Methodological Appendix

This appendix describes how I estimate the effects of a version of the Administration's HRA proposal that lacked one or more of the safeguards against worker-level shifting included in the proposed rule. The first section presents the modeling framework. The second section describes how I calibrate and solve the model. The final two sections discuss the methods used to process two of the data sources used in the calibration.

Modeling Framework

Consider two distinct populations $m \in \{E, I\}$ of size N_m, where population E consists of people receiving coverage from a large employer (without the HRA policy) and population I consists of people obtaining coverage on the individual market (without the HRA policy). A randomly drawn member of population m has expected claims risk R_m, which is assumed to be observable to both the individual and the employer, as well as an age rating factor A_m that would apply if that person enrolled in the individual market.[1] I assume that insurers offer a single plan design and that the plan's expected cost of covering a person with claims risk expected R_m is cR_m for some constant c; I discuss the implications of this simplifying assumption below.

Prior to implementation of the HRA policy, the individual market only contains people purchasing coverage on their own. Expected plan spending is thus given by $c\mathbb{E}[R_I]$. In equilibrium, insurers price to cover expected claims costs and thus set a base premium (i.e., the premium for a person with an age rating factor of one) of $B_{\text{pre}} = c\mathbb{E}[R_I]/\mathbb{E}[A_I]$.[2] When solving the model, it is convenient to normalize this quantity by the expected plan liability generated by covering the large employer population. Denoting the normalized quantity by b_{pre}, it is straightforward to see that

$$b_{\text{pre}} \equiv \frac{B_{\text{pre}}}{c\mathbb{E}[R_E]} = \frac{\mathbb{E}[R_I]/\mathbb{E}[R_E]}{\mathbb{E}[A_I]} = \frac{\phi}{\mathbb{E}[A_E]}, \tag{1}$$

where

$$\phi \equiv \frac{\mathbb{E}[R_I]/\mathbb{E}[A_I]}{\mathbb{E}[R_E]/\mathbb{E}[A_E]}$$

is the portion of the pre-policy difference in expected claims risk between the individual and large employer markets that insurers cannot price for via age rating.

After implementation of the HRA policy, the individual market contains both people purchasing coverage on their own and people purchasing individual market coverage via an HRA. The new equilibrium base premium in the individual market is thus given by

$$B_{\text{post}} = \frac{N_I c\mathbb{E}[R_I] + sN_E c\mathbb{E}[W]\mathbb{E}[R_E|W=1]}{N_I\mathbb{E}[A_I] + sN_E\mathbb{E}[W]\mathbb{E}[A_E|W=1]},$$

where s is the share of employers offering HRAs (assumed to be randomly selected from the overall population of employers), and W is an indicator variable that equals one if an individual worker would enroll in the individual market if offered an HRA.[3] The numerator of this fraction is plans'

[1] I assume age rating occurs in accordance with the federal default age rating curve in place for 2018 and later.

[2] I implicitly assume away administrative costs. This assumption is innocuous for the present purpose.

[3] The assumption that employers offering HRAs are randomly selected from the overall population of employers is surely not precisely correct. This would likely cause my estimates to understate the adverse effects of the HRA policy on the individual market risk pool.

aggregate expected claims liability, while the denominator is the sum of enrollees' age rating factors. Normalizing by the expected claims liability generated by the large employer population then yields

$$b_{\text{post}} \equiv \frac{B_{\text{post}}}{c\mathbb{E}[R_E]} = \frac{N_I \phi \frac{\mathbb{E}[A_I]}{\mathbb{E}[A_E]} + sN_E \mathbb{E}[W] \frac{\mathbb{E}[R_E|W=1]}{\mathbb{E}[R_E]}}{N_I \mathbb{E}[A_I] + sN_E \mathbb{E}[W]\mathbb{E}[A_E|W=1]}. \tag{2}$$

Equation (2) makes clear that post-policy premiums depend importantly on how many and what types of people associated with large employers switch into the individual market. As argued in the main text, if one or more of the safeguards against worker-level sorting were eliminated, workers or their employers would sort workers between individual market coverage and alternative coverage options priced based on each worker's own expected claims risk. I assume that workers would end up enrolled in the lower-cost option unless some "friction" kept them from doing so. That is, I assume that enrollment decisions are governed by

$$W = 1\{(1+f)A_E B_{\text{post}} < cR_E\} = 1\left\{(1+f)A_E b_{\text{post}} < \frac{R_E}{\mathbb{E}[R_E]}\right\}, \tag{3}$$

where f is a proportional friction amount.

As discussed in the main text, the friction f provides an ad hoc way of capturing a variety of factors that may affect enrollment decisions but are not directly modeled. One factor not directly accounted for in the model is the possibility that the relative price of more generous plan designs (e.g., broad-network plan designs) may be higher in the individual market due to adverse selection pressures that are not fully compensated for by risk adjustment. To the extent that more generous plan types are highly valued by enrollees who would have an incentive to leave employer coverage under the HRA policy, that could lead enrollees to be less willing to take up individual market coverage than my base behavioral assumption would imply; this tendency can be captured by a positive value for f. The friction f may also capture differences in hassle costs between the two enrollment modes, differences in loading factors between different types of coverage, or other factors that affect enrollment decisions.

Calibration and Solution Method

To calibrate the distributions underlying the various (conditional) expectations involving people in employer coverage that enter equations (1), (2), and (3), I construct a sample of enrollees in large employer coverage that includes data on both age and expected claims risk. Construction of this sample is described in detail below. I calibrate ϕ, the baseline relative claims risk of the individual and group markets, using CMS risk adjustment data and the method, processed as described in the final section. I calibrate $\mathbb{E}[A_I]$ using the age distribution of people with individual market coverage and no other coverage in the 2017 National Health Interview Survey.

Since my focus is on the long-run, I set N_I to 12.0 million, consistent with the Congressional Budget Office's projections of individual market enrollment in 2028 (CBO 2018). Similarly, I set N_E to approximately 135 million. This estimate is obtained by combining CBO's projection that 154 million non-elderly people will be enrolled in employer coverage in 2028 with an estimate based on the 2016 and 2017 Medical Expenditure Panel Survey, Insurance Component that 87.7 percent of people enrolled in employer coverage are enrolled through a large employer.

Pre-policy premiums and enrollment can be calculated directly from the calibrated parameters. I use an iterative method to solve for post-policy premiums and enrollment. Specifically, I start with an initial value for b_{post}. I use this initial value to compute W for each person in the employer sample and then compute $\mathbb{E}[W]$, $\mathbb{E}[R_E|W=1]$, and $\mathbb{E}[A_E|W=1]$ by averaging over the sample. I use these amounts to compute a new estimate of b_{post}. I repeat these steps until the estimates of b_{post} converge, which occurs quickly in practice.

The premium change under the policy can be calculated directly from the equilibrium premiums. Post-policy individual market enrollment is then given by $N_I + sN_E\mathbb{E}[W]$. The change in effective price metric reported in the text can be calculated as

$$\frac{(1-\mathbb{E}[W])c\mathbb{E}[R_E|W=0]+\mathbb{E}[W]\mathbb{E}[A_E|W=1]B_{\text{post}}}{c\mathbb{E}[R_E]}-1$$
$$=[1-\mathbb{E}[W]]\frac{\mathbb{E}[R_E|W=0]}{\mathbb{E}[R_E]}+\mathbb{E}[W]\mathbb{E}[A_E|W=1]b_{\text{post}}-1.$$

Construction of the Employer Sample

I construct a sample of enrollees in employer coverage by pooling the longitudinal files of the Medical Expenditure Panel Survey, Household Component (MEPS-HC) for 2009/2010 through 2015/2016. I limit the sample to people who meet the following inclusion criteria: (1) enrolled in employer coverage during all 24 months they appear in the panel; (2) under age 65 in the second year they appear in the panel; and (3) non-missing self-reported health status in the first year they appear in the panel. I do not limit the sample to people associated with large employers since enrollee characteristics are unlikely to differ substantially between large employers and small employers. I trend all spending amounts to 2016 based on the trend in average per enrollee spending in employer-sponsored coverage, as reported in the National Health Expenditure Accounts.

I construct an estimate of expected claims risk for each member of the sample. A significant challenge in constructing a suitable estimate is that the MEPS-HC is believed to understate the number of people with very high health care spending. In light of this shortcoming of the MEPS-HC, many natural approaches to estimating expected claims risk would likely also understate the number of people with very high *expected* claims risk, which would likely in turn lead me to understate the potential effects of relaxing the safeguards included in the proposed rule.

The ideal solution to this problem would be to perform this analysis in a database of health care claims rather than the MEPS-HC, but that was not feasible on the timeline under which the proposed rule is being considered. Therefore, I instead take the following three-step approach:

- **Construct a preliminary measure of expected claims risk in the MEPS-HC:** I first construct a preliminary measure of expected claims risk for each observation by estimating the following (non-linear) regression in the MEPS-HC sample:

$$S_2 = H\gamma + \sum_{k=1}^{5}\theta_k Y^k + \beta[1-F_1(S_1)]^{\xi} + \epsilon,$$

  where S_t represents spending in an individual's t^{th} year in the sample (normalized to have mean one), H is a vector of indicator variables for levels of self-reported health status, Y is age in years, F_1 is the cumulative distribution function of S_1, and ϵ is an error term.

- **Calibrate the marginal distribution of $\mathbf{R_E}$ from other evidence:** I next calibrate the marginal distribution of R_E to match the dispersion in ACG risk scores estimated in a database of health care claims by employees at small firms by Fleitas, Gowrisankaran, and Lo Sasso (2018), henceforth FGL. The ACG risk score, as calculated by FGL, incorporates enrollee age, prior year diagnoses, and prior year spending. The mean ACG risk score is normalized to one.

For calibration purposes, I assume that the distribution of R_E takes a generalized Pareto form, which does a good job of describing the distribution of similar risk scores constructed directly in the MEPS-HC. This distribution is a three-parameter family, so I use three moments for calibration: the minimum of the distribution's support, the mean, and the variance.

The mean is normalized to one, and I obtain the variance directly from Table 1 of FGL. FGL do not report the minimum ACG risk score in their sample, so I instead estimate a similar quantity using the MEPS-HC. Specifically, I estimate mean spending in the MEPS-HC sample among people who are: (1) under age 25; (2) have no spending in the prior year; and (3) report being in excellent health. I then divide this quantity by mean spending in the MEPS-HC sample, scaled up to reflect the percentage by which the MEPS-HC understates aggregate private health insurance spending as reported by Bernard et al. (2018). This scaling reflects the fact that the MEPS-HC understates average spending, but likely has a much more limited effect on expected spending at the bottom of the spending distribution.

- **Combine the preliminary measure and calibrated distribution:** The final step is to combine the results of the first two steps to impute a final measure of expected claims risk. Specifically, I obtain predicted values from the regression in the first step and then compute each observation's quantile within the distribution of predicted values, which I denote by $\hat{Q}$. I then assign each observation the expected claims risk associated with the corresponding quantile of the calibrated distribution for R_E. That is, the final estimate of expected claims risk is $R_E = F_{R_E}^{-1}(\hat{Q})$, where F_{R_E} is the calibrated distribution for R_E. This procedure aims to preserve the correlation between expected claims risk and enrollee age to the greatest extent possible, which is important since individual market premiums are age rated.

Estimating the Relative Claims Risk of Individual and Group Market Enrollees

Data in CMS' 2017 risk adjustment summary report can be used to estimate the relative claims risk of individual and small group market enrollees. Under the assumption that risk mix among small employers is similar to that among large employers, this information can be used to estimate the quantity ϕ that appears in equations (1) and (2).

CMS' risk adjustment reports provide three relevant pieces of information for each state (except Massachusetts and Vermont) for both the individual and small group markets: (1) the average plan liability risk score (PLRS); (2) the average actuarial value; and (3) the average rating factor. The average rating factor corresponds almost exactly to the average age rating factor that enters ϕ, leaving aside the small number of enrollees rated for tobacco use. Estimating the risk score components of ϕ is somewhat more complicated because the PLRS is a measure of expected claims liability to the *insurer* given the enrollee's health conditions, plan actuarial value, and cost-sharing reduction (CSR) enrollment status. To recover a measure that solely reflects health status and thus corresponds to the measure of expected claims risk that enters ϕ, I make two adjustments.

The first adjustment removes the effect of plan actuarial value. Following Owen (2016), I simply divide the average PLRS in each state and market by the average actuarial value in the state. The second adjustment, which is only relevant in the individual market, removes the effects of CSR enrollment status. Specifically, I divide the average PLRS by $1 + .12v$, where v is the share of a state's individual market enrollment receiving 87 percent or 94 percent actuarial value CSRs. This adjustment accounts for the fact that CMS increases the PLRS for recipients of these CSR types by a factor of 1.12 to account for the additional utilization induced by the lower cost sharing.[4] I estimate v using CMS data on effectuated enrollment and open enrollment plan selections. Unfortunately, the plan selection data are only available for states using the HealthCare.gov enrollment platform, so my estimate of ϕ uses data solely for these states.

Using the method described above, I compute a separate estimates of ϕ for each state, and I then compute a national average by weighting these state-level estimates by individual market enrollment. This generates an estimate of $\phi = 1.16$. I multiply this estimate by a factor of 1.125 to reflect future deterioration in individual market risk mix attributable to repeal of the individual mandate and liberalization of rules related to short-term coverage, based on estimates of the effects of those policy changes from CBO (2018).

References

Bernard, Didem M., Cathy Cowan, Thomas M. Selden, David Lassman, and Aaron Catlin. 2018. "Reconciling Medical Expenditure Estimates from the MEPS and NHEA, 2012." Agency f0r Healthcare Research and Quality Working Paper, No. 17003. https://meps.ahrq.gov/data_files/publications/workingpapers/wp_17003.pdf.

Congressional Budget Office (CBO). 2018. "Federal Subsidies for Health Insurance Coverage for People Under Age 65: 2018 to 2028." https://www.cbo.gov/system/files?file=2018-06/53826-healthinsurancecoverage.pdf.

Fleitas, Sebastian, Gautam Gowrisankaran, and Anthony Lo Sasso. 2018. "Reclassification Risk in the Small Group Health Insurance Market." National Bureau of Economic Research Working Paper, No. 24663.

Owen, Rebecca. 2016. "An Examination of Relative Risk in the ACA Individual Market." Society of Actuaries. https://www.soa.org/Files/Research/research-2016-examination-relative-risk-aca.pdf.

[4]CMS makes similar adjustments to the PLRS for certain other CSR types, but these CSR types have very low enrollment.