

RAS Affordable Care Act Microsimulation Model

*Brian Erard (Brian Erard and Associates), Layne Morrison, Robert Mueller (IBM), and Emily Heys
and Brock Ramos (IRS, Research, Analysis, and Statistics)*

Introduction

Implementation of the Patient Protection and Affordable Care Act (ACA) changes the health insurance landscape in the U.S., affecting businesses and individuals. ACA creates new responsibilities for taxpayers, employers, and State and Federal Governments. It also enhances the role of the Internal Revenue Service (IRS) in its administration of subsidies to purchase health insurance and payments in lieu of meeting coverage requirements. In carrying out its new administrative responsibilities, the IRS will interface with the Department of Health and Human Services (HHS) and the new Exchanges to share information. Quantitative estimates of employer-sponsored health insurance (ESI) offers and individual decisions to obtain health insurance coverage under the new ACA regime are essential inputs to guide the development of IRS systems, business processes, forms, and outreach programs to meet new responsibilities for administering subsidies and payments under the Act. Even more fundamentally, IRS seeks to understand behavior surrounding health insurance and respond with informed program design and administration.

To understand the potential implications of ACA on agency volume and workload, IRS Research Analysis and Statistics (RAS) undertook three progressive research steps:

- **Review of the Literature.** In 2010, RAS began a review of the relevant health insurance literature within the context of ACA. This review explored the key behavioral decisions for individuals in obtaining health insurance and how ACA sought to influence them, as well as driving factors in the success of credits and outreach.
- **Review Third-Party Models and Plausible Estimates.** Given the complexities of healthcare policy, several microsimulation models of the United States healthcare system have been developed. These models allow researchers to estimate outcomes based on employer, individual, and family behaviors as they relate to the decision to offer and take-up health insurance. Government agencies, including the Congressional Budget Office (CBO) and the Treasury Office of Tax Analysis (OTA), as well as private organizations such as The Lewin Group, the Urban Institute (UI), and RAND, have built major simulations over years. IRS explored these models and leveraged their estimates for initial planning purposes.
- **Design and Develop a Microsimulation Model.** In 2012, RAS collaborated with IBM and Brian Erard & Associates to build a microsimulation model for tax administration planning purposes. This model allows IRS to model administrative impacts of the new legislation and assess likely outcomes and workload volumes, supported by studies and research. This model fulfilled a unique role in two respects: first, it allowed IRS to perform ‘what if’ analysis; and second, it reported results at the tax return level. The RAS Model provides a return-based data file representing nonelderly taxpayers and their characteristics, including imputed health insurance and customer service usage. The model outputs statistics on the number of returns in different health insurance status categories (Employer Sponsored Insurance (ESI), Private, Public, Uninsured, and Exchange)—both before and after implementation of key ACA tax provisions—the number of employers offering ESI post-reform, the number of returns reporting the individual shared responsibility payment (ISRP), and the number of returns reconciling the premium tax credit (PTC).

The estimates will better prepare IRS to meet the workload demands, informing future implementation, budget, and resource allocation decisions.

This paper describes the model and foundational studies undertaken prior to design. We provide some short background on health insurance and its representation in microsimulation models. Then we discuss the model that was built for IRS and its outputs.

Microsimulation Models for Health Insurance

Over the last two decades, several microsimulation models have been built to describe the U.S. health insurance ecosystem. Starting from a microsimulation foundation with either tax or health data, these models predict future health-care coverage under ACA. Two are maintained by government agencies: The Congressional Budget Office's Health Insurance Simulation Model¹ (hereafter called simply CBO) projects impacts of current legislation, while Treasury Office of Tax Analysis (OTA) serves as the source for executive branch estimates. Both of these models focus on the implications of the law for public revenue. Alternative third-party microsimulation models also include those developed by Jonathan Gruber at MIT, the Urban Institute's Health Information Policy Simulation Model (which we refer to as UI), The Lewin Group's Health Benefits Simulation Model (Lewin);² and RAND's COMPARE Model (RAND).³

Microsimulation models are a popular technique to examine public policy impacts. They are classified by their focus on individual, representative records, to which rules and behaviors are applied. Figure 1 summarizes the four main components of models, as adapted from Abraham, 2012.⁴

FIGURE 1. Health Policy Microsimulation Model Components



Under this project, multiple third-party models were compared and contrasted across these dimensions. While there are similarities in terms of the underlying data sources and assumptions, the models also have important differences. These variations reflect differences in a number of design choices and assumptions, such as: the choice of base data, assumptions about exchange viability and availability (policy decisions as yet to be determined), participation parameter assumptions, and elasticity or utility-based frameworks. Authors exploring the prominence of analytic models in the health policy field suggest that such large variance in the models' point estimates is common. Glied and Tilipman (2010)⁵ conclude that models, reflecting the uncertainty of source surveys and modeling decisions, generally have proven accurate within 30 percent of actual program consequences. Despite this rather high level of uncertainty, the authors conclude that the models can still be quite helpful to policymakers and administrators. Namely, they provide an understanding about the range of results, while generating insights into key dynamics and the leverage of various assumptions.

In 2012, IRS began the development of its own administrative model for two primary reasons. First, the individual-based (rather than tax return-based) output and policy focus of existing models was not well suited for understanding the impact of changes in employer and taxpayer behavior on IRS workload volumes in administering tax-related ACA provisions (such as the PTC and ISRP) and assisting taxpayers with meeting their responsibilities under these provisions. Second, it was apparent from the existing models that plausible alternative behavioral assumptions sometimes led to widely different predicted outcomes. IRS therefore sought a simulation model that could perform 'what if' analyses to understand the potential impacts of alternative scenarios on key workload parameters. The resulting RAS-ACA Model is a microsimulation model that provides a flexible platform to support the IRS for future years as actual data is acquired.

¹ CBO technical documentation: <http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/87xx/doc8712/10-31-healthinsurmodel.pdf>.

² Lewin report publications can be found at <http://www.lewin.com/publications/?published=anytime&expertiseid={99087FB4-3394-4446-9BAE-980833611DE7}>.


³ RAND COMPARE model related publications can be found at <http://www.rand.org/health/projects/compare/publications/coverage.html>.

⁴ Adapted from Abraham, J. (2012). Predicting the Effects of the Affordable Care Act: A Comparative Analysis of Health Policy Microsimulation Models, State Health Reform Assistance Network Policy Brief.

⁵ Glied, Sherry and Nicholas Tilipman. Simulation Modeling of Health Care Policy. Annual Review of Public Health. 2010. n.p. <http://www.biomedsearch.com/nih/Simulation-modeling-health-care-policy/20235853.html>.

In designing the model, IRS relied on the best practices of existing models without replicating unnecessary elements. The final design is summarized across the four dimensions, as indicated in Figure 2. Each element is discussed in this paper.

FIGURE 2. IRS RAS-ACA Microsimulation Model Components



Data

Microsimulation involves studying the interaction of many complex, interrelated variables within one data file. To simulate health insurance policy, models rely on several types of data: demographic/population data; employer offering information; insurer offerings/premium prices and coverage, as well as medical expenditures. All models must integrate or match data from several sources since all of the variables required are seldom available solely from one source. Our model is simpler in that we focus on just three outcomes (exchange coverage, any other coverage, and no coverage) without a detailed behavioral model, and therefore we do not include medical costs or premium elasticities.

A critical feature of a microsimulation model is the quality of its input data source. This rich input file represents real, record-by-record data with a variety of fields, serving as the foundation onto which imputations and calculations are performed to create the final file. Successful models strike a careful balance between many of the desirable features of an input file:

- Representative of the underlying population—ideally at granular levels (ex., State-level, county-level) and subgroups of interest;
- Rich in relevant descriptive and predictive variables; and
- Usually a sample, for reasonable run times (small and tractable).

Common choices for the core data file are nationally representative data files such as the Current Population Survey's Annual Social and Economic Supplement (CPS ASEC—used by Urban Institute and Jonathan Gruber) or the Survey of Income and Program Participation (SIPP—used by CBO and RAND).

We focus on taxpayers, so we use IRS Compliance Data Warehouse (CDW) records as our base. To represent the complexity of the healthcare landscape, we select three key administrative data sources to capture the important linkages between employers and employees. The first source is employer records, public and private, from payroll records (Forms 941, 943, 944). Government entities file a payroll tax form and thus are included in our sample, although they do not file an annual tax return. The latest data year available was Tax Year 2010, which is used consistently for all data pulls. The employers in these files are linked to employees through a match to our second key data source: W-2 records. The employees (along with the information extracted about their employers) are then matched to our third key data source: individual income tax returns. Through this process, we match information from 131 million individual non-dependent returns, 5.1 million employer returns, and 216 million W-2 filings. Our starting population of individuals comes from tax returns, which restricts the population to filers. Nonfiling employees are identified by matching W-2 records to individual income tax returns, yet we focus only on filers in the model.

Sampling: To select a sample of employers and their employees, we undertake a two-stage cluster sampling approach. Under this approach, we begin by drawing a stratified random sample of employers from the relevant population of either private or government employers; very large employers are sampled with certainty. In the second stage we draw a sample of employees from each of the sampled employers. We include an increasing number of employees for the sample as the size of the employer increases. We sample all workers from smaller establishments in an effort to

carefully model participation in ESI and other health insurance options for such workers. In the case of larger establishments, we felt that the increased sample sizes would be sufficiently large to represent health insurance participation among all employees. Two main stratifiers are used in selecting employers; size and employer type. Size of the firm is derived using a combination of two measures: average employees (as reported on payroll forms) and a count of W-2 employees, which represent anyone employed throughout the year. Our sampling strata include 12 size categories and 3 employer types: Private, Federal Government, and State & Local Government. Once the selection of employees is complete, we match them to their individual income tax returns (when present). We then supplement our data file by including employer information for the spouse (where the primary filer is the employee matching to an employer from our sample) and for the primary filer (where the spouse is the employee matching to an employer from our sample). Finally, we exclude returns from our sample in which the primary filer (and the spouse, in the case of joint returns) age 65 or over.

To account for those without wage income (those taxpayers who did not have an employer), we draw a stratified random sample from the sub-population of individual income tax returns of nonelderly filers with no reported wage earnings. The strata are based on the Federal poverty level (FPL) measure of income, the presence of earnings from self-employment, and the presence of pension or social security earnings. We oversample those with FPL in the PTC range, especially those with retirement or self-employment income. As the majority of these individuals do not have ESI, they are likely candidates for the PTC.

TABLE 1. Combined Sample of Employers, Individuals and Returns

Employer Type	Number of Employers	Number of Individuals	Number of Returns
Private	19,052	360,085	–
Government	826	76,424	–
Nonwage Earners	–	–	33,380
Total	19,878	436,509	33,380

Imputations for Health Insurance: In the future, IRS will receive information from all employers regarding their ESI offer status. For the current project, however, it is necessary to impute this status.

Our approach relies on tabulations of ESI offer status by employer characteristics in the Medical Expenditure Panel Survey Insurance Component (MEPS-IC) for 2010. In the case of private employers, the MEPS-IC provides a breakdown of ESI offer rates by establishment size and industry. By applying these offer rates to the overall counts of employers within each size and industry category, we arrive at the target number of employers in each category to be assigned to offer ESI. In the case of State and local government employers, the MEPS-IC provides separate breakdowns of ESI offer rates by employer size and by Census region. Within each category, we begin by assigning ESI offer status to employers that appear to provide ESI based on their filing of a report of an employee health insurance plan on Form 5500, claim of a health insurance premium credit on Form 8941, or claim of a deduction (over \$100) on their income tax return for employee benefit program expenditures. We then randomly assign additional employers within that category to offer ESI to approximately achieve the target number of offering employers. We assign all Federal employers in our sample to offer ESI.

In addition to imputing ESI offer status for employers in our sample, we also impute current health insurance status to each return in our sample. For a small portion of the returns in our sample, we are able to predict the likely health insurance status based on claims of a health savings account or self-employed health insurance expense deduction. For the remaining returns, we have developed and implemented an econometric methodology for our imputations based on the 2011 CPS ASEC (which contains information about earnings in 2010).

For this project, we assign a return-based measure of health insurance status based on the status of the primary taxpayer and, in the case of joint returns, the spouse. The CPS ASEC includes an indicator for the imputed individual income tax return filing status for each household in the sample (one of the imputed statuses is nonfiler). We restrict our attention to imputed filers who are under age 65 (or in the case of married joint-filing status, who have a spouse under age 65). To derive our return-based measure of health insurance status, we begin by identifying the health insurance status of the primary filer and, in the case of joint returns, the spouse (some hierarchies are employed to resolve

multiple forms of health insurance) and whether ESI is attributed to the primary or spouse. This provides us with an imputed indicator for four categories of health insurance status to our return-level database:

1. ESI
2. Private
3. Public
4. Uninsured

We estimate separate models for wage and nonwage earners.

Wage Earners: For wage-earner returns, we begin by estimating a model for the probability of ESI coverage status among nonelderly ASEC returns with wage earnings. Although some taxpayers who currently work for an employer that does not offer ESI still receive ESI coverage (typically, through a former employer), ESI coverage tends to be much more likely among workers whose employers do offer ESI. Unfortunately, the ASEC does not identify whether the respondent's employer offers ESI. Rather, we only observe whether ESI is received. To address this problem, we employ a bivariate probit model with partial observability:

$$\begin{aligned} P_O^* &= \beta_O' x_O + \epsilon_O \\ P_{R|O}^* &= \beta_R' x_R + \epsilon_R \\ \Pr(R|\bar{O}) &= 0.20, \end{aligned}$$

where P_O^* is a latent variable for the propensity of the employer to offer ESI, $P_{R|O}^*$ is a latent variable for the conditional propensity for a return to have ESI coverage status given that the employer offers ESI, and $\Pr(R|\bar{O})$ represents the conditional probability that a return has ESI coverage status given that the employer does not offer ESI. This latter conditional probability is set equal to 0.20, which is roughly consistent with the likelihood of a nonwage earner return having ESI status in the ASEC return population. In this model, the x terms represent vectors of explanatory variables, the β terms represent vectors of coefficients to be estimated, and the ϵ terms are random disturbances assumed to be normally and independently distributed. We do not observe the latent variables, and we do not observe whether the employer offers ESI. Rather, we observe only an indicator for whether a return has ESI coverage status; this is the source of partial observability. Although this model can be estimated using only the ASEC data, accumulated experience with partial observability models suggests that the performance of such an approach in identifying the propensity for an employer to offer ESI is likely to be subpar. Consequently, we have employed an alternative approach that incorporates additional information about the likelihood of an ESI offer.

Recall that the imputed measure of employer ESI-offer status in our return-level database was derived at the employer level from a breakdown of ESI offer rates in the MEPS-IC for such factors as employer size, industry, and whether the employer is private or government. It was then established as a return-based measure by linking the employers and their assigned ESI offer status to individual returns in our database. Rather than estimate the employer ESI-offer probit specification

$$P_O^* = \beta_O' x_O + \epsilon_O$$

in our above model directly from the ASEC, we instead estimate this probit equation using the subsample of wage-earner returns in our return-level database. The dependent variable in our specification is the imputed ESI offer status indicator, and the explanatory variables include interactions between employer size and an indicator for a government employer and interactions between employer size and industry category. Effectively, our probit model serves as a convenient way of translating the original MEPS-IC tabulations of ESI-offer rates by employer size and employer type into a formula for predicting the likelihood that an employee works for an employer who offers ESI (based on the size and type of his or her employer).

Estimation of our specification yields an estimate of vector β_O , which is then substituted for the actual value β_O in our full bivariate probit model with partial observability. The remaining parameters of the model (in vector β_R) are estimated from the ASEC sample using the method of maximum likelihood. The explanatory variables x_R in our specification for the conditional probability that a return has ESI offer status given that the employer offers ESI include

the natural log of the age of the employee, the natural log of the number of child dependents, and indicators for part-time/full-time status, dual wage-earner returns, gender, married joint-filing status, the interaction between gender and married-joint filing status, the presence of self-employment earnings (Schedule C or Schedule F), the presence of unemployment insurance, and State of residence.

The results of the above estimation strategy provide us with predictive equations for imputing ESI coverage status to returns in our return-level data base. The next step is to develop additional predictive equations for imputing the other three health insurance coverage categories (private, public, and uninsured). To estimate the likelihood of a return falling into one of these three categories, we estimate a multinomial logit model using the subset of returns in the ASEC sample from these categories. The explanatory variables include all of the regressors in x_R above (which are used in the prediction of ESI coverage) as well as an indicator for whether the taxpayer is a local, State, or Federal Government employee.

Nonwage Earners: In the case of nonwage earners, a multinomial logit model for all four health insurance coverage statuses is estimated using the ASEC nonelderly nonwage earner subsample. For this analysis, the explanatory variables include the natural log of FPL, the natural log of age, the natural log of the number of dependent children, and indicators for nonjoint returns with male filers, married joint returns, the presence of self-employment earnings, the presence of unemployment insurance, the presence of Social Security benefits, the presence of pension earnings, and State of residence.

Imputations for Taxpayer Services: The complexity of ACA will drive taxpayers to use IRS services in greater numbers. However, IRS faces two challenges in estimating future service usage due to ACA. First, like many of the behaviors under ACA, there is little or no prior data from which to estimate marginal increases. Secondly, with respect to services, IRS captures identifying information related only to account issues. Consequently, the presence or absence of usage of the major service channels by specific taxpayers cannot be identified in most cases. In order to address this second challenge, prediction equations for taxpayer service usage were developed based on the Taxpayer Experience Survey (TES) and then applied to impute service usage for each return in the model input file. This was done for four service channels—phone, Taxpayer Assistance Center (TAC), mail and email.⁶

Our preliminary analysis of the four survey indicators in the TES indicated that the usage of the four categories of service was not independent. In other words, the likelihood of one service category being used depended on whether other service channels were used. It was therefore desirable to employ a statistical model that: (1) allowed for the possibility of a taxpayer using multiple service channels; and (2) accounted for dependencies across service channels. This led to the selection of a multivariate probit model. Under this model, we have:

$$\begin{aligned} P_1^* &= \beta_1' x_1 + \epsilon_1 \\ P_2^* &= \beta_2' x_2 + \epsilon_2 \\ P_3^* &= \beta_3' x_3 + \epsilon_3 \\ P_4^* &= \beta_4' x_4 + \epsilon_4, \end{aligned}$$

where the P^* terms represent latent variables for the propensity to use each service, the x terms represent explanatory variables, the β terms represent coefficients to be estimated, and the ϵ terms represent random errors. These errors are assumed to follow a multivariate normal distribution. Also estimated with the model are the correlations between these error terms (measures of unobserved dependencies among the different service channels). In this model, the latent P^* variables are not observed. Rather, one observes only the 0-1 indicator for whether a service channel was employed. With four service channels, traditional maximum likelihood estimation was not a tractable option. Instead, we employed a methodology known as simulated maximum likelihood, which relies on repeated random draws from truncated normal distributions to estimate the value of the likelihood function for alternative sets of parameter values until convergence is obtained. Provided that sufficient random draws are employed, the resulting estimates have desirable properties and are similar to traditional maximum likelihood estimation results.

We began by selecting candidate explanatory variables, limiting our selection to variables that were present in both the TES data file and our return level data base. These included indicators for four FPL categories, interactions between

⁶ While VITA may be considered the fifth relevant channel, a direct indicator of VITA participation is already available on the model input file.

each of these categorical indicators and an indicator for taxpayer receipt of a relevant IRS notice, and indicators for married-filing jointly status, the presence of a dependent, primary taxpayer under age 30, presence of self-employment income (Schedule C or Schedule F), and an EITC claim. After first estimating our model using all of these candidate variables, we performed Wald tests to identify jointly insignificant regressors. After removing such regressors, we estimated a more parsimonious final model. For all four equations, the explanatory variables in the final specification included the FPL indicators and their interactions with the notice indicator. The phone equation also included indicators for married-filing jointly status and the presence of a dependent. The mail equation also included indicators for a married-filing jointly status and an EITC claim, and the email equation also included an indicator for an EITC claim.

After estimating the weighted model using the TES subsample of returns filed by taxpayers under age 65, we next used the estimation results to predict the probability in our return-level database associated with each of the 16 possible combinations of services (1 possible way to not select services from any channel, 4 possible ways to select a service from only one channel, 6 possible ways to select services from two different channels, 4 possible ways to select services from three different channels, and 1 possible way to select services from all four channels). We then used a uniform random draw to impute different combinations of service indicators to returns in our sample based on the estimated probabilities associated with each combination of services.

Other Considerations

Weighting: Since the PTC is available to those without Medicare, we exclude tax units where both the primary and spouse are age 65 or older. Since the focus of our model is Form 1040 filers, we exclude the sampled employees who were nonfilers or who filed late (past our cut-off date of October 9, 2012). For modeling purposes, we weight by three key units of analysis: employers, employees, and individual income tax returns. We construct a set of firm weights (firmwt) that make the 19,476 sampled employers in our data file broadly representative of all employers in the population. For the individual income tax return, our weighting accounts differently for those with and without wage earnings. In the case of wage earning individuals, each sampled employee was associated with a unique Tax Year 2010 Form 1040 return. To transform our employee-based sample weights into return-based weights, we accounted for two cases where the probability of selection was higher: first, individuals who worked for more than one sampled employer, and second, married taxpayers filing jointly who each worked for a sampled employer. The sample weight was adjusted downwards in such instances to account for oversampling of such returns.

Aging: Our file is aged from 2010 by State-level growth factors for each of the years in 2014–2017 by applying forecasts from IRS Publication 6149. These account for growth of all returns, including elderly returns.

Final File: Ultimately, we arrived at two interrelated model files: (1) a nationally representative sample of tax returns (Return Level Weighted); and (2) a nationally representative sample of employers (which also includes additional employers who were not sampled, but are associated with sampled individual income tax returns; this file is named Employer List).

- We have a sample of 401,039 individual returns representing 114,628,010 nonelderly returns. Of those returns, weighted, the model estimated the prereform health insurance status statistics to be 65,587,248 returns on ESI, 8,935,391 returns on Private, 13,249,133 returns on Public, and 26,856,238 returns uninsured.
- We have a sample of 19,476 employers, representing a population of 4,769,044 employers (this is augmented with 78,328 employers included as employers of primary taxpayer's spouses).

Policies

Microsimulation models are most often used to inform policymakers and analysts of potential policy impacts. While policy rules seem straightforward to implement, they seldom are. Each model has a different set of assumptions and levels of granularity, often determined by the limitations of source data. As models were being specified for ACA, policy guidance was still in development. Major assumptions and policy provisions in the case of ACA have changed over the duration of this study. A few examples are:

- **Employee Choice Vouchers:** Free choice vouchers were eliminated in 2011 Budget Compromise passed April 15, 2011.

- **Medicaid Expansion:** The Supreme Court ruling in *National Federation of Independent Business, et. al., v. Sebelius* (June 28, 2012) limited the ability to withhold Federal Medicaid funds from States opting out of Medicaid expansion.
- **Employer Shared Responsibility:** In July 2013, implementation of the Employer Shared Responsibility provision was delayed until Tax Year 2015 and fully phased in for Tax Year 2016. Additional regulatory guidance clarified affordability guidelines, mandating that employee payments not exceed 9.5 percent of the W-2 wages, applied to individual (self-only, not family) plans.
- **Individual Shared Responsibility:** In 2014, third-party information returns will not be reported to the IRS, potentially decreasing the impact of the individual shared responsibility provision.

These changes illustrate one of the key motivators for having a microsimulation model—the ability to conduct ‘what if’ analysis when policy changes.

Our model represents basic eligibility rules for the PTC, the imposition of the ISRP for those without health insurance and Employer Shared Responsibility Payments (ESRP) for Applicable Large Employers (ALEs).

PTC and Reconciliation: IRS is responsible for administering the end-of-year reconciliation of APTC/PTC. The Exchange will determine upfront taxpayer eligibility, reflecting Modified Adjusted Gross Income (MAGI) or its approximation at enrollment. For a significant number of taxpayers, the applicable FPL will vary from the initial Exchange determination as income fluctuates throughout the year. When the PTC amount for the full year is greater than advance payments received, the excess is considered a credit for the taxpayer and is applied against their tax balance or added to their refund. However, if income increases during the year (or family size decreases), a taxpayer may have a tax liability from the APTC. This APTC repayment amount is based on FPL and premium amount. To simulate the reconciliation process, the model compares the actual 2010 and 2011 FPL (based on Modified Adjusted Gross Income and number of exemptions) for sampled returns. For taxpayers who do not have a return in 2010 and 2011, their FPL change is randomly assigned to fit the distribution of the observed FPL changes. Effectively, the 2010 FPL serves as the estimated FPL used to qualify for APTC. The 2011 FPL serves as the end-of-year FPL on which the PTC is based. Currently, the change in FPL occurs 6 months into the year. The model then compares APTC and PTC amounts based on this variation. Moving forward, the model will be updated to use distribution parameters of more recent tax years to estimate more recent FPL change scenarios. We will also explore FPL changes at different points in the year.

Individual shared responsibility payments: The ISRP amount for each person on a tax return will be calculated as the greater of a flat payment (legislation specifies \$95 in 2014, \$325 in 2015, and \$695 in 2016, adjusted for inflation thereafter) and a payment based on the percentage of applicable family income (1.0 percent in 2014, 2.0 percent in 2015, and 2.5 percent in 2016, adjusted for inflation thereafter), but will not exceed the applicable Bronze plan amount. In estimating these payments, we have wide confidence intervals since penalties are applied for individuals and our model simplifies all members of the household unit to one insurance status.

Employer shared responsibility payments are overstated by the model. Our overestimation derives from the uncertainty in modeling offer behavior between the employer and employees. First, it proved very challenging to enforce a notion of take-up rate for a given employer. Therefore, the behavioral transitions for the individuals are unable to aggregate up to a very meaningful result for an employer. For this reason, applying the level of rules required for an employer (offer value, number of full-time/part-time employees, etc.) entails high-level estimation and represents an overstatement of the number of employers likely to face responsibility payments. An ALE will be subject to the ESRP in two circumstances: (1) they do not offer insurance; or (2) they do not offer affordable insurance (both circumstances require at least one full-time employee to receive the PTC for the ESRP to apply). This model projects ESRP only for ALEs not offering insurance.

Behavior

Behavior, such as participation in public programs, can be modeled in various ways, including cell-based imputation, elasticity approaches and/or utility-based equations. Among microsimulation models, CBO and the Lewin Group use elasticity-based approaches, whereas RAND and Urban Institute use a utility-maximization approach to estimate behavioral change.


- **Elasticity-based models** exploit variation in prices over time to estimate changes in marginal preferences by the group. The disadvantage of elasticity-based models is that historical data are challenging to estimate, so measures of current elasticity vary substantially. Moreover, historical data may not be a good representation of decisions with such vast policy changes. Specifically, prediction in an elasticity-based model is based on empirical evidence from past experience. Therefore, when confronted with new experiences or policy with little to no empirical evidence, elasticities may provide limited information.
- In contrast, a **utility-maximization approach** permits individuals and firms to weight the benefits of an option (e.g., reduced out-of-pocket expenditure, lower risk) against the costs (e.g., higher premiums). Utility-based models allow greater flexibility to evaluate decisionmaking under new situations. However, the predictions of such a model are sensitive to the empirical specification of the utility function, and observed choices are not always consistent with the predicted outcomes.

By design, the RAS-ACA Model has no built-in behavioral model to drive post-ACA transitions. Rather the model represents these two interrelated behavioral decisions through employer / employee linkage, determining the provision of the largest source of health insurance, ESI (as shown in Figure 3).

Behavioral transition 1: Employer Choice to Offer ESI. Given uncertainty about employer behavior, employers keep, offer, or drop insurance based on parameters entered by the user (assigned by employer size for three types of employers: private, State or local government, and Federal Government). In the baseline, we assume employers who currently offer are likely to continue offering.

Behavioral transition 2: Choice of Health Insurance for the Return. The key factor in determining coverage will be the available choices to the individual. Therefore, the probability transitions are assigned based on the employer’s offering status. Decisions will be dependent on available family income, as expressed through FPL, as well as other characteristics not available for our model, such as health status, risk aversion, religious and social norms and values. Since current-day preferences express some of these unobserved characteristics, we assign new insurance status conditional on today’s insurance status. Transitions are entered by the RAS analyst to govern return-level insurance, by current health insurance status and FPL.

FIGURE 3. Behavioral Design for Health Insurance Decisions



Users rely on an external model, informed judgment, or intuition to choose behavioral parameters. One of the most important sources of the baseline parameters is the model used by OTA. They provided a framework for reasonable transitions from one insurance status to another for different FPL groups. RAS then explored the outcomes of various scenarios where those transitions were adjusted to obtain a better understanding of likely outcomes. In the future, this approach may be further refined.

The model does not currently account for major behavioral assumptions concerning inertia and psychological factors, ESI take-up rates and auto-enrollment provision, or the participation rates for public programs. These could be built into a behavioral engine. However, the analyst can also project the outcome and assign via cell-based imputation in SAS.


Outcomes

The RAS-ACA Model generates outputs at three levels. Level 1 estimates project post-reform insurance behavior by FPL. These estimates project the number of tax returns, grouped by FPL category, distributed across five insurance categories: ESI, Private, Public, Uninsured, and Exchange. These estimates are similar to outputs from other microsimulation models, but they are restricted to the population of Federal individual income tax filers and they are return-based rather than individual-based. At this level, the outputs we are most concerned with are the Exchange population, PTC eligible returns, the uninsured population, and employer offer decisions.

Level 2 estimates focus on how ACA provisions will impact tax returns. These estimates use Level 1 outputs as inputs and incorporate data from tax returns. For instance, the PTC population selected in Level 1 is used to determine the impacts of reconciliation. Tax data is used to determine FPL change, prereform balance due or refund, and how reconciliation impacts their balance due or refund. The key level 2 outputs are: refund/balance due change as a result of reconciliation, refund/balance due change as a result of ISRP, number of taxpayers exempt from ISRP due to unaffordable premiums, and the number of employers subject to ESRP.

Level 3 estimates project the impact of ACA tax provisions on customer service demand and compliance activity; Level 1 and 2 outputs serve as inputs for these projections. For instance, if Level 2 estimates project few returns required to repay a portion of their APTC as a result of reconciliation, projections of collection activity as a result of ACA will likely be low. Figure 4 illustrates what is estimated at each level and the relationship of the three levels.

FIGURE 4. Model Projection Hierarchy



Baseline Scenarios

We rely on external projections for guidance since we do not have a modeled behavioral component. Baseline scenarios are based on external projections from CBO and OTA. Therefore, our Level 1 baseline outputs approximate OTA and CBO outcomes on a return level. We are able to match these outcomes using behavioral transitions provided by OTA broken out by current health insurance status and FPL.


Establishing the relationship between returns and individuals requires a separate set of assumptions. The average tax unit consists of about 2.1 individuals. However, as FPL increases, the size of the household generally increases as

well. Therefore, we must make assumptions about the average tax unit size among the different insurance statuses. For instance, we can assume returns choosing Public insurance are likely to be low income, and thus, have a lower individual-to-return ratio.

Furthermore, CBO and OTA use average annual estimates for each of their insurance categories that do not account for churning (i.e. moving in and out of health insurance statuses). Churning is an important factor because everyone who receives APTC is required reconcile their advanced credit with the credit they should have received. Annual averages likely understate reconciliation activities. Therefore, we calibrate the behavioral transitions to approximate CBO and OTA estimates on a return level while allowing for additional churning into the Exchange and PTC populations. Lastly, outcomes must account for the lack of nonfilers. For instance, a portion of the uninsured population does not have a filing requirement and will not file. Therefore, the model projects fewer returns relative to the number of individuals projected by CBO and OTA.

Figure 5 provides a notional comparison of the RAS-ACA Model outputs compared to model outputs from the Urban Institute (UI), the Lewin Group, CBO, and RAND. The RAS-ACA Model outputs (Return) are generally in line with other model outputs; however, the population choosing Public insurance is noticeably smaller in most scenario runs compared to other models. This is due to the RAS-ACA Model’s focus on filers as opposed to individuals. Also, this notional comparison highlights the impact of assumptions regarding the average tax unit size across insurance categories. Larger populations of uninsured and Exchange returns relative to the other models can be attributed to a smaller return to individual ratio.

FIGURE 5. Notional Comparison of Outputs From Multiple Models



Our baseline outcomes are driven by the most recent data available. Therefore, as CBO and OTA projections change, our Level 1 baseline outputs change as well. Also, as we obtain actual data and a better understanding of behavior we will continue to update our baseline and alternative scenarios. For instance, HHS recently reported that 8 million individuals have enrolled in the Exchanges, which exceeds the April CBO estimate by 2 million individuals. Moreover, HHS reported the distribution of ages for Exchange enrollees and provided the ratio of applicants to applications. These are important inputs since age is a primary factor in determining premium prices and the applicant to application ratio gives us insight into the individual to return ratio for Exchanges. Changes to Level 2 and 3 estimates will also change due to the dependent nature of these downstream estimates. Once we have actual tax data, the baseline will be updated again.

Alternative Scenarios

Our key assumptions focus on employer offer rates, behavioral transitions, premium amounts, and FPL dynamics. Available data provides insight and guidance to our assumptions and model inputs. Premium amounts are now set for 2014; we do not expect FPL dynamics to deviate drastically from current trends, and most models agree that employer


offer rates will experience little to no change in the early years of ACA implementation. Behavioral transitions, on the other hand, must account for so many factors that uncertainty is expected.

The flexible nature of the model allows IRS the opportunity to explore various “what-if” scenarios to better explore the inherent uncertainty of taxpayer behavior. Recent data from HHS shows about 85 percent of Exchange enrollees will receive the APTC. This indicates a higher concentration of Exchange enrollees in the 100 to 400 percent FPL group than we initially expected. The flexible nature of the model allows us to change this behavioral assumption easily and analyze the downstream impacts of this new scenario on taxpayers and the IRS. We can examine scenarios where large numbers employers drop coverage or premiums spike in future years. This provides us a tool to analyze the sensitivity of our assumptions, and provides us with a range of possible outcomes based on realistic scenarios.

Model Outputs

Each model output for individual returns can be broken into FPL groups. Segmentation at this more refined level can inform taxpayer behavior in response to ACA tax provisions, potential outreach strategies, and more awareness of how ACA tax provisions are impacting different taxpayer groups. Figure 6 provides a more granular view of post-reform insurance decisions. As expected, Public insurance makes up a large portion of coverage among the lower FPL groups. Conversely, ESI is less popular at lower FPL levels but becomes more prominent as FPL increases.


FIGURE 6. Notional Distribution of Post-Reform Insurance Outcomes by FPL



The model then applies taxpayer data to these Level 1 projections to estimate the impact of ACA tax provisions on taxpayer returns. Specifically, the model examines the implications of the ISRP and PTC for filers of Federal individual income tax returns. Some nonfilers, who fall outside of our model, will also have ACA-related reporting requirements. In particular, uninsured nonfilers are required to calculate and report their ISRP or claim an exemption unless they do not have a filing requirement. There are nine exemptions from the ISRP, and only a few are administered solely by IRS. The model projects volumes for two of the more prominent IRS-administered exemptions: no filing requirement and unaffordability.

The model uses current income and projected income changes to estimate income at filing for sampled returns. The income and other return characteristics allow the model to determine whether a given taxpayer has a legal filing obligation. Taxpayers who are not required to file are exempt from the ISRP. Furthermore, age-weighted premium rates by State provided by HHS allow the model to roughly estimate the lowest cost Bronze plan for a given return. The model uses this information to determine which returns qualify for the unaffordability exemption. If 8 percent of household income is less than the lowest cost Bronze plan, the return is exempt. Figure 7 shows the distribution of Filing Requirement and Unaffordability exemptions by FPL. As expected, a large majority of these exemptions are available to the lowest FPL brackets, but some taxpayers in higher FPL brackets may be eligible for an unaffordability exemption, especially if Bronze premiums in their rating area are high.

FIGURE 7. Filing Requirement and Unaffordability Exemptions to ISRP by FPL




Additionally, the model projects the liability for those subject to the ISRP. Using income data from the tax return, the model determines whether a taxpayer pays the flat-rate payment or percent payment. These penalties are set for 2014, 2015, and 2016 and tied to inflation for 2017 and beyond. The payment is compared to the national average Bronze plan offered through the Health Insurance Marketplaces for the coverage year of 2014 and adjusted using the CPI inflator for Tax Years 2015 and beyond. Payment amounts exceeding the national average Bronze plan are capped at the annual Bronze-premium amount.

The model addresses multiple aspects of the PTC and reconciliation. Reconciliation could have significant impacts on a taxpayer's refund, especially in rating areas with high premiums (i.e. the more expensive the second lowest cost Silver plan, the larger the PTC and the more potential for volatile reconciliation outcomes). Figure 8 provides a notional projection of the impacts of reconciliation on taxpayer refunds or balances due. The figure shows about half of PTC recipients will experience an FPL increase and about half will experience an FPL decrease. Very few are expected to have an even reconciliation because small changes in income or family size should result in at least a marginal change in PTC at reconciliation. This outcome is the result of one possible FPL change scenario. For instance, FPL change estimates in a recession year would likely result in fewer repayments and higher tax credits.

The model's access to taxpayer data gives us the ability to explore outcomes relevant to taxpayers and the IRS in more detail. In addition to the high-level outcomes provided by Figure 8, we can identify the actual additional credit or repayment required, whether the repayment amount is limited by statutory caps, whether the taxpayer currently has refund or balance due (not considering ACA), and whether reconciliation changes their refund or balance due status. This adds considerable value to the IRS as it identifies potential increases in collection cases, notice volumes and demand for live assistance.

FIGURE 8. Notional PTC Reconciliation Implications



For employers, the model outputs the number and percentage of employers offering insurance and the number not offering insurance after ACA implementation. Figure 9 shows offer rates by employer size (number of employees). The figure mirrors current-day offer behavior. Small firms are less likely to offer, most firms larger than 50 are likely to offer, while almost 100 percent of the largest firms offer. Analysts can manipulate scenarios to change ESI offer outcomes and observe the downstream impact on taxpayers and the IRS. Moreover, the model sums the number of employers with an ESRP. For the first 2 years of ACA, RAS analysts will make off-model adjustments to consider implementation delays for employers.

FIGURE 9. Notional ESI Offer Rates by Employer Size



Finally, the model uses historical data from the CDW and survey data from the Taxpayer Experience Survey to output results on service and compliance behavior. Using a combination of CDW and TES data, the model develops a customer service profile. The model outputs by FPL range and by post-reform insurance status, counts of returns and their IRS service channel preference of phone, Taxpayer Assistance Center, mail, and email. Moreover, the model uses CDW data to determine characteristics of returns that currently experience compliance actions to identify existing and potential new compliance cases as a result of ACA. The compliance actions include Math Error, Automated Under-reporter, and Collection.

Future Development and Model Maintenance

In the future, RAS plans to expand the model to include additional capabilities and update the model as more information surrounding the assumptions becomes available. New developments will include additional functionality and capabilities regarding service and compliance as well as an updated approach to modeling FPL dynamics. Additionally, RAS will maintain the model by updating the input data as internal administrative data and new external research data become available.

Currently, the model assigns Customer Service and Compliance outcomes to emulate current taxpayer behavior. In the future, RAS plans to model the impact of ACA on Service Demand and Compliance and observe the impacts on IRS workstreams. New data from TES must be incorporated to capture additional digital channels, such as social media or smart phone applications, and how preferences may have shifted from live service channels to lower cost digital service channels. Moreover, expanded capabilities will allow RAS to shift preferences among different channels and determine the impact to IRS workload.

The model currently uses changes in FPL from Tax Year 2010 to Tax Year 2011. The change represents a recovering and growing economy. RAS will include multiple years of changes in income and family size: Tax Years 2008 to 2009, 2009 to 2010, 2011 to 2012, and 2012 to 2013. RAS will run the model on different scenarios to determine the sensitivity of reconciliation outcomes to changes in FPL.

RAS will continue to monitor external models and new data that impact ESI offerings, transition rates to new insurance, the average individual/family premium cost by State, changes in compliance, services, and paid preparer and software use. Less frequent updates include adjusting age or return demographics for future years, adding new variables from administrative data as they becomes available and adjusting weights in the sample.

RAS will incorporate statistics or calibrate the model to align with updates as external research data become available. Health and Human Services (HHS) released a final enrollment report in May 2014 that summarizes enrollment at the Health Insurance Marketplaces and provides statistics by Marketplace and age. The model will be calibrated to ensure returns selected for Health Insurance Marketplace coverage follow the same State and age distributions reported by HHS. The National Health Interview Survey, Current Population Survey, Kaiser Family Foundation, American Community Survey, and Medical Expenditure Panel Survey are additional resources that provide useful insights to health insurance in the United States and as these surveys are updated, the model will calibrate to known behaviors.

Current and Future Applications

The primary use for the RAS-ACA Model is to inform the IRS of impacts from the legislation on workload volumes. RAS generates a high-level summary report derived from outputs of the model to inform IRS executives of key volumes critical for resource planning. These volumes include the number of returns claiming the PTC, reporting the ISRP, and receiving ISRP exemptions; the number of ALEs subject to the ESRP; and the number of ACA-related information returns filed, for Tax Years 2014 through 2017. The model is used to run various scenarios and output the same estimates using internally consistent assumptions.

The model results inform various IRS operational questions on compliance and service. Specifically, IRS analysts used model outputs and applied historical notice rates to predict the volume of additional notices sent out for the first tax year of ACA implementation. The projections include both new and existing notices as a result of ACA. Moreover, PTC reconciliation outcomes were used by the IRS ACA Customer Service and Stakeholder Relations team to inform external messaging strategies for both the taxpayer population and government partners.

Level 1 and 2 estimates provide an understanding of the number of people impacted by ACA. These high-level counts are used to scope potential customer service demand. The model is being enhanced to provide insight into the types of questions taxpayers may need to resolve, whether it be form-related, account information, tax law, or payments. The results of this analysis will be used to determine any potential increases in demand for live assistance. This will inform risk analysis, potential mitigation strategies, and workforce plans.

The RAS-ACA Model is also being enhanced to estimate the impact on compliance activities. Some examples include notice projections, collection workload, math error, and impacts to Automated Underreporter (AUR). The IRS is able to use the information to make decisions about resource allocation, including print capabilities, reprogramming AUR, and sampling for the National Research Program.