 0.85. Parsons (1978) studies white male wage and salary earners and estimates the autocorrelation of earnings for different age categories (45-54 and 55-64) and three categories of educational achievement. For the 55-64 age category—the most relevant for CogEcon sample demographics—the first and second year autocorrelation are above 0.925 across education categories. and 0.970 for the lowest and highest education categories. The autocorrelation across years two and three is significantly different across the education groups, with autocorrelation of 0.674 for education=12, 0.863 for education=13-15 and 0.968 for education=16. The CogEcon sample of older and better educated workers suggests that 0.9 is a conservative estimate of the correlation of true income. However, the results are robust to both alternative assumptions about this correlation as well as estimating the correlation directly.

5 Parameter Estimates

Table 5 presents the estimated structural parameter values ($\theta = (\mu_{y^*}, b_a, b_m, \sigma_{y^*}, \sigma_e, \sigma_a, \sigma_m, \rho_a, \rho_m, \rho_{am})$) and standard errors for the full model of average tax rates, marginal tax rates and income (Equations 4, 5 and 6), and alternative specifications. This includes three means, four standard deviations and three correlations. Clustered robust standard errors are constructed from the asymptotically robust variance-covariance matrix using the outer product of the score to approximate the Hessian, clustering at the household level.

The first column of Table 5 reports the estimates for the full model. All of the parameters are estimated quite precisely. The mean of the ATR bias (b_a) is almost 6 percentage points. A typical person in CogE-

con pays 11% of their income in federal income taxes, but perceives it as 17%. There is also substantial heterogeneity in these beliefs. The estimated standard deviation (σ_a) is 8.8, meaning there is substantial unobserved heterogeneity even accounting for measurement error in income. With the mean perceived rate of 17%, two standard deviations would go from around 0 to 30, which spans the range of true ATR that could reasonably be observed. Having the two waves of data allows me to estimate the correlation of ATR errors across years (ρ_a), which is estimated to be 0.25.

The estimated parameters for the MTR errors show marked differences from the parameters for the ATR errors. There is no systematic difference between the survey and true marginal tax rates, as the mean of the MTR bias (b_m) is statistically indistinguishable from zero. The estimated standard deviation (σ_m) is 9.8, which is over one percentage point larger than the standard deviation of the ATR errors. Finally, the estimated correlation of MTR errors across years (ρ_m) is around 0.13, which is smaller in magnitude than the correlation of ATR errors but still significantly different from zero.

Estimates of the standard deviation of the ATR errors (σ_a) and MTR errors (σ_m) are economically large and significant. Heterogeneity in reported rates comes from unexplained variation in systematic errors, yet also reflects various measurement problems. In terms of the reported tax rates, there are both small scale errors such as approximating, guessing the truth, rounding despite being well-informed, and larger random noise. The positive correlation suggests that the large variation is not simply a failure to properly measure perceptions of tax rates. For example, a person who has a large itemized deduction could have lower taxable income and lower MTR and ATR than is accounted for in the model.

These correlations of the MTR errors across waves (ρ_m) and ATR errors across waves (ρ_a) are small compared to the within-year correlation of MTR and ATR errors (ρ_{am}) of over 0.5, which is both statistically and substantively significant. Someone who overestimates their ATR is more likely to overestimate their MTR. However, I cannot rule out the possibility that the positive correlation is the result of having accurate measures of the survey tax rates and errors in the computed rates. Regardless of the interpretation, the large positive estimate of ρ_{am} provides strong evidence that tax rate errors are not purely random noise.

Though not my primary focus, the estimated parameters on the income distribution are interesting in their own right, as this is the first paper to identify income measurement error from reported tax rates. True log income y_i^* has an estimated mean of 11.1, corresponding to median income \$66,171. The standard deviation of true log income is $\widehat{\sigma}_{y^*} = 0.78$, the standard deviation of income measurement error of $\widehat{\sigma}_e = 0.23$. These estimates imply a reliability ratio of 0.92, which is defined as the (estimated) variance of the latent true income divided by the variance of the noisy measure.⁹ Reliability ratios provide a way to compare the extent of measurement error across studies, wherein a higher reliability measure is associated with less survey noise. This estimate is slightly higher than what is typically found in validation studies of income.

9

$$\frac{\widehat{\sigma}_{y^*}^2}{(\widehat{\sigma}_e^2 + \widehat{\sigma}_{y^*}^2)} = \frac{(0.78)^2}{(0.23)^2 + (0.78)^2} = 0.92$$

Bound and Krueger (1991) find slightly higher reliability of CPS earnings, in the low 0.80 and 0.90 ranges for men and women, respectively. In contrast, Duncan and Hill (1985) estimate a ratio of 0.76 when comparing reported earnings from employees at a particular company with the company's payroll records.

Table 5: Maximum Likelihood estimates of measurement model parameters

	(1)	(2)	(3)	(4)	(5)
ATR errors (a-a*)					
Mean (b_a)	5.719 (0.436)	5.955 (0.474)		5.890 (0.442)	
Standard Deviation (σ_a)	8.798 (0.384)	9.758 (0.680)		8.993 (0.469)	
Correlation across waves (ρ_a)	0.247 (0.044)	0.439 (0.076)		0.406 (0.065)	
MTR errors (m-m*)					
Mean (b_m)	0.682 (0.459)		0.038 (0.460)	0.715 (0.471)	0.039 (0.454)
Standard Deviation (σ_m)	9.767 (0.454)		9.814 (0.532)	9.759 (0.565)	9.796 (0.752)
Correlation across waves (ρ_m)	0.128 (0.046)		0.264 (0.060)	0.313 (0.064)	0.264 (0.059)
Correlation of ATR & MTR errors (ρ_{am})	0.537 (0.045)				
Income					
Mean (μ_{y^*})	11.100 (0.045)	11.096 (0.045)	11.096 (0.045)	11.096 (0.041)	11.096 (0.045)
Standard Deviation (σ_{y^*})	0.783 (0.038)	0.778 (0.037)	0.781 (0.038)	0.760 (0.035)	0.779 (0.056)
Error Standard Deviation (σ_e)	0.230 (0.041)	0.248 (0.044)	0.244 (0.041)	0.273 (0.040)	0.264 (0.059)
Correlation across waves (ρ_{y^*})	0.9	0.9	0.9	0.9	0.906 (0.107)
Log likelihood	-2495.29	-1602.6	-1647.2	-2592.9	-1647.2
Parameters	10	6	6	9	7

Notes: The parameter estimates from the baseline model and alternative specifications are reported above. Column (1) is the baseline three equation model of MTR, ATR and income. Column (2) is the two equation model of ATR and income and column (3) is the two equation model of MTR and income. Column (4) restricts the correlation of ATR and MTR errors to be zero. In columns (1) thru (4), I calibrate the correlation of true income at $\rho_{y^*} = 0.9$. Finally, column (5) is the two equation model of MTR and income, but relaxing the assumption on the correlation of true income across waves. Estimation of the model in (1), (2) and (4) is done using the method of maximum simulated likelihood. In practice, I use 50 Quasi-random Halton sequence draws over two dimensions per individual to simulate the likelihood. The rates are scaled by 10 for the estimation and then rescaled back up. The value of the log likelihood function are not directly comparable across the models.

5.1 Specification checks and further analyses

These estimated parameters from the full model of ATR, MTR and income are robust to alternative specifications. First, columns (2) and (3) of Table 5 present estimated parameters of a two-equation model of ATR and income (Equations 4 and 6) and a two-equation model of MTR and income (Equations 4 and 5), respectively. The estimates are similar to those in the full model, but with larger standard errors. There are also some noticeable differences between the estimates. The correlation of ATR errors across waves (ρ_a) and MTR errors across waves (ρ_m) are twice as large, and the standard deviation of the ATR error (σ_a) is slightly larger ($\hat{\sigma}_a = 9.758$ versus $\hat{\sigma}_a = 8.798$). It is important to note that I estimate the two-equation model of MTR and income without simulating the likelihood function because I do not need to integrate out the income measurement error. The consistency of the estimates across models provides important evidence that the simulated likelihood estimation strategy is not impacting the results.

Column (4) in Table 5 presents estimates of the full model (Equations 4, 5 and 6) but assuming the MTR and ATR errors are uncorrelated ($\rho_{am} = 0$). As expected, the estimated correlation of ATR errors across waves (ρ_a) is similar to in a model of only ATR and income, and the estimated correlation of MTR errors across waves (ρ_m) is similar to what it is when estimated in a model of only MTR and income. Marginal and average tax rates are mechanically linked, so it's important to estimate the equations jointly to account for their dependence. Once ATR errors are included in the model the correlation of the MTR errors across waves is close to zero. These correlations highlight the importance of jointly modeling ATR and MTR.

Column (5) in Table 5 presents estimates for the two-equation MTR model (Equations 4 and 5) when relaxing the assumption that the correlation of true income is $\rho_{y^*} = 0.9$. The estimate of this correlation is slightly larger, $\hat{\rho}_{y^*} = 0.91$, but the assumed value is well within the 95% confidence interval for the estimated correlation. Under reasonable assumptions about the cross-wave correlation of true income, I get estimates of the income measurement error that are similar to what I find using marginal tax rate responses and are consistent with what has been previously found in the literature. Since my goal is to account for measurement error in income when comparing reported tax rates with measures computed using reported income, the exact estimate of income measurement error is of lesser importance. These results suggest that assuming $\rho_{y^*} = 0.9$ is a reasonable assumption in this context.

This exercise also provides indirect evidence that simulating the likelihood function for the full three-equation model in fact limits what I can identify using my current set of numerical routines. As discussed in Appendix C, estimation of the three equation model introduces computation issues that make it sometimes difficult to determine convergence. It raises questions about identification of the model. However, the results from this exercise provide suggestive evidence that estimation problems are computational in nature and are stemming from simulating the likelihood function rather than problems with identification.

Parameter estimates are generally robust to varying assumptions about the correlation of true income. A large correlation of true income is needed to justify the substantial variation in tax perceptions conditional on reported income. This is because the a higher correlation of income means that variation in the observed

values gets attributed to measurement error rather than variation in income. The higher the correlation of true income, the larger the estimate of the income measurement error. Intuitively, if the true income is highly correlated across waves then variation in observed income across waves must be explained by random noise (measurement error) because true income acts as a fixed component and does not vary much.

Table 6 presents the estimated structural parameter values and standard errors when allowing the mean and standard deviation of true income and tax perceptions to differ across waves. Column (1) is repeated from Table 5 for reference. In column (2), the distribution of true income and the tax rate errors can be different across waves. The estimation routine appears unstable, so I calibrate the income measurement error at $\sigma_e = 0.3$. The same patterns show up in the individual waves and the estimated parameters when pooling the observations across waves is in between the estimated parameters for the individual waves. While the standard deviation of ATR errors and MTR errors are both around 9.5 percentage points in 2010, but diverge in 2012. In particular, the standard deviation of the ATR errors falls from 9.3 to 8.3 percentage points, while the standard deviation of the MTR errors increases from 9.5 to 10.5 percentage points the standard deviations diverge in 2012. The mean ATR error hardly changes, while the mean MTR error increases from -0.01 to 1.5 . These changes are noteworthy in light of differences in the question wording across waves.

Table 6: Maximum Likelihood estimates of measurement model parameters

	(1)		(2)		(3)	
	Coef.	SE	Coef.	SE	Coef.	SE
ATR errors (a-a*)						
Mean (b_a)						
Pooled	5.719	0.436				
Wave 1			5.904	0.535	6.153	0.577
Wave 2			5.590	0.488	5.742	0.538
Standard Deviation (σ_a)						
Pooled	8.798	0.384				
Wave 1			9.165	0.585	10.103	0.786
Wave 2			8.210	0.355	9.316	0.908
Correlation across waves (ρ_a)	0.247	0.044	0.257	0.045	0.248	0.054
MTR errors (m-m*)						
Mean (b_m)						
Pooled	0.682	0.459				
Wave 1			0.041	0.568	0.002	0.588
Wave 2			1.455	0.6617	1.313	0.627
Standard Deviation (σ_m)						
Pooled	9.767	0.454				
Wave 1			9.280	0.494	10.219	0.468
Wave 2			10.390	0.965	10.770	0.822
Correlation across waves (ρ_m)	0.128	0.046	0.100	0.041	0.120	0.049
Corr. MTR & ATR errors (ρ_{am})	0.537	0.045	0.525	0.047	0.475	0.049
Income						
Mean (μ_{y^*})						
Pooled	11.100	0.045				
Wave 1			11.106	0.047	11.106	0.045
Wave 2			11.086	0.049	11.086	0.049
Standard Deviation (σ_{y^*})						
Pooled	0.783	0.038				
Wave 1			0.723	0.034	0.936	0.052
Wave 2			0.762	0.043	0.980	0.064
Error Standard Deviation (σ_e)	0.230	0.041	0.3			
Observations		348		348		348
Log likelihood		-2495.29		-2491.76		-2539.25

Notes: Column (1) displays the baseline estimates to aid comparison. Column (2) displays estimates from a model in which I calibrate the standard deviation of income measurement error at $\sigma_e = 0.3$ and calibrate the correlation of true income at $\rho_{y^*} = 0.9$. True income, MTR errors and ATR errors can have different distributions in wave 1 and 2. In (3), income is assumed to be without error. For the sake of comparison across the models, I still calibrate the correlation of true income at $\rho_{y^*} = 0.9$. The value of the log likelihood function are not directly comparable across the models. Estimates in Column (2) suggest that it is reasonable to assume the distributions are the same across waves. And comparing estimates in columns (2) and (3) shows the impact of income measurement error on the estimated tax rate errors.

Ignoring survey response error overstates the heterogeneity in tax beliefs. Column (3) imposes the restriction that income is measured without error ($\sigma_e = 0$). Comparing columns (2) and (3) shows how income measurement error affects other estimated parameters. As noted, this induces downward bias in the estimated mean error in tax beliefs. When incorporating income measurement error there appears to be a significant decline in the standard deviation of the tax rate errors. Without accounting for measurement error in income, the standard deviation of the difference between survey and reported rates hover around 10 percentage points, the exact value depending on the wave and whether it is ATR or MTR. These standard deviations are reduced to around 9 percentage points once accounting for measurement error, which is approximately 1 percentage point smaller, or 10 percent of the overall variation. But the decline is not economically significant, as the magnitude of these declines is miniscule compared to the overall level of these rates. While the fully structural model requires strong assumptions, it allows me to rigorously account for income measurement error and establish that the descriptive patterns cannot be explained by errors in the computed rates. This modest impact of income measurement error reinforces the conclusion that there are systematic biases and substantial heterogeneity in perceptions of tax rates.

6 Systematic heterogeneity in tax perceptions

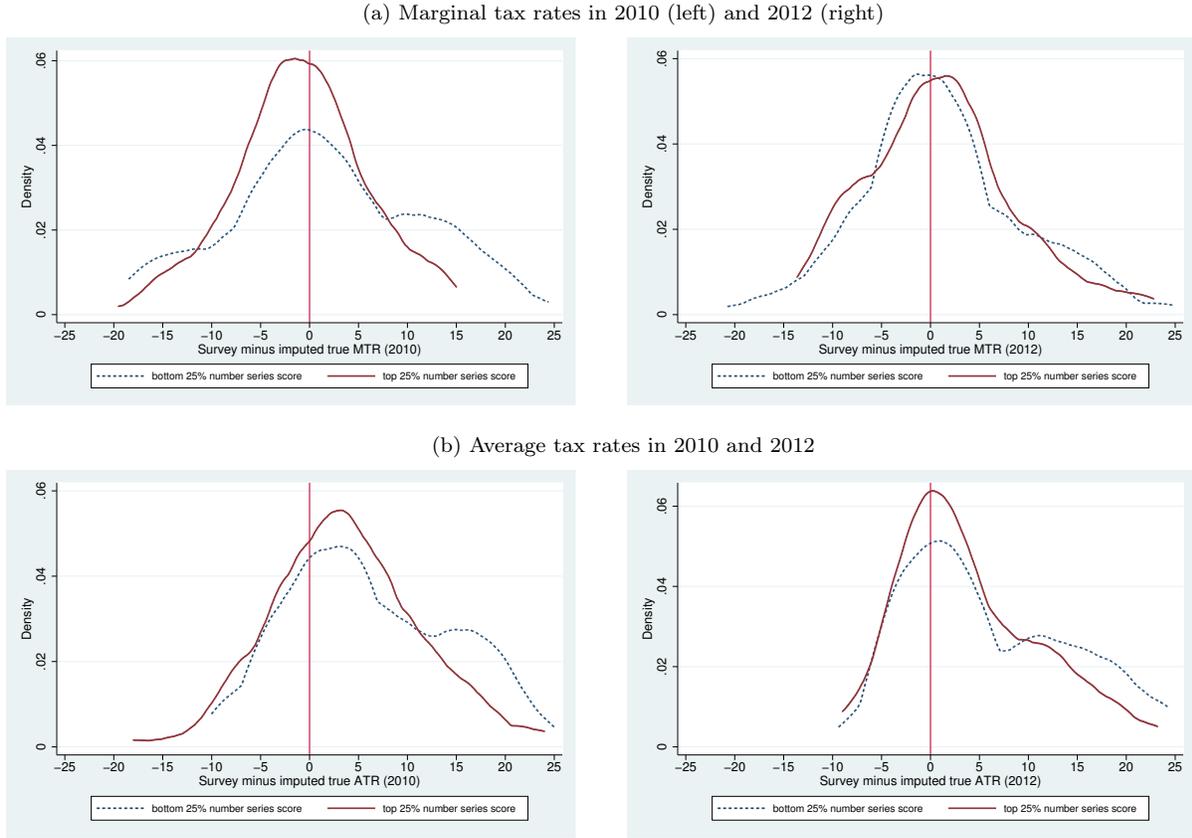
In this section I explore how misperceptions vary by observable characteristics. I focus on how the systematic biases are related to demographic characteristics, cognitive ability, general financial sophistication and the use of paid tax preparers.¹⁰

To motivate this analysis, consider Figures 4 and 5 comparing the distribution of misperceptions across sub-groups in the sample. Figure 4 compares the distribution of tax rate misperceptions for respondents in the bottom and top quarters of cognitive ability (number series score). The graphs are of the kernel density estimates of the difference between reported rates and the imputed true rates. Imputations come from the conditional expectations $\widehat{m}_{i,w}^* = \widehat{E} [m_{i,w}^* | \widehat{\theta}, \Omega_{i,w}]$ and $\widehat{a}_{i,w}^* = \widehat{E} [a_{i,w}^* | \widehat{\theta}, \Omega_{i,w}]$, which will be described more fully in Section 7. Respondents in the bottom quarter of cognitive ability (number series score) display much greater variation in their errors than do respondents in the top quarter. In particular, the right-tail of the distribution is larger for the low cognitive ability group.

Figure 5 compares the distribution of tax rate errors for respondents who used paid tax preparers and those who did not. There is a similar pattern when comparing respondents who used tax preparers and those who did not, with the later exhibiting less dispersion and smaller upward bias. However, respondents who use tax preparers are less likely to answer the question about tax rates, suggesting self-selection might make these numbers appear more accurate than one might expect overall.

¹⁰Number series is a measure of fluid intelligence, or IQ, whereas the financial sophistication score is interpreted as a measure of crystallized intelligence.

Figure 4: Distribution of tax rate errors: top 25% versus bottom 25% of number series score



Notes: These figures show the distribution of the difference between survey tax rates minus imputed true rates. I use the Epanechnikov kernel and the plugin estimator of the asymptotically optimal constant bandwidth (Fan and Gijbels (1996)).

6.1 Model specification

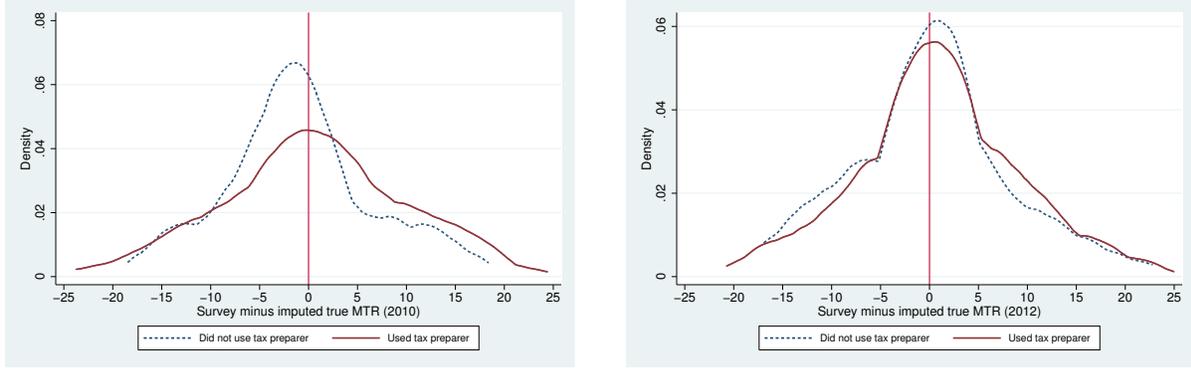
I model systematic heterogeneity by specifying the mean log income and mean tax rate misperceptions as linear indices. This is defined as $b_x^m = \mathbf{x}_i \cdot \beta^m$ for marginal tax rate misperceptions, $b_x^a = \mathbf{x}_i \cdot \beta^a$ for average tax rate misperceptions, and $\mu_{y^*} = \mathbf{x}_i \cdot \beta^y$ for true (latent) income. Estimates of parameters in β^m and β^a tell us how tax rate misperceptions vary by other observable characteristics.

The likelihood function is the same as in Section 4.4 but is now conditional on covariates \mathbf{x}_i . Let Ω denote the set of observed data, which is a vector $\Omega_i = (\mathbf{y}_i, \mathbf{r}_i, \mathbf{x}_i)'$ of reported income and tax rates for individuals $i = 1, \dots, N$.

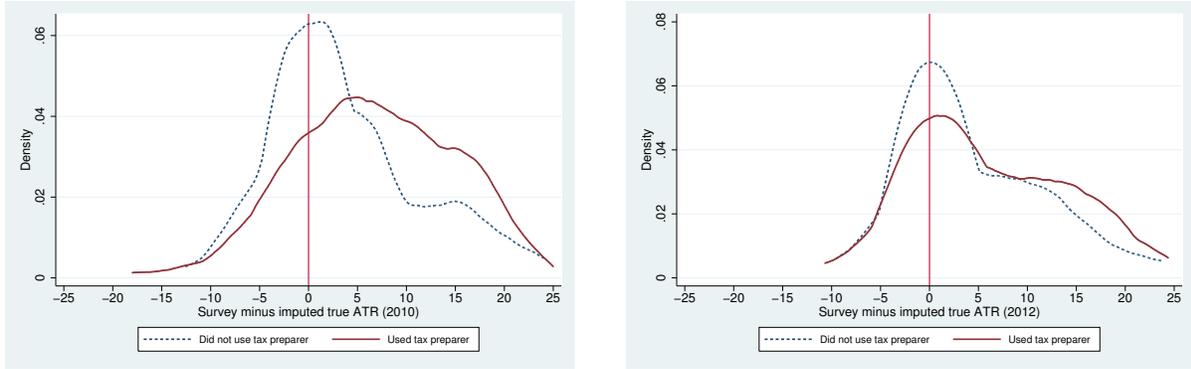
I include basic demographic variables to condition on marital status, age, gender, education and race. Particularly interesting are the measures of ability and the use of professional tax preparers. The number series, verbal analogies and financial sophistication scores are all standardized and the coefficients should be interpreted in terms of a one standard deviation change in the scores. I also include log wealth and an indicator variable for living in a state that has a state income tax.

Figure 5: Distribution of tax rate errors: Used tax preparer versus no tax preparer

(a) Marginal tax rates in 2010 (left) and 2012 (right)



(b) Average tax rates in 2010 and 2012



Notes: These figures show the distribution of the difference between survey tax rates minus imputed true rates. I use the Epanechnikov kernel and the plugin estimator of the asymptotically optimal constant bandwidth (Fan and Gijbels (1996)).

6.2 Results

Table 7 presents estimated parameters and standard errors for the full structural model incorporating systematic heterogeneity in true income and in tax rate misperceptions. These values are the re-scaled (multiplied by ten) values of the estimates and standard errors. The main take-away from analyzing heterogeneity is the general cluelessness about marginal tax rates. Respondents who use paid tax preparers report over 3 percentage points higher average tax rates and over 2 higher marginal tax rates, holding their true tax rates fixed. Number series and financial sophistication scores are associated with ATR misperceptions; I gain something from using ability measures above and beyond what is explained by education.

Table 7: Maximum Likelihood estimates of systematic heterogeneity

Parameters	Latent variable		
	log(true income)	ATR-ATR*	MTR-MTR*
Used paid tax preparer (in 2011)	0.005 (0.061)	3.269 (0.958)	2.266 (1.007)
Number Series score	0.060 (0.040)	-1.463 (0.653)	-0.238 (0.780)
Verbal Analogies score	0.058 (0.037)	0.156 (0.589)	0.127 (0.655)
Financial sophistication score	0.062 (0.038)	-0.828 (0.603)	0.855 (0.680)
State income tax (=1 if yes)	-0.061 (0.073)	0.346 (1.086)	0.877 (1.136)
Education : 13-16 years	0.129 (0.067)	1.158 (1.226)	-0.490 (1.304)
Education: >16 years	0.403 (0.084)	0.517 (1.522)	-2.236 (1.661)
Occupation deals with finance/investment	-0.204 (0.110)	-1.755 (1.424)	-1.319 (1.699)
Occupation deals with budgets	0.052 (0.088)	1.526 (1.086)	1.203 (1.291)
Male (=1 if yes)	0.083 (0.056)	-0.171 (0.848)	0.007 (0.967)
Nonwhite (=1 if yes)	0.158 (0.109)	2.169 (1.873)	2.545 (1.972)
Age/10 (in 2011)	-0.016 (0.003)	-0.115 (0.050)	0.083 (0.062)
Married (=1 if yes in 2011)	0.252 (0.077)	1.366 (1.140)	1.100 (1.264)
log(wealth)	0.264 (0.044)	-0.624 (0.430)	-0.717 (0.540)
Constant	8.347 (0.588)	18.148 (6.140)	2.401 (7.644)
Standard deviation of the errors	0.311 (0.031)	8.476 (0.394)	9.507 (0.479)
Standard deviation of true income	0.437 (0.028)		
Correlation across waves		0.194 (0.046)	0.108 (0.047)
Correlation of ATR & MTR errors			0.545 (0.049)
Log likelihood	-2179.3678		
Parameters	52		
Number of obs	328		

Notes: Estimated parameters in bold are statistically significant at the 5% level. The optimization routine iterated between BHHH and BFGS quasi-Newton algorithm, converging after 29 iterations (with Stata convergence criteria of $qtolerance=1e-4$).

ATR misperceptions are inversely related to number series scores but there is no detectable relationship with verbal ability. Respondents with higher financial sophistication report lower average tax rates and higher marginal tax rates. In both equations the coefficient on financial sophistication is not statistically significantly different from zero.

True income is strongly correlated with being married, younger and having more education. It is also positively associated with cognitive ability and financial sophistication scores, but the relationship is small and is not statistically significant at the 5 percent level. There is no relationship between true income and using a paid tax preparer.

I also analyze whether experience in two types of occupations could explain heterogeneous misperceptions. I use occupation variables constructed by McFall, Kapinos and Willis (2011), who find that occupation is highly correlated with financial sophistication. The first indicator is for finance occupations, defined as working in occupations that do rates of return, cost-benefit analysis and investment decisions. The other indicator is for budgeting occupations, defined as occupations with resource management or budgeting content. The magnitudes of these coefficients are interesting. The direction of the estimates make intuitive sense, but they are unfortunately measured imprecisely. Having worked in a financial industry is associated with reporting lower MTR and ATR, relative to what is implied by income. People in finance occupations likely make more savvy investment decisions and strategically avoid taxes. This would lead to smaller reported MTR and ATR because there are aspects of true income (such as the type of income) that are not accounted for and would show up and underestimating tax rates (when in fact the respondent might be exactly correct but we do not capture it because of the imperfect measurement instrument). In contrast, having worked in a budgeting-focused job is associated with reporting higher MTR and ATR, relative to what is implied by income. And working in budgeting occupations might lead to habits associated with overestimating taxes. While working in these occupations might give someone greater tax knowledge, it is also possible that someone might develop the habit of taking conservative estimates in order to satisfy budgets. I am better off assuming my taxes will be higher than they are and being pleasantly surprised at the end-of-the-year “surplus” than to end up short because of underestimating the amount of one’s income going to federal taxes.

Differences across marital status and age could reflect imperfect modeling of the “true” tax system, so it’s important to look for systematic differences along these dimensions when analyzing heterogeneity in tax perceptions. The statistically significant coefficient on age is consistent with these respondents having additional tax credits (such as for being over 65) that I have not accounted for.

There is little evidence that taxpayers in states with income taxes have systematically different perceptions than people without state income taxes. This suggest that systematic overestimates of ATR are not driven by people including state income taxes. Log wealth is associated with smaller tax errors. There are two possible channels. People with lots of wealth are more likely to be getting income that is taxed at lower rates than wage and salary income. For these people my estimate of their “true” income may be problematic

because of crude distinction between income (and mostly treating all income the same way).

Respondents with tax preparers report systematically larger MTR and ATR. Number series and financial sophistication scores are associated with ATR misperceptions, even while conditioning on education, wealth and having a financial occupation. These measures of ability capture something above and beyond what is explained by differences in educational attainment.

7 Ex-post predictions of true income, MTR, and ATR

This section illustrates how to use the estimated population parameters to account for income measurement error when imputing latent true income, marginal tax rates and average tax rates. The estimated population parameters provide a way to combine information from noisy measures of income and perceived tax rates to refine the imputation of the true rates.

Now suppose the distribution of income measurement error and tax rate errors are taken as given and there is only a single wave of survey data. In particular, the econometrician knows the distribution of tax rates and income across the population, and observes data $\Omega_{i,w} = (y_{i,w}, a_{i,w}, m_{i,w})'$ for an individual i , which includes reported income, average tax rate and marginal tax rate in wave w .

7.1 Imputing true income

The estimated population parameters can be used to compute the conditional expectation of true log income $(\widehat{y_{i,w}^*})$, defined as

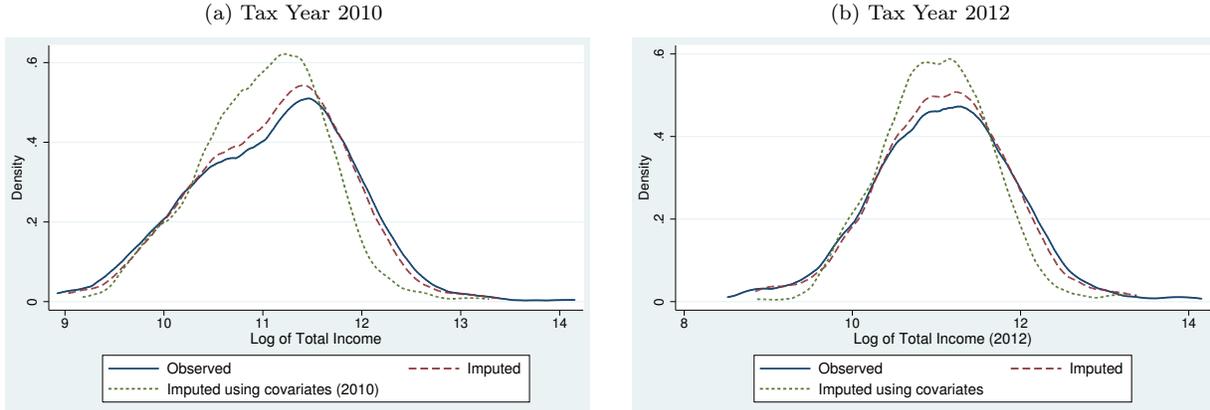
$$\widehat{y_{i,w}^*} = E \left(y_{i,w}^* \mid \widehat{\theta}, \Omega_{i,w} \right) = \int y_{i,w}^* \cdot \Pr \left(y_{i,w}^* \mid \widehat{\theta}, \Omega_{i,w} \right) \cdot dy_w^* \quad (13)$$

Figure 6 compares the distribution of observed log income with the distribution of imputed true log income $(\widehat{y_{i,w}^*})$. It presents the distribution of imputed log income from the baseline structural model without covariates, and imputed log income when using covariates. The top figure is for tax year 2010 and the bottom figure is for tax year 2012. In both years, accounting for income measurement error compresses the distribution of imputed true log income and does so more in the case with covariates. In particular, it is a few observations at the very top and bottom of the distribution of observed log income that get adjusted the most.

To better understand the value added from asking for marginal and average tax rates, I explore three assumptions about what data is available. First, conditional on $y_{i,w}$. Second, conditional on $y_{i,w}$ and $a_{i,w}$. Third, conditional on $y_{i,w}$, $a_{i,w}$ and $m_{i,w}$.

Figure 7 displays a scatterplot of observed log income on the x-axis and imputed log income on the y-axis. This measure is only conditional on reported income. In other words, this plot shows how the reported income gets transformed when trying to account for measurement error. Figure 7a and Figure 7b are for tax years 2010 and 2012, respectively. The diagonal line is at 45 degrees; observations above the line have

Figure 6: Distribution of the Log of Income: Observed vs. Imputed (2010)



Notes: These figures show the distribution of observed log income (reported by survey respondents), imputed log income from the full model without covariates and imputed log income from the full model with covariates. I use the Epanechnikov kernel and the plugin estimator of the asymptotically optimal constant bandwidth (Fan and Gijbels (1996)).

imputed income above the observed values ($\widehat{y_{i,w}^*} > y_{i,w}$) and observations below the line have imputed values less than the observed values ($\widehat{y_{i,w}^*} < y_{i,w}$).

Figure 8 displays the same scatterplots of observed log income on the x-axis and imputed log income on the y-axis, but this measure is conditional on reported income, as well as reported MTR and ATR. In each scatterplot, the size of the circle is proportional to the frequency of observation. Second set shows slightly increased variation in imputed rates. This is evidence that conditioning on tax rates does change imputed income, although the changes appear small in magnitude.

7.2 Imputing true tax rates

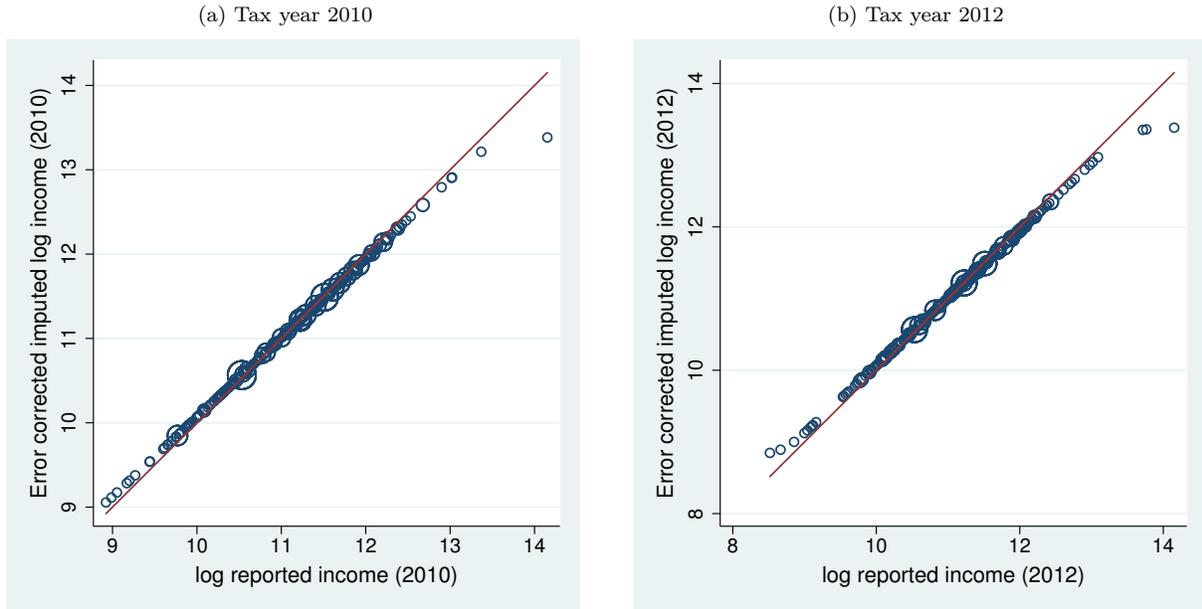
I impute true marginal and average tax rates using two approaches. The first approach puts imputed income directly into the known tax rate functions. The second approach takes the conditional expectation of the tax rates. Differences between the two values comes from non-linearities in the tax rate schedule. I also use a third approach to imputing marginal tax rates. I select the modal statutory tax rate, or the one that is most probable.

Rates associated with imputed income are directly comparable to the computed rates that comes from ignoring measurement error altogether. This provides the simplest way to understand the impact of income measurement error on the distribution of tax rates. The second approach is what should be used when using the imputed rates in regression analyses. Finally, the third approach provides the clearest way to understand the value added from imputing true tax rates using tax rates along with income.

7.2.1 Using imputed income

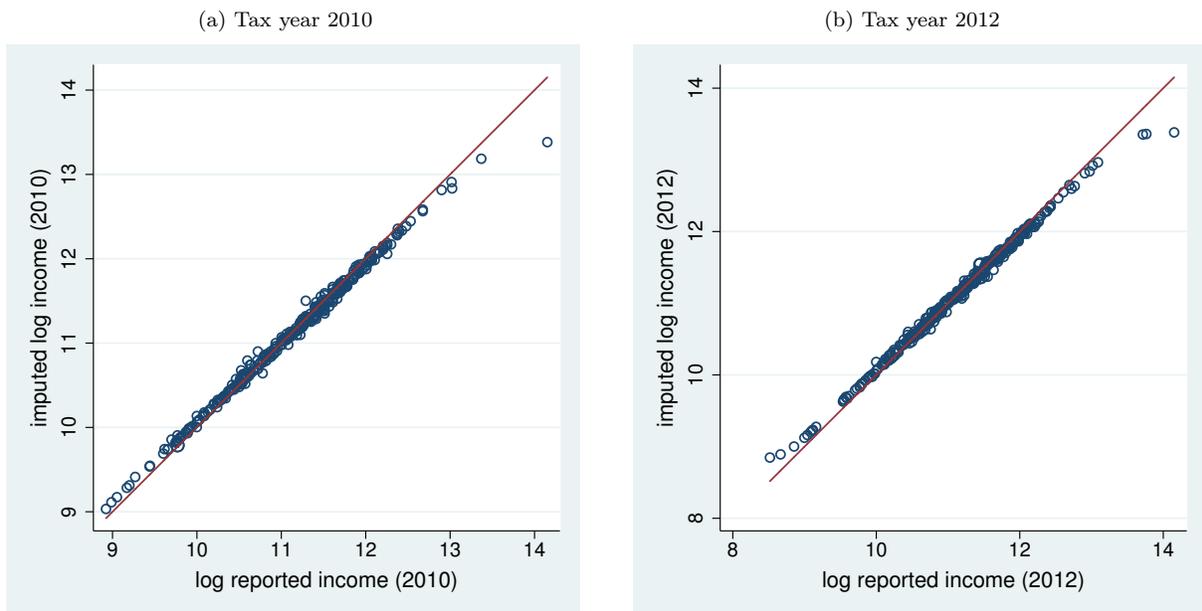
Imputed income $\widehat{y_{i,w}^*}$ is transformed into true taxable income,

Figure 7: Scatterplot of reported and imputed log income (not conditioning on ATR, MTR)



Notes: 7a is a scatterplot of log reported income and imputed log income in 2010 and 7b is the same thing for 2012. The imputations are the ex-post expected values conditional on reported income. This corrects for measurement error in income. The area of the symbol is proportional to the frequency of observation. The diagonal red line is at 45%, which corresponds to imputed income being equal to reported income.

Figure 8: Scatterplot of reported and imputed log income (conditioning on ATR, MTR)



Notes: 8a is a scatterplot of log reported income and imputed log income in 2010 and 8b is the same thing for 2012. The imputations are the ex-post expected values conditional on reported ATR, MTR and income. The area of the symbol is proportional to the frequency of observation. The diagonal red line is at 45%, which corresponds to imputed income being equal to reported income.

$$\widehat{I}_{i,w}^* = \max \left\{ s_{i,w} \cdot \exp \left(\widehat{y}_{i,w}^* \right) - D_{i,w}^* - E_{i,w}^*; 0 \right\} \quad (14)$$

which is plugged into the known marginal tax rate function $M_w(\cdot)$ to compute the marginal tax rate $\widehat{m}_{i,w}^* = M \left(\widehat{I}_{i,w}^* \right)$ and in the tax liability function $T_w(\cdot)$ to get average tax rate

$$\widehat{a}_{i,w}^* = \frac{T_w \left(\widehat{I}_{i,w}^* \right)}{s_{i,w} \cdot \exp \left(\widehat{y}_{i,w}^* \right)} \quad (15)$$

The same procedure is used to impute income when using different scenarios concerning the observed data. Note that all derivations for imputed rates can be adapted for when only income or only income and ATR are observed. In such cases, $\Omega_{i,w} = (y_{i,w}, a_{i,w})'$ or $\Omega_{i,w} = (y_{i,w})$. In particular, I also imputed marginal tax rates assuming that only income or only income and ATR are observed, rather than income, ATR and MTR.

Figure 9 presents the scatterplot of marginal tax rates computed from income, before and after accounting for income measurement error. Computed MTR is on the x-axis and measurement error-corrected imputed MTR is plotted on the y-axis. Figure 9a is for tax year 2010 and Figure 9b is the same thing for 2012. In both years, most observations are on the 45% degree line, which means their imputed MTR equals the computed MTR. Observations below 45% have MTRs edited downward when accounting for measurement error in income, while observations above 45% have MTRs edited upward.

7.2.2 Conditional expectation

The conditional expectation of the true MTR is the weighted average across statutory tax rates

$$\widehat{m}_{i,w}^* = \widehat{\mathbb{E}} \left[m_{i,w}^* \mid \widehat{\theta}, \Omega_{i,w} \right] = \sum_{j=1}^7 \left[\tau_j \cdot \Pr \left(m_{i,w}^* = \tau_j \mid \widehat{\theta}, \Omega_{i,w} \right) \right] \quad (16)$$

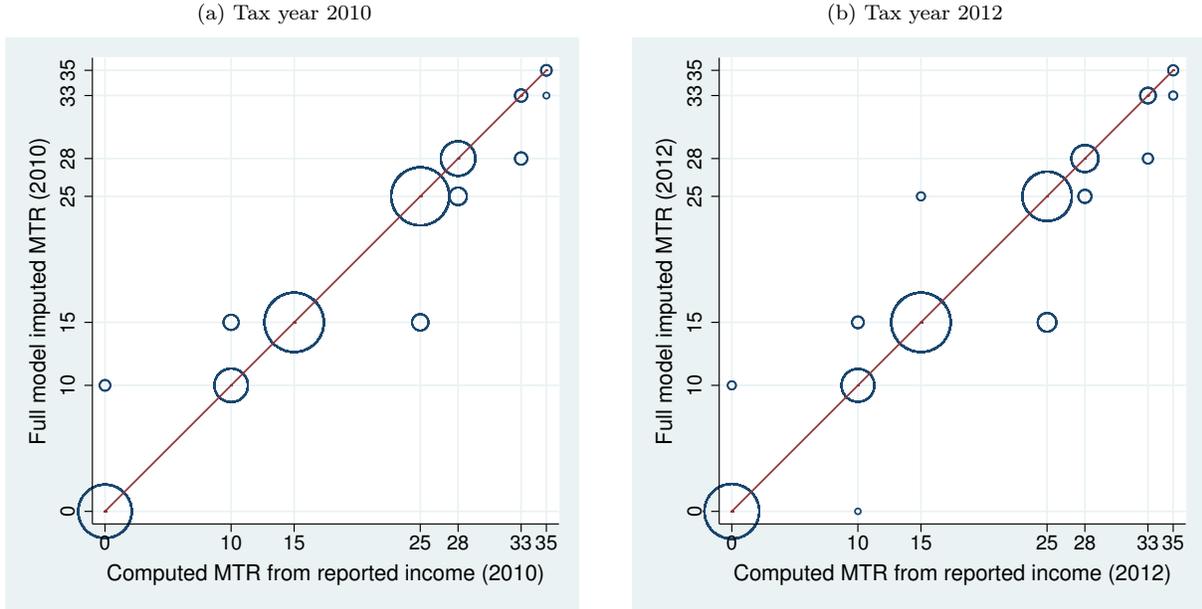
where $\Pr \left(m_{i,w}^* = \tau_j \mid \widehat{\theta}, \Omega_{i,w} \right)$ is the conditional probability that individual i has tax rate τ_j , given the observed data and structural parameter estimates. The conditional expectations are relatively easy to compute by Bayes' law, using the parameter estimates $(\widehat{\theta})$, survey measures of tax rates (\mathbf{r}_i) and log income (\mathbf{y}_i) . This corrects for random noise in the survey income and tax rates, and is based on the procedure Kimball et al. (2008) use to impute risk tolerance. The vector of predicted true rates are denoted $\widehat{\mathbf{m}}_i^*$ and $\widehat{\mathbf{a}}_i^*$, and they come from the following conditional expectations (where hats on the expectation operator mean estimates).

I also define the best predictor of true average tax rates in terms of the conditional expectation

$$\widehat{a}_{i,w}^* = \widehat{\mathbb{E}} \left[a_{i,w}^* \mid \widehat{\theta}, \mathbf{r}_i, \mathbf{y}_i \right] = \int a_{i,w}^* \cdot f \left(a_{i,w}^* \mid \widehat{\theta}, \mathbf{r}_i, \mathbf{y}_i \right) da_{i,w}^* \quad (17)$$

This computation is more intensive than the one for marginal tax rates. I use a change-of-variables formula

Figure 9: Scatterplot of computed MTR and imputed MTR (from full model)



Notes: These figures are scatterplots of computed MTR (from reported income) on the horizontal axis and imputed MTR on the vertical axis. The imputations are the marginal tax rate associated with ex-post expected log income, which is conditional on reported ATR, MTR and income. The area of the symbol is proportional to the frequency of observation. The diagonal red line is at 45%, which corresponds to imputed MTR being equal to computed MTR.

to cast this in terms of unobserved income measurement error, numerically integrate over the distribution of income measurement error and then compute $f(a_{i,w}^* | \hat{\theta}, \mathbf{r}_i, \mathbf{y}_i)$ using Bayes' Theorem.

Figure 10 presents the scatterplot of average tax rates computed from income and the conditional expectation of true ATR. Computed ATR is on the x-axis and measurement error-corrected imputed ATR is plotted on the y-axis. Most observations only change slightly when accounting for income measurement error.

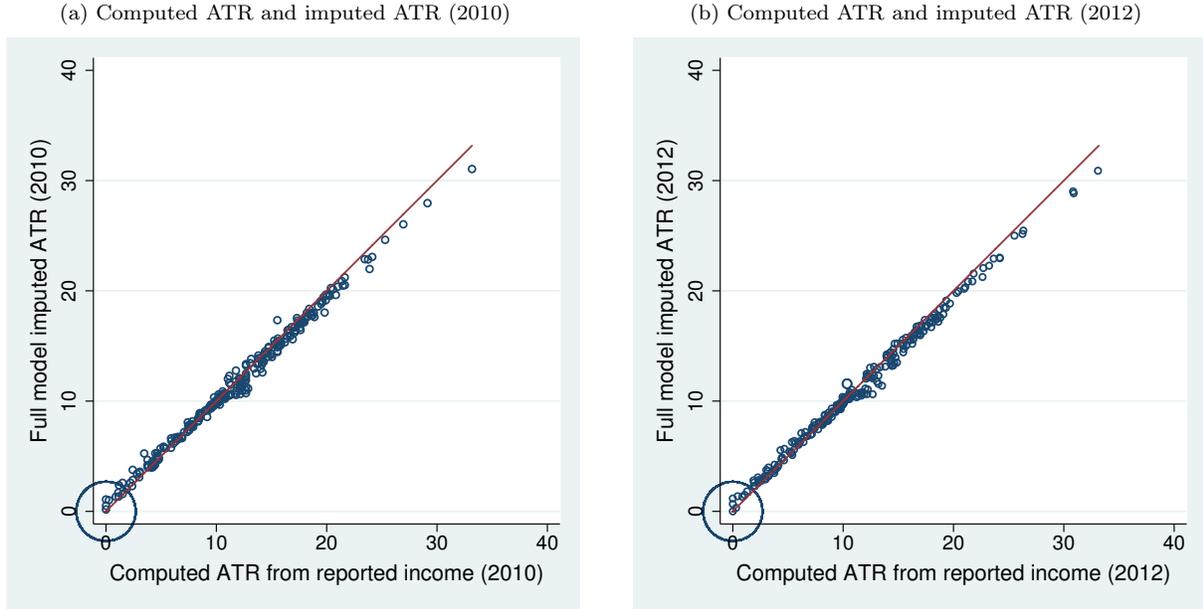
7.2.3 MTR associated with maximum probability

Finally, given the finite number of marginal tax rates, one could argue that the most reasonable imputed value is the modal tax rate, or that which is “most probable.” Using this approach, I compute the probability of having each statutory rate τ_j , $\Pr(m_i^* = \tau_j | \hat{\theta}, \Omega_{i,w})$, and select the one with the highest probability of being the true rate:

$$\widetilde{m}_{i,w}^* = \arg \max_{\tau_j} \left\{ \Pr(m_i^* = \tau_j | \hat{\theta}, \Omega_{i,w}) \right\} \quad (18)$$

This is equivalent to estimating the probability that true taxable income is within income bracket associated with rate τ_j , $\Pr(\overline{I}_w^{j-1} \leq I_{i,w}^* < \overline{I}_w^j | \hat{\theta}, \Omega_{i,w})$. The probability associated with the modal MTR provides information about the confidence with which I assign statutory marginal tax rates.

Figure 10: Scatterplot of computed ATR and imputed ATR from full model



Notes: 10a is a scatterplot of computed ATR and imputed ATR in 2010 and 10b is the same thing for 2012. The imputations are the ex-post expected values conditional on reported ATR, MTR and income. The area of the symbol is proportional to the frequency of observation. The diagonal red line is at 45%, which corresponds to computed ATR being equal to imputed ATR.

7.3 Do tax rate perceptions help improve imputed true rates?

Looking at a few individual cases can help show the structure of the measurement model and imputation procedures. There were seven respondents in 2010 who were married, reported income of \$100,000, had no dependent exemptions, and all of their income included in AGI (share equal to 1). Taking this income data as given, these respondents had computed ATR of 12.7 and computed MTR of 25. The imputation procedure is an attempt to correct for measurement error in income when computing the rates. This provides a clear way to observe how the survey measures of MTR and ATR impact imputed income and true tax rates. Consider the observations listed in Table 9 :

Table 8: Examples of individual cases

R's ID	Observed (2010)			Imputed Y			Cond. Expect.		Modal MTR	
	Income	MTR	ATR	$\widehat{y}_{i,1}^*$	$\widetilde{m}_{i,1}^*$	$\widetilde{a}_{i,1}^*$	$\widehat{m}_{i,1}^*$	$\widehat{a}_{i,1}^*$	$\widetilde{m}_{i,1}^*$	$\Pr(\widetilde{m}_{i,1}^*)$
R1: 7013920010	100,000	3	15	85,053	15	10.7	17.8	11.2	15	0.72
R2: 7013920020	100,000	7	7	87,780	25	11.0	19.3	11.5	15	0.57
R3: 7004980020	100,000	10	10	89,823	25	11.3	19.9	11.7	15	0.51
R4: 7020620010	100,000	15	28	94,043	25	11.9	20.2	12.2	25	0.51
R5: 7005380020	100,000	18	18	95,748	25	12.1	21.5	12.3	25	0.65
R6: 7007370010	100,000	25	9	98,397	25	12.5	23.5	12.5	25	0.84
R7: 7013590010	100,000	35	30	105,728	25	13.4	24.1	13.3	25	0.89

For these observations, differences in reported tax rates results in variation in imputed income, which ranges from \$85,053 to \$105,728. Respondent R2 reports $ATR=MTR=7$ and gets imputed income of 87,780, whereas R3 reports $ATR=MTR=10$ and gets imputed income of 89,823 and R5 reports $ATR=MTR=18$ and gets imputed income of 95,748. Respondents R6 and R7 correctly report $MTR>ATR$ and give statutory marginal tax rates, and their probabilities of their modal MTR are by far the largest among the seven respondents listed.

Imputed tax rates change as a result of accounting for income measurement error. For example, respondent R1 has an imputed MTR (from imputed Y) different from the computed MTR (of 25). Accounting for income measurement error pushes imputed rates from the conditional expectations toward the mean. Overall, 8.3% of respondents in each wave have imputed MTR different from their computed MTR, and 15% had an imputed MTR different from computed MTR in at least one wave.

There are sometimes large differences between the imputed rates across the imputation procedures. The marginal tax rate function is concave. The conditional expectations depend on how far the imputed income is to the tax bracket thresholds and on the conditional distribution of income overall. Taking the ex-post expected value has a substantial impact on the imputed tax rates relative to the other approaches. By taking the weighted average, the discrete marginal tax rates are transformed into a continuous variable. Because the tax functions are non-linear, the tax rates computed from predicted income will typically be biased. Even when observed income is an unbiased prediction of true income, using an unbiased prediction of income can induce systematic bias in predicted tax rates.

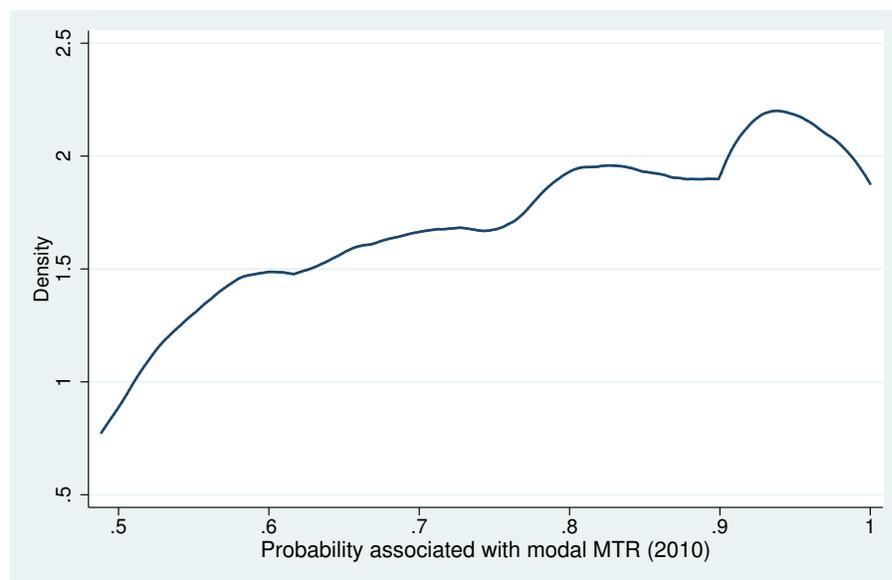
Even though tax rate perceptions vary greatly, they can still be useful when imputing true rates. While I do not know who has accurate tax rate beliefs and who does not, I can still use the information from the consistency of the reported tax rates with reported income. If tax rates and income are consistent, then I am confident about the true tax rate. If tax rates are consistent with one another, but not with income, this places greater weight on the tax rates than cases in which the ATR and MTR are inconsistent with one another. Therefore, having both ATR and MTR provides information about the accuracy of the rates.

7.3.1 Analyzing distribution of modal MTR probabilities

Incorporating additional information from reported tax rates generates more confidence about how the model classifies people into statutory marginal tax rate categories. This result comes from analyzing the empirical distribution of the probability mass on the modal tax bracket. For each observation, I compute the probability that true income is within each of the statutory marginal tax rate brackets, then select the modal MTR and the probability associated with it. Figure 11 displays the empirical distribution of the probabilities of the modal MTR for 2010 when conditioning on income and accounting for measurement error. This is meant to aid intuition. If MTR were known with certainty, the entire probability mass would be on the true MTR for each individual and the empirical distribution would be a mass at probability 1. If MTR were completely random, the probabilities would be evenly distributed across the seven tax rates,

with 0.14 probability each, and the empirical distribution would be a mass at 0.14.

Figure 11: Distribution of the modal MTR probabilities in 2010



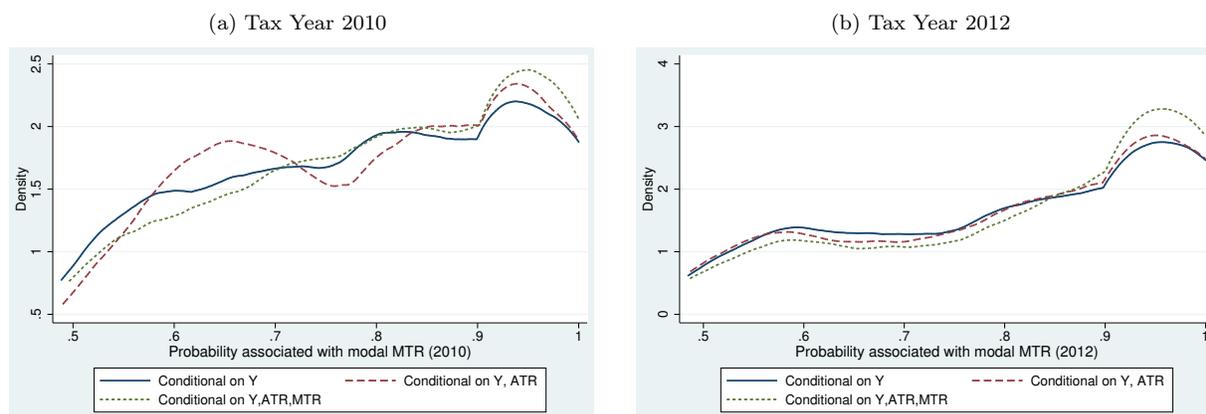
Notes: These figures show the distribution of observations across estimated probabilities of having the modal marginal tax rate. This figure is constructed using the modal MTR when conditional on log income and accounting for measurement error. I use the Epanechnikov kernel and the plugin estimator of the asymptotically optimal constant bandwidth (Fan and Gijbels (1996)).

Figure 12 compares the empirical distribution of the probabilities of the modal MTR when conditioning on income only; on income and ATR; and then on income, ATR and MTR. Changes in the distribution reflect changes in the probabilities associated with the modal MTR. Figure 12a plots these distributions for tax year 2010 and Figure 12b plots these distributions for tax year 2012. A noticeable difference between the imputations for 2010 and 2012 is the bi-modal distribution of the probabilities when conditioning on income and ATR in 2010 but not in 2012. In both waves, including both ATR and MTR results in greater mass toward probability one. Conditioning on survey measures of ATR and MTR result in greater likelihood of classifying respondents by their true MTR.¹¹

The share of adjusted gross income (AGI) in total income has a big effect on the probability of the modal marginal tax rate. Figure 13 shows a scatterplot and a fitted third-degree polynomial of this relationship between AGI share and the probability of the modal MTR. Intuitively, the tax rate is known to be zero for non-taxed income. As the share of AGI in total income decreases, the share of non-taxed income increases and the probability of being in the zero tax bracket increases until the probability approaches one. More

¹¹This ad hoc comparison could be addressed more formally in the context of statistical hypothesis testing. The empirical question would be how additional information from reported ATR and MTR change the probability of having type 1 and type 2 errors in the imputed MTR. From a Bayesian perspective, type 1 error occurs when information should not substantially change one's prior estimate of the probability, but does. Similarly, a type 2 error is one that looks at information that should change one's estimate, but does not. In the problem at hand, a type 1 error means that the ex-post imputed value is different from the computed rate when the computed rate is in fact correct. A type 2 error means that the ex-post imputed value fails to correct the tax rate when the computed rate is incorrect. While this analysis is beyond the scope of this paper, it is a potentially interesting way to assess the various marginal tax rate imputation procedures.

Figure 12: Distribution of the modal MTR probabilities: conditioning on different information



Notes: These figures show the distribution of observations across estimated probabilities of having the modal marginal tax rate. They are constructed using the modal MTR when conditioning on log income only, when conditioning on income and ATR perceptions, and then on income, ATR and MTR. I use the Epanechnikov kernel and the plugin estimator of the asymptotically optimal constant bandwidth (Fan and Gijbels (1996)).

generally, the share of taxable income $s_{i,w}$ has an important effect on imputations.

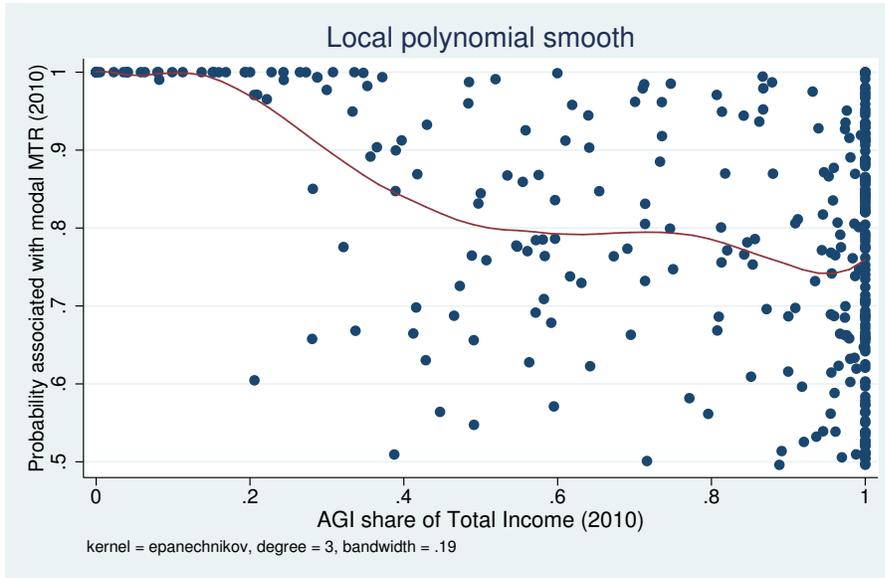
Figure 14 recreates Figure 12a and Figure 12b but this time restricting the sample to respondents who have a share of AGI in total income that is above 0.8. The distributions are strikingly different when limiting the sample in this way. The patterns are similar to when using the unrestricted sample. For 2010, conditioning on only ATR and income can reduce the probability of modal MTR than when using income alone.

Figure 15 shows the distribution of modal MTR probabilities varies across subgroups. Comparing the distribution of modal MTR probabilities across subgroups illustrates important heterogeneity in the quality of the survey tax rate measures. In particular, Figure 15a compares respondents who were working in 2010 with those who were not working. It shows the distribution of modal MTR probabilities shifts to the right for respondents who were working.

In Gideon (2014), I classify respondents into four types of tax perceptions based on their income and tax rate responses. Figure 15b compares the distribution across these four types. Type A respondents are those who distinguish between ATR and MTR and report statutory marginal tax rates. Type B respondents distinguish between ATR and MTR and do not report statutory marginal tax rates. Type C respondents do not distinguish between ATR and MTR and report statutory marginal tax rates for both. Type D respondents do not distinguish between ATR and MTR and report average tax rates for both. There are noticeable differences across tax perceptions but I do not have a clear interpretation of these patterns.

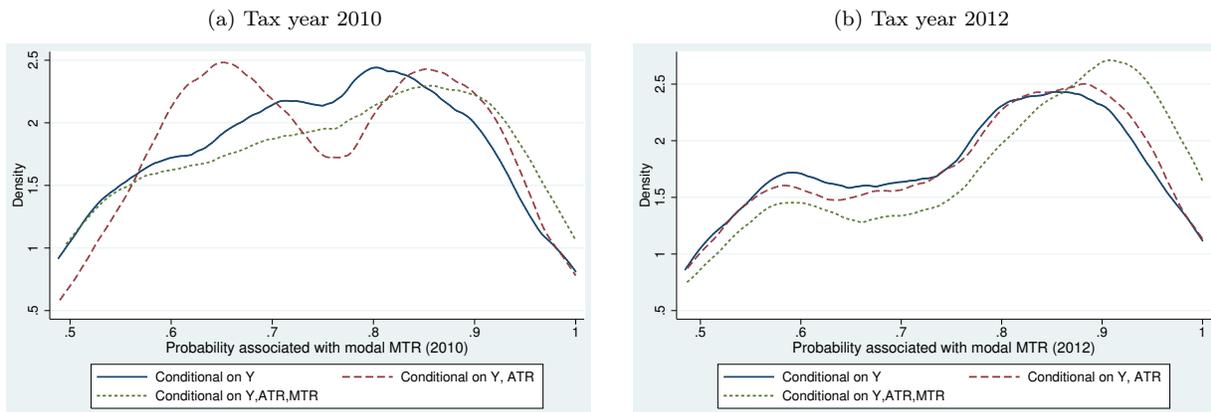
Figure 16 presents results by marital status. In particular, Figure 16a shows the distribution for married households exhibits more dispersion than for single households. Figure 16b unpacks the difference by marital status, comparing the distribution of each household's financial respondent with his or her spouse. The financial respondent is assigned based on self-assessment of one's own versus one's partners knowledge of

Figure 13: Share of AGI in total income and the modal MTR probabilities (in 2010)



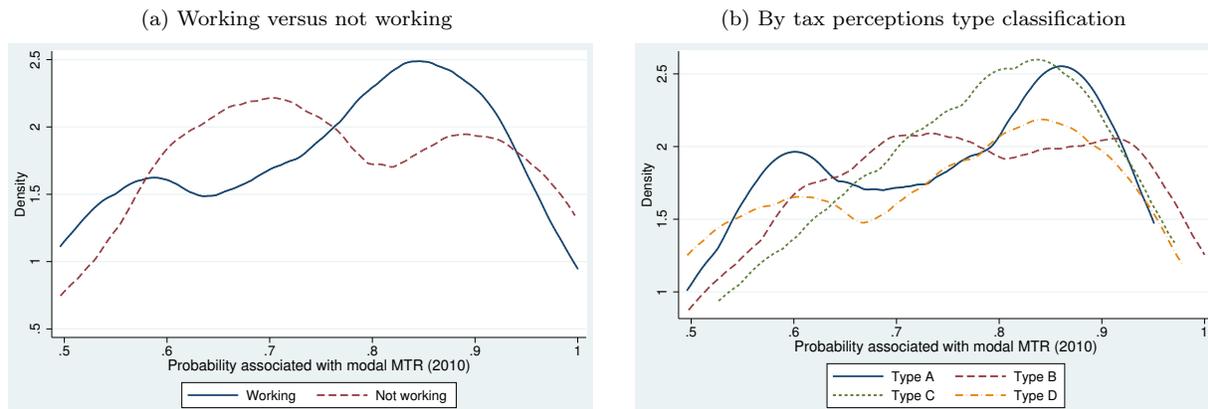
Notes: This figure shows the relationship between the share of total income included in adjusted gross income (AGI) and the modal MTR probabilities. These probabilities are computed using the distribution of true income conditional on income, ATR and MTR.

Figure 14: Distribution of modal MTR probabilities: restricted to sample with AGI share of total income above 0.8



Notes: These figures show the distribution of observations across estimated probabilities of having the modal marginal tax rate. The sample is restricted to households with adjusted gross income (AGI) that was at least four-fifths of their total income. This recreates Figure 12a and Figure 12b but with the restricted sample.

Figure 15: Distribution of modal MTR probabilities in 2010: by working status and classification of tax perceptions



Notes: These figures show the distribution of modal MTR probabilities in 2010 across subgroups of the sample. The density is conditional on income, ATR and MTR. The top figure compares respondents who were working in 2010 and those who were not. The top bottom figure compares the distribution across Chapter 2 classification of the mental models. Type A respondents are those who distinguish between ATR and MTR and report statutory marginal tax rates. Type B respondents distinguish between ATR and MTR and do not report statutory marginal tax rates. Type C respondents do not distinguish between ATR and MTR and report statutory marginal tax rates for both. Type D respondents do not distinguish between ATR and MTR and report average tax rates for both. The sample is restricted to households with adjusted gross income (AGI) that was at least four-fifths of their total income.

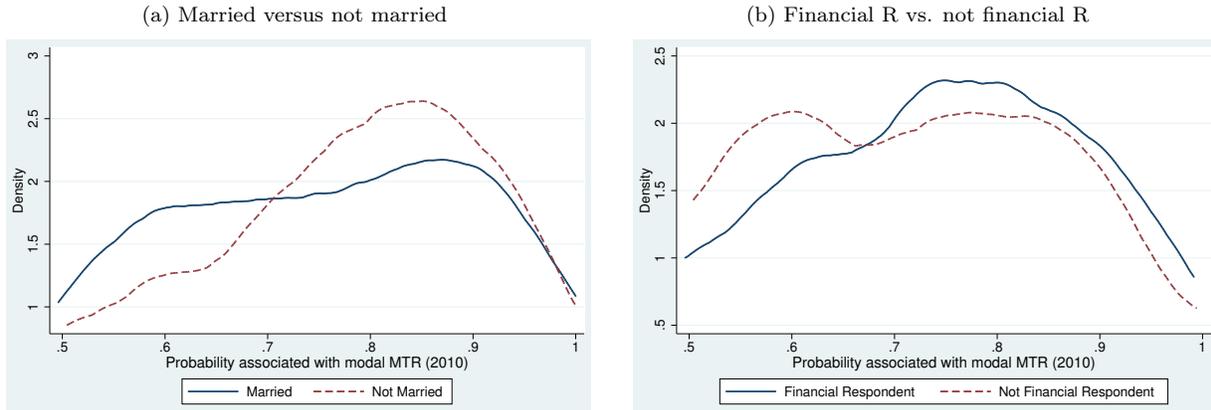
their finances. Data from the designated “financial respondent” generates higher modal MTR probabilities. This provides evidence that differences by marital status is partly due to specialization within married households. The probabilities of the modal MTR are in general larger for financial respondents than his or her spouse. This is consistent with tax rate perceptions being informative about true income among respondents who are more knowledgeable about their financial situation.

8 Discussion and conclusion

This paper describes the results of an exploratory study of income tax perceptions. I advance new survey methodology for measuring marginal and average income tax rates. Questions asking about average and marginal federal income tax rates was fielded on CogEcon 2011 and asked again on CogEcon 2013. Using a panel of survey measures which combines marginal tax rates, average tax rates and income, I find evidence of systematic bias and significant heterogeneity in how people perceive income tax rates. I find compelling evidence of systematic errors in perceived average and marginal tax rates that cannot be explained by measurement error in income. Respondents tend to overestimate average tax rates and underestimate marginal tax rates. These survey measures, however imperfect, shed light on heterogeneous perceptions and can be combined with the typical approach of using income to impute better measures of true rates.

It is perhaps surprising that many people answered the question about ATR but not about MTR, and that perceptions of ATR have both a larger mean bias and smaller standard deviation than perceptions of

Figure 16: Distribution of modal MTR probabilities in 2010: by marital and financial R status



Notes: These figures show the distribution of modal MTR probabilities in 2010. The density is conditional on income, ATR and MTR. The top figure compares respondents who were married in 2010 and those who were not married. The bottom figure compares the distribution for the household financial respondents and their spouse or partner. The sample is restricted to households with adjusted gross income (AGI) that was at least four-fifths of their total income.

MTR. A potential explanation is that household financial decisions might depend more on budgeting and allocation of resources rather than optimizing marginal decisions. This could also rationalize the systematic overestimation of perceived ATR, as it's better to have resources left over due to overestimating tax liability than to come up short on April 15th.

People might not understand the definition of marginal tax rate yet still make decisions about labor supply with income tax consequences in mind. While we attempt to elicit perceptions of these perceived tax consequences, an important limitation of this survey approach is that I cannot confirm that reported rates are the same as what people use when making decisions. Several respondents provided feedback at the end of the study that they consult an accountant when they have questions about taxes, suggesting that important decisions would be based on their accountant's perceptions rather than their own, and these need not coincide.

This exploratory study motivates several next steps. From a methodological perspective, there is a need for further work to refine the way in which the tax rate questions are posed to respondents. Respondents seem willing and able to respond meaningfully to questions eliciting their average tax rate but not questions about marginal tax rates. If privacy were a concern, people should be more willing to answer about statutory MTR than about ATR. This is because MTR conveys information about a range of taxable income, whereas the ATR conveys more precise information about taxable income and the fraction of untaxed income. This suggests that people are willing to answer personal questions about income taxes, but only when they understand what is being asked. Respondents might also better understand the concept of marginal tax rate if they are provided with a concrete example. Yet the differential response rate for the ATR and MTR questions was almost as common in CogEcon 2011, when the preamble included an example of MTR, as in CogEcon 2013, when there was no example. Since the example does not seem to be enough, the question

would need to more clearly convey the intuition of a “marginal tax rate.” In fact, it might make sense to not use the word “marginal” at all, but explain what it means to be on the “margin.” However, we pilot tested similar questions but faced the same confusion. We should not marginalize respondents for not understanding the standard economic terminology. That said, many respondents provided sufficiently precise tax rates that it appears they were able to understand the questions. More importantly, self-assessed knowledge of taxes was highly predictive of who responded to the questions. This suggests the questions were adequately clear to a knowledgeable respondent, but too confusing for people with less technical knowledge of the tax system.

Perhaps the best solution is to ask for a categorical answer for the marginal tax rate. This would reduce confusion about whether the question is asking about a statutory or effective marginal tax rate. It would also reduce the cognitive burden on the respondent. While the reported rates no longer convey information about whether the respondent knows the statutory rate schedule, it should drastically reduce non-response and would provide a more unified interpretation of the marginal tax rate errors in terms of how people map their subjective taxable income into the tax rate categories.

It would be advantageous to conduct this study using a more representative sample. The CogEcon sampling frame includes Americans age 50 and older. Moreover, respondents have higher cognitive ability and income than the typical person their age. Older, better educated, more intelligent households are likely more knowledgeable about tax rates than the population overall. This should be kept in mind when interpreting the results. Evidence of substantial heterogeneity and systematic misperceptions in the CogEcon sample suggests this is an important phenomena overall. That said, the problems respondents had understanding marginal tax rates will likely be even more pronounced in a truly nationally representative sample, so refining and improving the survey questions is of first-order importance.

Taxpayers cannot be expected to perfectly understand the behemoth U.S. tax system. My results suggest that people are not perfectly knowledgeable about their income tax rates, but are also not completely ignorant. Understanding heterogeneity in perceptions is key to improving inference in models estimated using tax rates and key to understanding responsiveness to tax incentives. Despite these caveats and qualifications, I am optimistic that collecting survey measures of tax rates can be used as both an instrument for behavioral studies and as a way to refine measures of true tax rates when administrative data is unavailable.

References

- [1] Abeler, Johannes and Simon Jager. 2013. "Complex Tax Incentives: An Experimental Investigation." IZA DP No. 7373.
- [2] Abowd, John M. and Martha H. Stinson. 2013. "Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data." *The Review of Economics and Statistics*, forthcoming
- [3] Auld, D. A. L. 1979. "Public Sector Awareness and Preferences in Ontario," *Canadian Tax Journal* 27: 172-83.
- [4] Bhat, C.R. 2001. Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model. *Transportation Research* 35B: 677-695.
- [5] Bollinger, Christopher. 1998. "Measurement Error in the CPS: A Nonparametric Look" *Journal of Labor Economics* 16(3): 576-594.
- [6] Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Measurement error in survey data." In *Handbook of labor economics*, vol. 5, ed. J. Heckman and E. Leamer, 3707-3843. Amsterdam: Elsevier.
- [7] 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, no. 1 (January): 1-24.
- [8] Brown, C.V. 1968. Misperceptions About Income Tax Incentives. *Scottish Journal of Political Economy*, February: pp. 1-12.
- [9] Carroll, RJ, D. Ruppert and LA Stefanski. 2006. *Measurement Error in Nonlinear Models: A Modern Perspective*, Second Edition. Chapman & Hall: New York.
- [10] Duncan, Greg, and Daniel Hill. "An Investigation of the Extent and Consequences of Measurement Error in Labor-Economic Survey Data." *Journal of Labor Economics* 3, no. 4 (1985): 508-32.
- [11] Dungan, Adrian and Michael Parisi. Individual Income Tax Rates and Shares, 2010. *Statistics of Income Bulletin*, Winter 2013, pp. 18-63.
- [12] Enrick, Norbert Lloyd. 1963. "A Pilot Study of Income Tax Consciousness." *National Tax Journal* 16: 169-73.
- [13] Enrick, Norbert Lloyd. 1964. "A Further Study of Income Tax Consciousness." *National Tax Journal* 17: 319-21.
- [14] Fan, J. and I. Gijbels. 1996. *Local polynomial modelling and its applications*. Chapman & Hall.
- [15] Feenberg, Daniel, and Elizabeth Coutts. 1993. "An Introduction to the Taxsim Model." *Journal of Policy Analysis and Management* 12(1): (<http://www.nber.org/taxsim>)
- [16] Fisher, Gwenith, Michael Gideon, Joanne W. Hsu and Brooke McFall. 2011. "Cognitive Economics Study: Development and Methodology." November 2011. Available at <http://cogecon.isr.umich.edu/survey/docs/dev+meth.pdf>
- [17] Fujii and Hawley. 1988. "On the Accuracy of Tax Perceptions." *The Review of Economics and Statistics*.

- [18] Gideon, Michael. 2014. "Uncovering Heterogeneity in Tax Perceptions." Mimeo.
- [19] Gideon, Michael, Seth Koch and Brooke Helpie McFall. 2013. "Cognitive Economics Study: CogEcon 2011 Documentation." December 2013. Available at <http://cogecon.isr.umich.edu/survey/docs/datadoc2011.pdf>
- [20] Haider, Steven and Gary Solon. 2006. "Life-Cycle Variation in the Association between Current and Lifetime Earnings." *American Economic Review*. Vol 96, No. 4, pp. 1308-1320.
- [21] Hudomiet, Peter. 2013. "Dynamic survey measurement error in annual earnings: The role of cognition, earnings shocks and untaxed earnings." Working paper (version: May 28, 2013).
- [22] Hoopes, Jeffrey, Daniel Reck and Joel Slemrod. 2013. "Taxpayer Search for Information: Implications for Rational Inattention." NBER working paper 19482.
- [23] 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.
- [24] Kimball, Miles S., Claudia R. Sahm, and Matthew D. Shapiro. 2008. "Imputing Risk Tolerance from Survey Responses." *Journal of the American Statistical Association* 103(September): 1028-1038.
- [25] Kolenikov, Stanislav. 2001. "Review of Stata 7," *Journal of Applied Econometrics*, Vol. 16, Issue 5: 637-646.
- [26] Lewis, Alan. 1978. "Perceptions of Tax Rates," *British Tax Review*. 6: 358-66.
- [27] McFall, Brooke, Kandice A. Kapinos and Robert J. Willis. 2011. "Occupational Learning, Financial Knowledge, and the Accumulation of Retirement Wealth." Mimeo, University of Michigan.
- [28] Parsons, Donald O. 1978. The Autocorrelation of Earnings, Human Wealth Inequality, and Income Contingent Loans. *Quarterly Journal of Economics*, Vol. 92, No. 4, pp. 551-569.
- [29] StataCorp. 2013. *Stata 13 Base Reference Manual*. College Station, TX: Stata Press.
- [30] Statistics of Income Publication 1304, "Individual Income Tax Returns 2010." Table 1.1: All Returns: Selected Income and Tax Items, by Size and Accumulated Size of Adjusted Gross Income."
- [31] Sullivan, Paul. 2009. "Estimation of an Occupational Choice Model when Occupations are Misclassified," *Journal of Human Resources* 44(2): 495-535.
- [32] Train, Kenneth. "Halton Sequences for Mixed Logit." Working paper, University of California, Berkeley, 1999.
- [33] Train, Kenneth. 2003. *Discrete Choice Methods with Simulation*. Cambridge: Cambridge University Press.
- [34] Wagstaff, J.V. 1965. "Income Tax Consciousness Under Withholding." *Southern Economic Journal* XXXII: 73-80.

A Survey Instruments

A.1 CogEcon 2011 tax rate questions

Instructions for the following questions:

*These questions focus on current and future federal income tax rates, both in general and for you personally. The **marginal tax rate** is the tax rate on the last dollars earned. For example, if a household's income tax bracket has a marginal tax rate of 15%, then a household owes an extra \$15 of taxes when it earns an extra \$100.*

Answer each question with a percentage between 0 and 100. Please provide your best estimate of the marginal tax rate even if you are not sure.

*These questions are about **federal income taxes only**; please do not include state or local taxes, or payroll taxes for Social Security and Medicare.*

Q1. *For a household in the **highest tax bracket** in 2010:*

The marginal tax rate on wage and salary income was _____% and the marginal tax rate on dividend income was _____%.

Q2. *Tax rates may change in the future. I think that for a household in the **highest tax bracket** in 2014:*

The marginal tax rate on wage and salary income will be _____% and the marginal tax rate on dividend income will be _____%.

We now want to ask you about your household's federal taxes. Please use the same definitions of federal income tax and marginal tax rate as on the previous page.

Q3. *Please think about your household's income in 2010 and the amount of federal income tax you paid, if any.*

Approximately what percentage of your household income did you pay in federal income taxes in 2010?
_____%

Q4. *Now we want to ask about **your households' marginal income tax rate**. Please think about your household's federal income tax bracket and the tax rate on your last dollars of earnings.*

In 2010, my household's marginal tax rate was _____%.

Q5. *Suppose that in 2014 your household receives the same income you had in 2010. However, the federal income tax schedule might change.*

I expect that my household's marginal tax rate **in 2014** would be _____%.

A.2 CogEcon 2013 tax rate questions

Instructions for the following questions:

*These next two questions focus on your federal income tax rates. These questions are about **federal income taxes only**; please do not include state or local taxes, or payroll taxes for Social Security and Medicare.*

Answer each question with a percentage between 0 and 100. Please provide your best estimate of the marginal tax rate even if you are not sure.

Q1. *Please think about your household's income in 2012 and the amount of federal income tax you paid, if any.*

Approximately what percentage of your household income did you pay in federal income taxes in 2012?

_____%

Q2. *Now we want to ask about **your households' marginal income tax rate**. The marginal tax rate is the tax rate on the last dollars earned. Please think about your household's federal income tax bracket and the tax rate on your **last** dollars of earnings.*

In 2012, my household's marginal tax rate was _____%.

B Data construction and measurement

B.1 Constructing baseline taxable income measure

The objective is to measure taxable income for each taxable unit in 2010 and 2012 using self-reported information about household income from CogEcon 2011 and CogEcon 2013, respectively.

B.1.1 Measuring gross income: Total and components

CogEcon measures household-level gross income for each respondent as follows. First, respondents are asked for the total combined income of all members of their family (living in same household) during the past 12 months (Question C2). This includes wages or salary, net income from business, farm or rent, pensions, dividends, interest, Social Security payments, and any other money or income. Then respondents are asked about four particular sources of income: (i.) wages and salary, (ii.) employer-provided pensions, (iii.) Social Security and (iv.) distributions from retirement accounts. The questions were identical across the two waves of the study, although some question numbers and the reference year changed. Respondents who do not provide an exact value are then asked to select a range of values from the list of ranges presented in the question.

For income categories (i.), (ii.) and (iii.), there were separate questions about the respondent's own income and other questions in which the respondent is supposed to answer on behalf of his or her partner. Distributions from retirement accounts are only asked about in terms of the household's distributions overall, not distinguishing between distributions from accounts held by one partner or the other.

The survey asked respondents whether they currently received Social Security or Railroad Retirement benefit payments. If yes, they were asked to provide the age at which benefits were first received and the amount received each month. The same set of questions was asked about the respondents' spouse or partner, when applicable.

Respondents were then asked whether they currently received payments from an employer- or union-provided (defined benefit) pension plan. If yes, they were asked to provide the amount received each month. Again, the same set of questions was asked about the respondents' spouse or partner, when applicable.

Next, they were asked about total earnings before taxes from all jobs during the previous calendar year. Question C33 in CogEcon 2011 asked for the respondent's total earnings before taxes from all jobs during 2010; question C28 in CogEcon 2013 asked the same thing but about earnings during 2012. Again, the same set of questions was asked about the respondents' spouse or partner, when applicable. Question C42 in CogEcon 2011 asked for the spouse's/partner's total earnings before taxes from all jobs during 2010; question C36 in CogEcon 2013 asked the same thing but about earnings during 2012.

Later in the survey, respondents are asked whether they (or their spouse/partner) withdrew money or received payments from tax-advantaged retirement accounts in the previous calendar year. This includes 401(k) plans, 403(b) plans, Keoghs, traditional IRAs and Roth IRAs. If yes, they are asked for the amount withdrawn, before taxes and other deductions.

Using these data, I construct two measures of household-level gross income. The first equals the self-reported sum of all income:

$$Inc_{total} = C2_{val}$$

The second measure equals the sum of the component parts of income (that is, based on the questions discussed above):

$$Inc_{sum} = wage + pension + SocSec + retdist$$

There are a few issues of measurement that must be accounted for when comparing the income components and the self-reported total household income.

First, I made the following edits to the income data. For reported total income (Inc_{total}), question C2, I make the following edits. If the yearly income was reported as less than \$200, I assume it was written in thousands (e.g., 75 rather than 75,000) and I multiply the reported income by 1000. It was apparent that these respondents wrote numbers in thousands, which could have been unintentional shorthand for their full answer. Looking at individual cases—comparing reported income across waves and reported wealth amounts across waves—corroborates this interpretation of the data.

Next, if they clearly gave their monthly income (e.g., they reported the number that they gave as their monthly Social Security retirement benefit) then I multiply the amount by 12 to convert it into yearly. If the reported monthly Social Security benefit is greater than \$5,000 per month, I assume this was meant as the yearly amount (since it is more than any household could receive in a month) and divide it by twelve. If the monthly pension is reported as greater than \$20,000 per month, I assume this was meant as the yearly amount and divide it by twelve. Finally, I divide the reported amount by 100 for a particular respondent who provided other information that led to the conclusion that what the respondent reported as \$10000000 was intended as \$100000.¹²

If the reported monthly Social Security benefit is less than \$10 per month and includes decimal places, then I assume the value was keyed in with a period rather than a comma and multiply it by 1000. If the reported monthly Social Security benefit is less than \$50 per month, I assume this was meant as the yearly amount and report monthly value after multiplying by 1000 and divide it by 12. If the yearly wage/salary value is reported as less than \$200, I assume this was written in thousands and multiple by one thousand.

Second, Social Security retirement benefits and pensions are reported as a monthly value but I do not know how many months the respondent (or spouse) received these sources of income. I impute Social Security and pension income during 2010 and 2012 using the reported monthly value and an imputed number of months they were received.

I calculate the Social Security retirement benefits received during the specified tax year (either 2010 or 2012) in the following steps. I impute the number of months using self-reported age when benefits were first received, along with information about birth month and year (collected by CogUSA in prior waves). The birth month is missing for a handful of respondents. In such cases I assume it was in July. I assume benefits were claimed immediately when the respondent turned the specified age and then received beginning in the subsequent month. For example, if someone claimed benefits at age 62 and was 63 in January 2010, assume they received benefits for all twelve months of 2010. If they turned 62 in May 2010, then assume they received benefits for seven months, June thru December. If the respondent received benefits at all during 2010 then I assume they received benefits for all 12 months of 2012. I then multiply the number of months times the reported monthly amount to get the imputed social security benefits received during the specified year. I follow the same procedure for the spouse's Social Security benefits. If the respondent claims to be married or partnered (with a financial future) but the person is not in the sample, I assume that the spouse is the same age as the respondent. The respondent and spouse's Social Security retirement benefits are added together to get the value of all household Social Security retirement benefits.

The benefits received during the specified tax year (either 2010 or 2012) were calculated in the following steps. If the respondent worked during 2010 (2012) but was retired by the time of the CogEcon 2011 (CogEcon 2013) survey, I impute the number of months that the pension was received during 2010 (2012). I use self-reported retirement age, current employment status, and number of weeks working in 2010 and 2012 to impute the number of months the respondent received pension payments in 2010 and 2012. I assume the respondent received the pension for non-work weeks. I assign zero months for respondents who worked all of 2010 (or 2012) but then retired by the time they completed the survey during 2011 (or 2013). I assume they received for 12 months if they were over 65 years old in January 2010 (or January 2012) with the exception of one respondent below 65 but is clearly receiving pension for the whole year.¹³ The imputed number of months is multiplied times the reported monthly amount to get the imputed pension received during the specified year. I follow the same steps to calculate the spouse's pension amount. The respondent and spouse's pensions are added together to get the value of all household pensions.

My preferred measure of gross income combines information from these two preceding measures. When computing tax rates for respondents we want to use the components of income in order to account for the fact that not all income sources are taxed in the same way. However, the income components are not exhaustive and therefore we would underestimate total income if we took these measures at face value.

¹²One respondent (sampid=7017080020) reported \$100,000 in 2011 and then \$10,000,000 in 2013. This R is over 80 years old and generally provided reasonable numbers; e.g., their SS benefits were \$855 a month in one year and \$850 a month in the other. Aggregated parts of income summed to \$34,400 in 2011 and \$53,700 in 2013, suggesting that the "other income" would, if anything, be larger in 2011 than in 2013, which would mean the total income reported in 2011 should be larger than in 2013. Hence, I divided the reported income by \$100.

¹³The respondent with sampid=7003600010 was 60 years old and claimed to be retired at 55 but was still working 52 weeks in 2010. However, as he also reported collecting a pension in an earlier wave of the survey, I determine that he was likely to have been receiving said pension for the entire year.

$$GrossIncome = Inc_{sum} + \widehat{OthInc}$$

where

$$\widehat{OthInc} = Inc_{tot} - Inc_{sum}$$

I assume $\widehat{OthInc} = 0$ if $Inc_{tot} < Inc_{sum}$ or if the respondent has less than fifty thousand dollars in financial assets outside of tax-advantaged retirement accounts. I also use total income when the components are missing.

Respondents who are partnered and planning a financial future together were asked to give information about their own and their partner's income. However, for tax purposes I want only their own income and not that of their partner. In such cases I use only the respondents amounts for (i) thru (iii), half of the retirement distributions (iv) and half of this residual.

B.1.2 Measuring Adjusted Gross Income (AGI)

I use NBER's Taxsim tax calculator to determine adjusted gross income for each respondent.¹⁴ This program allows me to distill taxable Social Security benefits from the total amount. The survey did not collect information about above-the-line deductions, so the difference between gross income and AGI will be the same as going from gross income to total income.

$$AGI = \text{Total Income} - \text{Above Line Deductions}$$

Total Income is gross income minus tax exempt income. This includes tax exempt interest, qualified dividends, the part of Social Security benefits, pensions and annuities that are not taxable.

Above-the-line deductions include IRA contributions, student loan interest, self-employed health insurance contributions, etc. For traditional plans these withdrawals are subject to taxation and therefore should be included in our measure of taxable income. I do not know whether these withdrawals were from traditional or Roth-designated accounts. I assume that distributions come from traditional accounts. The rules concerning distributions from traditional versus Roth designated accounts would make it more likely that distributions are from traditional accounts. This is corroborated by the fact that over half of respondents claimed they took out the required minimum distribution.

Tax treatment of Social Security retirement benefits:

- Depends on total income and marital status
- Generally, if Social Security benefits were only income, benefits are not taxable.
- If the household received income from other sources, benefits will not be taxed unless modified adjusted gross income is more than the base amount for the household's filing status.
- Quick computation:
 - Add one-half of the total Social Security benefits the household received to all other income, including any tax exempt interest and other exclusions from income.

$$\text{Combined Income} = AGI + \text{Nontaxable Interest} + \text{Half of SS benefit}$$

- Then, compare this total to the base amount for the household's filing status. If the total is more than the household's base amount, some of the benefits may be taxable.
- If Combined Income < \$32,000 for married filing jointly, none of benefits taxable
- If Combined Income < \$25,000 for single, none of benefits taxable

¹⁴The citation for the NBER Taxsim calculator is Feenberg and Coutts (1993).

B.1.3 Measuring Taxable Income (TI)

Taxable Income (TI) is the amount of income that is actually subject to federal income taxation. It is adjusted gross income minus all deductions and exemptions:

$$\text{Taxable Income} = \text{AGI} - \text{Deductions} - \text{Exemptions}$$

The amount of exemptions comes from marital status and the number of dependents that reported in the CogEcon survey. In tax year 2010 the personal exemption was \$3,650 for each qualified dependent. For tax year 2012 this amount increased to \$3,800 per dependent exemption.

I do not know whether respondents itemized deductions and for simplicity assume all respondents receive the standard deduction. The standard deduction for households who filed as single was \$5,700 in 2010 and \$5,950 in 2012. The standard deduction for taxpayers who were married filing jointly was \$11,400 in 2010 and \$11,900 in 2012. Taxpayers who turned 65 on or before January 2nd of the subsequent year (either 2011 or 2013) were eligible for an additional deduction. The standard deduction increases by \$1,400 in 2010 and \$1,450 in 2012 for single filers, and increases by \$1,100 in 2010 and \$1,150 in 2012 for each elderly member for those who are married filing jointly. I use the information on birth date for the respondent and his or her spouse (described above) to determine the number of members over age 65.

B.1.4 Limitations

Tax liability also gets adjusted by tax credits, and refundable credits can generate a net transfer to the household, as is typical with the Earned Income Tax Credit (EITC). Tax credits are largely ignored in my analyses. First, we do not have information about which credits people filed for. Second, the sample consists of older and higher income taxpayers, many of whom no longer have children living with them, so will not be eligible for many of the largest tax credits (e.g., the EITC).

It is important to recall that the individual components and reported total household income are reported over different time horizons. There is no obvious way to improve upon this ad hoc approach.

Another drawback of using total income 1 is that it is given for the “past 12 months,” which do not coincide with the tax years that I focus on. CogEcon 2011 was fielded at the end of 2011, so the total income is for 2011 rather than 2010.

Because surveys were fielded in late 2011 (and 2013), C2 provides 2011 (and 2013) income rather than 2010 (and 2012) income. It is important to note that this question asks about income in the past 12 months but for the purpose of the tax rate measurement we want income in the preceding calendar year.

B.2 Description of variables

- Used paid tax preparer (in 2011): Indicator variable equal to one if the respondent used a paid tax preparer the last time she filed a tax return. Assigned yes if any of the following were selected as answers to question C28: Commercial tax preparation company (like H&R Block); Financial planner or advisor; Accountant; Lawyer. The exact wording and answer options are the following: The last time you or your spouse filed a tax return, did you receive assistance or use tax software? If yes, please check all that apply. If no, please check “Did not receive assistance or use software.” If you are not married, please answer only for yourself. [Family member, friend or colleague; Tax software (like TurboTax); Commercial tax preparation company (like H&R Block); Financial planner or advisor; Accountant; Lawyer; Did not receive assistance or use software]
- Cognitive ability is capture by two variables: (i.) standardized number series score, (ii.) standardized verbal analogies score. These are both measures of fluid intelligence and come from the face-to-face CogUSA cognitive assessments.
- Financial sophistication is the average of the CogEcon financial sophistication scores computed from CogEcon 2008 and CogEcon 2009.¹⁵ The score is standardized after taking the average score from

¹⁵See CogEcon documentation for details about their construction.

CogEcon 2008 and CogEcon 2009 (when both were available).

- State income tax: Indicator variable for living in a state that has a state income tax
- Financial occupation dummy variables: For the occupation deals with finance/investment, defined as working in occupations that do rates of return, cost-benefit analysis and investment decisions. Budgeting occupations are defined as occupations with resource management or budgeting content. See McFall et al (2011) for details about these occupational measures.
- Log(wealth): This is the log of total household wealth. Total household wealth is the average from wealth measured in CogEcon 2011 and CogEcon 2013. It includes financial, housing and miscellaneous assets and debts. Of the sample of 348, there are 11 respondents for whom we have no information about their wealth. It is implausible that they had zero wealth in both CogEcon 2011 and CogEcon 2013. Because zeros are problematic when using logs, I assign the mean of log wealth (12.57) to these 11 respondents with zero wealth in both CogEcon 2011 and CogEcon 2013.
- Male: Indicator variable equal to one for men.
- Nonwhite: dummy variable equal to one if race is non-white or hispanic, based on the CogUSA race variable.
- Education is a categorical variable: less than 12 years, 12-16 years and over 16 years; the category for less than 12 years is excluded.
- Age/10: Age on date the survey was completed. See CogEcon documentation for more information.
- Working: indicator for being employed at all during 2010 (or 2012).

C Maximum Likelihood Estimation

C.1 Estimation

Estimation of $\theta = (\mu_{y^*}, b_a, b_m, \sigma_{y^*}, \sigma_e, \sigma_a, \sigma_m, \rho_a, \rho_m, \rho_{am})$ is done using the method of maximum simulated likelihood. I use Stata's built-in maximum likelihood estimation routines, which utilize particularly robust optimization algorithms (Kolenikov, 2001). I use standard Monte Carlo methods to simulate the likelihood function by taking draws over the two-dimensional income measurement error. I reduce computational burden by using Monte Carlo integration with Halton draws to simulate the likelihood function.

This approach is straightforward to understand but does not do as well at maintaining the smoothness of the likelihood function. I can achieve convergence using an estimated Hessian, but getting my likelihood function to converge when using the NR algorithm would require a more complicated simulation method. Using simpler numerical integration with the quasi-Newton methods works well at much lower computational cost.

C.2 The simulation algorithm

The estimation strategy uses the following algorithm to simulate the likelihood function for each observation i . It requires integrating over the bivariate normal income measurement error distribution.

1. I use Stata plug-in "mdraws" to generate K quasi-random Halton draws (from the unit interval) of a two-dimensional vector for each individual $i = 1, \dots, N$. Define these as $\eta_{i(k)} = (\eta_{i(k),1}, \eta_{i(k),2})'$ for $k = 1, \dots, K$, where $\eta_{i(k),1}$ is associated with income error in wave 1 and $\eta_{i(k),2}$ is associated with income error in wave 2. The vector $\eta_{i(k)} = (\eta_{i(k),1}, \eta_{i(k),2})'$ is only generated once, before the estimation routine. These draws are made before estimation to reduce the chance of convergence failures that arise because of noise created by draws at each iteration in simulation-based estimation. Train (2003) argues that Halton draws are more effective for MSL estimation than pseudo-random draws because they provide the same accuracy with fewer draws. The draws for one observation tend to be negatively correlated with those of the previous observation, which reduces error in the simulated likelihood function. Because each observation gets its own draws, changing the order of observations in the data can impact estimation results.
2. These draws are used to generate K two-dimensional income errors for each individual. These are defined as $e_{i(k)} = (e_{i(k),1}, e_{i(k),2})'$ for each individual $i = 1, \dots, N$ with $k = 1, \dots, K$, where $e_{i(k),1}$ is the income error in wave 1 and $e_{i(k),2}$ is the income error in wave 2. Because income errors in waves 1 and 2 are uncorrelated, these draws $e_{i(k),w}$ are calculated using the inverse transform method

$$e_{i(k),w} = \tilde{\sigma}_e \cdot \Phi^{-1}(\eta_{i(k),w})$$

where $\tilde{\sigma}_e$ is the latest estimate of the standard deviation of the measurement error.

3. For each draw $e_{i(k),w}$, calculate the implied true income $Y_{i(k),w}^*$ as

$$Y_{i(k),w}^* = \exp(y_{i(k),w} - e_{i(k),w})$$

This is transformed into adjusted gross income $AGI_{i(k),w}^* = s_{i,w} \cdot Y_{i(k),w}^*$, where $s_{i,w}$ is the share of observed adjusted gross income as a fraction of observed total income. This is used to calculate true taxable income

$$I_{i(k),w}^* = \max \left\{ AGI_{i(k),w}^* - D_{i,w}^* - E_{i,w}^*; 0 \right\}$$

with known deductions $D_{i,w}^*$ and exemptions $E_{i,w}^*$. Taxable income is then plugged into the known tax functions $M_w(\cdot)$ and $T_w(\cdot)$ to calculate the true tax rates associated with this realization of income,¹⁶

¹⁶NOTE: To improve the numerical stability of the optimization routine, tax rate variables are scaled by ten to put them between 0 and 10, which is closer to the scale for log income.

which implies values of the tax errors of

$$\begin{aligned}\varepsilon_{i,w}^m &= m_{i,w} - M_w \left(I_{i(k),w}^* \right) \\ \varepsilon_{i,w}^a &= a_{i,w} - \frac{T_w \left(I_{i(k),w}^* \right)}{Y_{i(k),w}^*}\end{aligned}$$

Define $z_w = \frac{y_w - \mu_{y_w}^*}{\sqrt{\sigma_{y_w}^{*2} + \sigma_e^2}}$ and the joint density of $(y_{i,1}, y_{i,2})$ can be written as

$$f(\mathbf{y}_i) = \frac{1}{2 \cdot \pi \cdot \sigma_y^2 \cdot \sqrt{1 - \rho_y^2}} \cdot \exp \left(\frac{z_1^2 + z_2^2 - 2 \cdot \rho_y \cdot z_1 \cdot z_2}{2 \cdot (1 - \rho_y^2)} \right)$$

Calculating $\Pr(\mathbf{r}_i | \mathbf{y}_i, e_{i(k)})$ is simplified by using Bayes law to write this joint distribution as a marginal distribution multiplied by a conditional distribution

$$\Pr(\mathbf{r}_i | \mathbf{y}_i, e_{i(k)}) = \Pr(\mathbf{r}_{i,2} | \mathbf{y}_{i(k)}^*, \mathbf{r}_{i,1}) \cdot \Pr(\mathbf{r}_{i,1} | \mathbf{y}_{i(k)}^*)$$

I write the distribution of the tax rate responses in the second wave conditional on the tax rate responses in the first wave, because joint normality of the tax rate responses implies that the conditional distributions are also normally distributed.

$$\mathbf{r}_2 | \mathbf{y}_i^*, \mathbf{r}_1 \sim N(\mathbf{b}_{\mathbf{r}_2 | \mathbf{r}_1}, \Sigma_{\mathbf{r}_2 | \mathbf{r}_1})$$

The multivariate normal distribution has conditional distributions that are also normally distributed, so

$$\mathbf{b}_{\mathbf{r}_2 | \mathbf{r}_1} = \mathbf{b}_{\mathbf{r}_2} + \Sigma_{21} \Sigma_{11}^{-1} (\mathbf{r}_{i,1} - \mathbf{b}_{\mathbf{r}_1}) = \begin{bmatrix} b_{a_2} + \frac{\rho_a \sigma_a}{1 - \rho_{am}^2} \cdot \left[\left(\frac{a_1 - b_a}{\sigma_a} \right) - \rho_{am} \cdot \left(\frac{m_1 - b_m}{\sigma_m} \right) \right] \\ b_{m_2} + \frac{\rho_m \sigma_m}{1 - \rho_{am}^2} \cdot \left[\left(\frac{m_1 - b_m}{\sigma_m} \right) - \rho_{am} \cdot \left(\frac{a_1 - b_a}{\sigma_a} \right) \right] \end{bmatrix}$$

and

$$\Sigma_{\mathbf{r}_2 | \mathbf{r}_1} = \Sigma_{22} - \Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12} = \begin{bmatrix} \sigma_a^2 \cdot \left(1 - \frac{\rho_a^2}{1 - \rho_{am}^2} \right) & \rho_{am} \sigma_a \sigma_m \cdot \left(1 + \frac{\rho_a \rho_m}{1 - \rho_{am}^2} \right) \\ \rho_{am} \sigma_a \sigma_m \cdot \left(1 + \frac{\rho_a \rho_m}{1 - \rho_{am}^2} \right) & \sigma_m^2 \cdot \left(1 - \frac{\rho_m^2}{1 - \rho_{am}^2} \right) \end{bmatrix}$$

4. Compute the likelihood function $L_i^k(\theta)$ given income error $e_{i(k)}$ associated with draw k ,

$$L_i^k(\theta) = f(\mathbf{r}_{i,2} | \mathbf{y}_{i(k)}^*, \mathbf{r}_{i,1}) \cdot f(\mathbf{r}_{i,1} | \mathbf{y}_{i(k)}^*) \cdot f(\mathbf{y}_i)$$

The densities $f(\mathbf{r}_{i,1} | \mathbf{y}_{i(k)}^*)$ and $f(\mathbf{r}_{i,2} | \mathbf{y}_{i(k)}^*, \mathbf{r}_{i,1})$ are written similarly to the density $f(\mathbf{y}_i)$.

5. Repeat this algorithm K times, and compute the simulated likelihood function as the average of the K error probabilities over the K draws of the vector of errors $e_{i(k)}$,

$$\check{L}_i(\theta) = \frac{1}{K} \sum_{k=1}^K L_i^k(\theta)$$

6. Use $\check{L}_i(\theta)$ to estimate the maximum simulated likelihood parameters.

C.3 Gradient-based optimization methods

Stata's default method of finding the maximum uses a modified Newton-Raphson (NR) algorithm. While the NR algorithm is typically effective, it requires calculation of the Hessian matrix, which can be particularly time consuming when maximizing a simulated likelihood function.

Stata's default optimization routines have three requirements for declaring convergence. First, the tolerance for changes in the coefficient vector from one iteration to the next or the tolerance for changes in the likelihood from one iteration to the next must be sufficiently small. Second, the second criterion is having a sufficiently small gradient relative to the Hessian (`nrtolerance()`). This is formally based on $g'\mathbf{H}^{-1}g$, where the gradient g and Hessian matrix \mathbf{H} are calculated at the parameter vector $\hat{\theta}$. Finally, the Hessian must be concave.

I use a different convergence criterion when simulating the likelihood function. When using algorithms `bhhh` and `bfgs`, the convergence criterion `qtolerance()` in Stata 13 checks the modified Hessian matrix calculated at the parameter vector. The convergence criterion does not require calculating the Hessian matrix and therefore reduces computational burden. See Stata 13 Manual for more details.

When using maximum simulated likelihood methods, there is a tradeoff between unbiasedness and computational burden. The main issue with MSL arises because the log operator can induce bias in the MSL estimator. Even if $\check{L}_i(\theta)$ is an unbiased simulator of $L_i(\theta)$, the non-linear log transformation can make $\ln \check{L}_i(\theta)$ a biased simulator of $\ln L_i(\theta)$, which translates into bias in the MSL estimator. Raising the number of draws used to simulate individual-level likelihood functions improves the accuracy of the results by smoothing the objective function (by increasing simulation accuracy). This bias diminishes as more draws are used.¹⁷ However, estimation is numerically intensive and convergence can be slow when the number of draws is large and especially when the number of equations is large.

In light of this tradeoff, I select an optimization routine and convergence criterion in order to maintain robustness without unnecessary computation burden. My optimization routine switches between the Berndt-Hall-Fletcher-Powell (BHHH) algorithm (for 5 iterations) and the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (for 10 iterations). Using these quasi-Newton methods reduce the computational burden relative to using the standard Newton-Raphson (NR) algorithm by replacing the computed Hessian matrix with something easier to calculate.

There are two reasons I use the BHHH algorithm rather than the standard Newton-Raphson (NR), both related to the BHHH using the score to approximate the Hessian rather than computing it directly. First, computing the Hessian matrix is computationally intensive, as it requires computing the second derivatives at every iteration. In the context of simulated likelihood estimation this greatly reduces computational efficiency. Second, simulating the likelihood function can result in an inverted Hessian that is not positive definite, which creates convergence problems. In particular, simulating the likelihood function can generate non-concavities of the log-likelihood function or lead to the numerical issues from having a computed Hessian that is a poor approximation to the analytical one.

¹⁷In fact, if the number of draws K rises at any rate with the number of observations N , then MSL is consistent, and if K rises faster than \sqrt{N} , then MSL is asymptotically equivalent to ML (Train, 2003, p. 273-4). Increasing the number of draws with the sample size can reduce simulation bias to negligible levels, which can be ensured by having the ratio of the number of draws to the square root of sample size sufficiently large (Hajivassiliou and Ruud, 1994, 2416-2419).

D Imputation derivations

D.0.1 Derivation of conditional expectations

In this section I work through the steps needed to impute true income

$$\widehat{y_{i,w}^*} = \mathbb{E} \left(y_{i,w}^* \mid \widehat{\theta}, m_{i,w}, a_{i,w}, y_{i,w} \right) = \int y_{i,w}^* \cdot \Pr \left(y_{i,w}^* \mid \widehat{\theta}, m_{i,w}, a_{i,w}, y_{i,w} \right) \cdot dy_{i,w}^*$$

Everything is conditional on parameter vector $\widehat{\theta}$ and this derivation ignores this term for simplicity. From Bayes' Theorem,

$$\Pr \left(y_{i,w}^* \mid m_{i,w}, a_{i,w}, y_{i,w} \right) = \frac{\Pr \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_{i,w}^* \right) \cdot \Pr \left(y_{i,w}^* \right)}{\Pr \left(m_{i,w}, a_{i,w}, y_{i,w} \right)}$$

The denominator $\Pr \left(m_{i,w}, a_{i,w}, y_{i,w} \right)$ is the unconditional probability and can be re-written by integrating over the distribution of possible true income

$$\Pr \left(m_{i,w}, a_{i,w}, y_{i,w} \right) = \int f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_{i,w}^* \right) \cdot f \left(y_{i,w}^* \right) \cdot dy_{i,w}^*$$

It does not depend on the realization of the true values. The expected value can now be re-written as

$$\widehat{y_{i,w}^*} = \int y_{i,w}^* \cdot \frac{f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_{i,w}^* \right)}{W_i} \cdot f \left(y_{i,w}^* \right) \cdot dy_{i,w}^*$$

where

$$W_i = \int f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_{i,w}^* \right) \cdot f \left(y_{i,w}^* \right) \cdot dy_{i,w}^*$$

The integrals in this equation do not have a closed form but can be approximated by simulation. With M simulation draws y_k^* from the distribution of $y_{i,w}^*$ the approximation can be written as

$$\begin{aligned} \widehat{y_{i,w}^*} &= \int y_{i,w}^* \cdot \frac{f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_{i,w}^* \right)}{W_i} \cdot f_{y^*} \left(y_{i,w}^* \right) \cdot dy_{i,w}^* \\ &\approx \frac{1}{C_i} \cdot \sum_{k=1}^M y_k^* \cdot f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_k^* \right) \end{aligned}$$

where C_i is a normalization factor

$$C_i = \sum_{k=1}^M f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_k^* \right)$$

Take 100 draws of taxable income from $y_{i,w}^* \sim N \left(\widehat{\mu}_y, \widehat{\sigma}_y^2 \right)$, where $\widehat{\mu}_y$ and $\widehat{\sigma}_y^2$ were estimated above. Draws y_k^* are implicitly defined in terms of the percentage in the distribution of all values, which satisfies $p_k = \Phi \left(\frac{y_k^* - \widehat{\mu}_y}{\widehat{\sigma}_y} \right)$ where $p_k = k/100$ for $k = 1, \dots, 100$. Draws are computed as $y_k^* = \widehat{\mu}_y + \widehat{\sigma}_y \cdot \Phi^{-1} \left(p_k \right)$. Calculate $f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_k^* \right)$ for each observation i and for each draw y_k^* .

$$f \left(m_{i,w}, a_{i,w}, y_{i,w} \mid y_k^* \right) = f \left(m_{i,w}, a_{i,w} \mid y_k^* \right) \cdot f \left(y_{i,w} \mid y_k^* \right)$$

Plug in to the formula for the approximation

$$\widehat{y_{i,w}^*} \approx \frac{\sum_{k=1}^M y_k^* \cdot f(m_{i,w}, a_{i,w}, y_{i,w} | y_k^*)}{\sum_{k=1}^M f(m_{i,w}, a_{i,w}, y_{i,w} | y_k^*)}$$

Imputed income $\widehat{y_{i,w}^*}$ is transformed into taxable income,

$$\widehat{I_{i,w}^*} = \max \left\{ s_{i,w} \cdot \exp \left(\widehat{y_{i,w}^*} \right) - D_{i,w}^* - E_{i,w}^*; 0 \right\}$$

where $s_{i,w}$ is the share of observed adjusted gross income as a fraction of observed total income. This is plugged into the known marginal tax rate function $M_w(\cdot)$ to compute the marginal tax rate $\widetilde{m_{i,w}^*} = M \left(\widehat{I_{i,w}^*} \right)$ and in the tax liability function $T_w(\cdot)$ get average tax rate $\widetilde{a_{i,w}^*} = \frac{T_w \left(\widehat{I_{i,w}^*} \right)}{\exp \left(\widehat{y_{i,w}^*} \right)}$.

D.0.2 Predicting ATR and MTR directly

Because of the mapping from true income to tax rates, the above algorithm can also be used to predict average and marginal tax rates. I define the best predictor of true average tax rates in terms of the conditional expectation

$$\widehat{a_{i,w}^*} = \widehat{\mathbb{E}} \left[a_{i,w}^* | \widehat{\theta}, \mathbf{r}_i, \mathbf{y}_i \right] = \int a_{i,w}^* \cdot f \left(a_{i,w}^* | \widehat{\theta}, \mathbf{r}_i, \mathbf{y}_i \right) da_{i,w}^*$$

This computation is more intensive than the one for marginal tax rates. I use a change-of-variables formula to cast this in terms of unobserved income measurement error, numerically integrate over the distribution of income measurement error and then compute $f \left(a_{i,w}^* | \widehat{\theta}, \mathbf{r}_i, \mathbf{y}_i \right)$ using Bayes' Theorem.

This is plugged into the known tax function $T_w(\cdot)$ to calculate the true (implied) ATR $a_{k,w}^*$ and then the imputed ATR $\widehat{a_{i,w}^*}$:

$$a_{k,w}^* = \frac{T_w \left(I_{k,w}^* \right)}{\exp \left(y_{k,w} - e_{k,w} \right)}$$

$$\widehat{a_{i,w}^*} \approx \frac{\sum_{k=1}^M a_{k,w}^* \cdot f \left(m_{i,w}, a_{i,w}, y_{i,w} | y_k^* \right)}{\sum_{k=1}^M f \left(m_{i,w}, a_{i,w}, y_{i,w} | y_k^* \right)}$$

Similarly, for marginal tax rates, where $m_{i,w}^* = M \left(y_{i,w}^* \right)$ is a known deterministic function.

$$\widehat{m_{i,w}^*} \approx \frac{\sum_{k=1}^M m_{k,w}^* \cdot f \left(m_{i,w}, a_{i,w}, y_{i,w} | y_k^* \right)}{\sum_{k=1}^M f \left(m_{i,w}, a_{i,w}, y_{i,w} | y_k^* \right)}$$