Does the individual mandate affect insurance coverage?
Evidence from tax returns

Ithai Lurie, Daniel W. Sacks, and Bradley Heim

Online Appendix

Appendices

A Further policy details

A.1 More details on the individual mandate

We restrict our sample to single, childless adults because the mandate penalty is more complicated for larger families. Here we describe these complications. Recall from Section 1 that the monthly mandate penalty is

$$Penalty = 1/12 \max\left\{\min\left\{[A + .5C]F, 3F\right\}, S(\text{MAGI} - \text{tax filing threshold})\right\},$$

where F is the flat fee, A the number of uninsured adults that month, C the number of uninsured children, and S the percent of income. In 2016, F was \$325 and S was 0.02. For a given filing threshold, the penalty is therefore a function of income and the number of uninsured adults and children. Each child counts as half an adult for the purposes of determining the mandate penalty so we refer to the number of uninsured "adult equivalents", equal to $A + .5C$. Note that the number of uninsured adult equivalents affects the flat fee (the first term in brackets) but not the percent of income (the second term).

Appendix Figure A.1 plots the monthly mandate penalty in 2015 as a function of income and the number of uninsured adult equivalents, for a married, filing jointly tax return (which has a filing threshold of \$20,600). There are six kink points, equal to each intersection of $F(A + .5C)$ with the percent of income, for $A + .5C \in \{0.5, 1, 1.5, 2, 2.5\}$. At each of the kink points, the mandate penalty increases for some margins of coverage, but not for all margins. For example, consider a household with two adults and two children, and income of \$68,000. For this household, the percent of income payment is \$79. If there were one uninsured household member, the penalty would be \$79, because the percent of income exceeds the flat fee. A second uninsured household member would not increase the penalty because the percent of income exceeds the flat fee for two uninsured members. Only if the entire household were uninsured would the flat fee, \$81.25, exceed the percent of income. However, if the household's income were \$70,000, the percent of income would always exceed the flat fee, regardless of the number of uninsured adult equivalents. Thus there is a kink in the incentive to have the entire family covered, but no kink in the marginal incentive to cover the first, second, or third family member.

This example shows that, for multiperson households, the mandate penalty creates complex and fairly subtle incentives to increase coverage. In principle it is possible to examine coverage responses at each of the six coverage kinks, focusing on the relevant margin of coverage (e.g. a kink in the probability of having three or more uninsured adult equivalents at \$69,200). In practice we are concerned that households may not understand the specific incentives for monthly coverage generated by the individual mandate. We therefore focus on single households. For these households the penalty is relatively simple–it is linear in their number of uninsured months.

A.2 More details on other income-linked insurance inducements

The Premium Tax Credit and the Advanced Premium Tax Credit: The Premium Tax Credit (PTC) is a subsidy which may be used for purchasing an insurance plan on the

Health Insurance Marketplaces. The PTC is equal to the difference between a household's "benchmark premiums"—the second lowest-cost silver-tier health insurance plan available to it in the Health Insurance Marketplaces—and its expected contribution, a percent of income specified by law that ranges from 2 to almost 10 percent. The expected contribution is a kinked and discontinuous function of income, so for the PTC is also a kinked and discontinuous function of income, with potential discontinuities or kinks at 100, 133, 150, 200, 250, 300, and 400 percent of the FPL. There is also a kink in the PTC at the income level at which the expected contribution exactly equals the benchmark premium; this kink point varies across markets, since the benchmark premium varies across markets.[25] To help taxpayers manage liquidity, the PTC is paid in advance, throughout the coverage year, in the form of the Advanced Premium Tax Credit (APTC). APTC payment amounts are based not on realized MAGI, but on project income, which Marketplace enrollees report to the Marketplace at the time of signing up for insurance. If APTC payment are too high (because realized income exceeded projected income), taxpayers must repay the excess, with repayment limits that depend on realized income. These repayment limits are discontinuous functions of income, with discontinuities at 200, 300, and 400 percent of FPL. Heim et al. (2017) provide more detail on the PTC, APTC, and repayment requirements, and document income responses to the premium tax credit at the 400 percent of FPL discontinuity.

Cost-sharing reductions: Whereas the PTC helps subsidizes premiums, cost-sharing reductions (CSRs) subsidize out-of-pocket health care expenses. For every standard silver plan that insurers offer on the Marketplace, they must offer three additional CSR plans, which are identical in all aspects except their cost sharing. A standard silver plan has an actuarial value of 70 percent, meaning it covers 70 percent of expected health care costs. The CSR plans have actuarial value of 73 percent, 87 percent, and 93 percent. Insurers must charge the same premium for these more generous plans as they do for the base silver plan; the government paid for the additional cost-sharing until 2018. Only low-income people are eligible to purchase these more generous plans. People with income between 100 and 150 percent of FPL may purchase the 93 percent actuarial value plans; people with income between 150 and 200 percent of FPL may purchase the 87 percent actuarial value plans; and people with income between 200 and 250 percent of FPL may purchase the 73 percent actuarial value plans.[26]

Other policies: We believe that the PTC and CSRs are the most important potential threats to identification, in the sense that they create meaningful nonlinearities in the incentives to obtain insurance near the mandate kink point. Several other programs might also be relevant. Medicaid eligibility of course depends on income, although eligibility is determined in terms of rolling income throughout the year, rather than realized income. The eligibility levels vary across states and people. In some states single, childless adults are not eligible at any income levels, whereas children can be eligible up to 300 percent

[25]The PTC and the mandate penalty are assessed using slightly different modifications of AGI. The definition of MAGI for the purpose of PTC is similar to the mandate penalty, but also includes the non-taxable Social Security income. We focus on a population aged 27-64 with MAGI between $29,425 and $47,080, so we expect that non-taxable Social Security income is zero for nearly all our sample.

[26]DeLeire et al. (2017) show that these subsidies influence plan choice on the Exchanges.

of FPL or above.[27] Importantly for our analysis, however Medicaid eligibility is assessed on a different income basis than is mandate penalty. Medicaid eligibility depends on a rolling average of income, assess over the previous several months. The mandate penalty depends on realized taxable income. In practice, therefore, a household's income for assessing Medicaid eligibility can be quite different from its income for determining the mandate penalty (or the PTC for that matter). For example, a temporarily low income household can qualify for Medicaid, but have a high enough annual income that it would be subject to the penalty if it were uninsured. This Medicaid qualification would also not disqualify them from receiving a PTC.

Another kink in the incentive to obtain insurance comes from the tax deductibility of employer sponsored insurance, which creates a kink in the incentive to obtain ESI at each kink point in the income tax code. These kinks turn out not to be close to our nonlinearity points, because the income tax applies to taxable income rather than MAGI. For example, for a single taxpayer in 2016, the 15 percent tax bracket ran from $9,275 to $37,650 of taxable income. For a single tax payer with one exemption claiming the standard deduction, this works out to $19,625 to $48,000 in MAGI. Other programs such as SNAP, TANF, and the EITC may affect insurance demand through income effects. These programs do not have discontinuities at 138 percent of FPL, however, and at higher income levels the benefits are generally small.

Summary: The PTC, APTC, repayment requirements, and CSRs all create kinks or discontinuities in the incentive to obtain health insurance coverage.[28] These nonlinearities occur at even increments of the FPL: 100, 133, 150, 200, 250, 300, and 400 percent. The mandate discontinuity and kink points occurs between these critical values, as shown in Figure 1. It is therefore possible to separately identify the coverage effect of the individual mandate from the coverage effects of these policies by looking within narrow windows of FPL (133-150 percent, 200-250 percent, and 300-400 percent).

[27] See `https://www.kff.org/health-reform/state-indicator/medicaid-income-eligibility-limits-for-adults-as-a-percent-of-the-federal-poverty-level/` and `https://www.kff.org/medicaid/fact-sheet/where-are-states-today-medicaid-and-chip/`.

[28] Tebaldi (2017) studies coverage responses to the PTC in the California Marketplace, and Frean et al. (2017) study coverage responses at a national level using the in the ACS. DeLeire et al. (2017) study coverage responses to the CSRs.

Figure A.1: Monthly mandate penalty for multi-person households, 2015

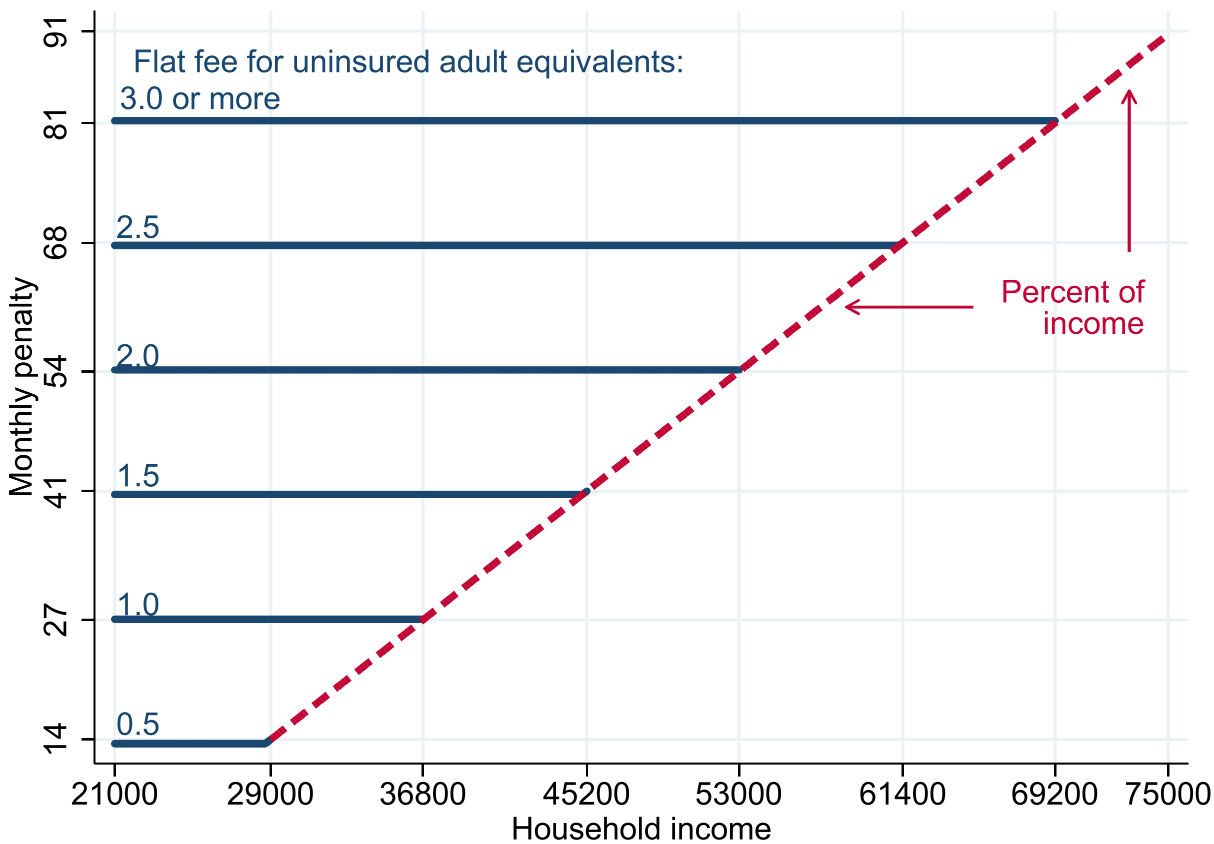

Notes: Source: Figure shows the monthly mandate penalty in 2015 as a function of income and the number of uninsured adult equivalents, for a household with married filing jointly tax return. The number of uninsured adult equivalents is the number of uninsured adults plus half the number of uninsured children. For some income levels and numbers of uninsured adult equivalents, insuring an additional adult equivalent does not change the mandate penalty, because the percent of income penalty is the same for any number of uninsured adult equivalents.

B Monte Carlo study of RKD estimators

We conducted a Monte Carlo simulation study to assess the performance of alternative RKD estimators. The canonical RKD specification is

$$y_i = \sum_{d=0}^{D} \alpha_d (v_i - c)^d + \sum_{d=0}^{D} \beta_d (v_i - c)^d \times D_i + \varepsilon_i,$$

where v_i is the running variable, D_i is indicator for the running variable exceeding the cutoff, and $\hat{\beta}_1$ is the kink estimate (Card et al., 2015). Our estimating equation is identical to this, with $D = 1$. In estimating a regression kink design, researchers must make several specification choices: the choice of degree D, the bandwidth h, the kernel, and whether to allow for a discontinuity.

The theoretical econometric literature recommends using a triangular kernel for boundary estimation problems such as this one. For estimating a kink, the theoretical literature also recommends $D \geq 2$, and it has developed plug-in estimators for bandwidth choice based on minimizing asymptotic mean squared error of the kink estimate (Imbens and Kalyanaraman, 2012; Calonico et al., 2014). However applied researchers have favored the uniform kernel—as the regression can then be estimated with OLS—and have found that high degree terms and asymptotically optimal bandwidths do not necessarily perform well in finite samples. Applied researchers also sometimes impose continuity (i.e. dropping the $(v_i - c)^0 D_i$ term).

To determine our baseline specification choices, we conducted a Monte Carlo following the suggestions in Card et al. (2017). The overall idea is to simulate many data sets using a data generating process that closely resembles our data, and then compare the performance of alternative RKD estimators across the data sets. Our data generating process is based on a high-order polynomial approximation to the data, with a true kink imposed. To do so, we first "dekink" the data by estimating the following regression, separately for 2015 and 2016:

$$y_i - \hat{\tau}_t D_i v_i = \sum_{d=0}^{5} \beta_d v_i^d + \sum_{d \neq 1}^{5} \theta_d v_i^d D_i + \epsilon_i. \tag{3}$$

where $\hat{\tau}_t$ is the estimated kink in year t, 0.05 in 2015 and 0.02 in 2016. Let $\hat{y}(v)$ be the predicted value from this regression when the running variable is v.

We simulate data with a known kink τ. We consider two cases: τ corresponding to a semi-elasticity of 0.5, which we consider to be the middle of past estimates, and τ corresponding to a semi elasticity of 0.2, which is low but consistent with the Massachusetts evidence. For each year and each of 1000 simulation data sets, we sampling with replacement from the empirical distribution of v and ϵ in that year. Each simulation dataset is the same size as our estimation dataset. Given the draw of v and ϵ, we form the outcome y as $\hat{y}(v) + \epsilon + D\tau v$, where τ is the assumed kink. We then estimate several different RKD specifications on the simulated data. For each simulated data set, we considered the power set of the following specification choices: bandwidth equal to the full range of income, the Fan-Gijbels bandwidth selector (as proposed by Card et al. (2015)), or Calonico et al. (2014) bandwidth selector (without scale regularization); polynomial degree $D = 1$ or $D = 2$; and imposing continuity or not. Throughout we use a uniform kernel, for consistency with the applied literature. We do not consider the bias-corrected estimator of Calonico et al. (2014) because it is computationally costly and initial simulations suggested that it lead to dramatically higher variance without large reductions in bias or improvements in coverage rates (a result also reported by Card et al. (2017)).

Appendix Table B.1 summarizes the performance of the various estimators in the 2015 sample. The linear estimator performs well: it achieves its nominal coverage rate, and rejects a false null 96-97 percent of the time. The Fan-Gijbels and CCT bandwidth selectors choose fairly small bandwidths, $1,248 to $1,806, relative to a maximal bandwidth of about $2,900. Relative to using the full range of the data, they have a higher RMSE and a lower rejection rate; the coverage rate is slightly higher for the discontinuous estimator and slightly lower for the continuous estimator. The linear estimator using the full range of the data has the lowest RMSE. The estimators that use only relatively local information give up some power without reducing bias. Allowing for a discontinuity results in slightly higher bias and variance. The quadratic estimators perform substantially worse than the linear estimators: they have higher (absolute) bias, much higher variance, and worse coverage.[29] We conclude that the linear estimator using the full range of data is likely to perform better than the alternatives, although none of the estimators achieves the nominal coverage rate, and this estimator has the worst coverage.

Appendix Table B.2 summarizes the performance of the estimators in the 2016 sample. Here too we find that the linear specification using the full range of the data has the lowest mean squared error, again with somewhat higher confidence intervals. In this case the coverage rate of the linear estimator is below the nominal rate when we impose continuity.

Because our 2016 estimates were statistically insignificant, we also investigated the power of our estimator to detect small kinks. Specifically, we re-ran our Monte Carlo simulations, but assuming a semi-elasticity of 0.2 instead of 0.5, and assuming a semi-elasticity of 0.14. The 0.2 semi-elasticity corresponds to the estimate that Hackmann et al. (2015) find using the Massachusetts mandate. They look at a sample of relatively high income adults, with income above 300 percent of FPL, so we believe this is a useful benchmark. We report the results of this simulation in Appendix Tables B.3 and B.4. The semi-elasticity of 0.14 corresponds to our main estimate. Consistent with our other simulation results, we find that the linear estimator using the full range of data outperforms estimators with higher order terms or tighter bandwidths. However, even for this estimator, we find somewhat limited power. When we do not impose continuity, we reject a false null in only 73 percent of iterations. Imposing continuity improves power. At the smallest semi-elasticity we considered, 0.14, we find limited power even when imposing continuity; we reject the false null in 74 percent of iterations. Without continuity we reject in less than half of all iterations.

[29]The FG bandwidth usually ends up exceeding the range of data in the quadratic case, so its performance is the same as the estimator using the full range of data.

Table B.1: Summary of performance of RKD estimators in Monte Carlo Simulation, 2015

Estimator	Median Bandwidth	$\frac{RMSE}{\tau}$	Coverage Rate	$\frac{Bias}{\tau}$	$\frac{Variance}{\tau^2}$	Rejection Rate
	(1)	(2)	(3)	(4)	(5)	(6)
A. Linear estimators						
BW = full, continuous	–	0.264	0.939	0.009	0.264	0.963
BW = FG, continuous	1786	0.346	0.947	-0.001	0.346	0.826
BW = CCT, continuous	1204	0.431	0.951	0.010	0.431	0.673
BW = full, discontinuous	–	0.271	0.943	0.012	0.271	0.958
BW = FG, discontinuous	1786	0.353	0.948	0.002	0.353	0.816
BW = CCT, discontinuous	1204	0.441	0.946	0.009	0.441	0.665
B. Quadratic estimators						
BW = full, continuous	–	1.103	0.880	0.803	1.103	0.418
BW = FG, continuous	4823	1.103	0.880	0.803	1.103	0.418
BW = CCT, continuous	1492	1.566	0.901	0.837	1.566	0.254
BW = full, discontinuous	–	1.131	0.867	0.851	1.131	0.405
BW = FG, discontinuous	4823	1.131	0.867	0.851	1.131	0.405
BW = CCT, discontinuous	1492	1.598	0.897	0.867	1.598	0.259

Table summarizes the performance of 12 RKD estimators, which differ in the degree of the underlying polynomial (linear or quadratic), bandwidth selector (full range of data, Fan-Gjbels, or Calonico et al. (2014)), and whether a discontinuity is imposed. The data are generated using a true kink of $\tau = 6.33 \times 10^{-5}$, corresponding to a semi-elasticity of 0.5 at the mandate kink point. The coverage rate is the fraction of confidence intervals containing this kink.

Table B.2: Summary of performance of RKD estimators in Monte Carlo Simulation, 2016

Estimator	Median Bandwidth	$\frac{RMSE}{\tau}$	Coverage Rate	$\frac{Bias}{\tau}$	$\frac{Variance}{\tau^2}$	Rejection Rate
	(1)	(2)	(3)	(4)	(5)	(6)
A. Linear estimators						
BW = full, continuous	–	0.139	0.898	0.069	0.139	1.000
BW = FG, continuous	2392	0.194	0.913	0.067	0.194	0.997
BW = CCT, continuous	1303	0.272	0.936	0.060	0.272	0.954
BW = full, discontinuous	–	0.170	0.943	0.034	0.170	1.000
BW = FG, discontinuous	2392	0.242	0.936	0.027	0.242	0.977
BW = CCT, discontinuous	1303	0.336	0.943	0.022	0.336	0.841
B. Quadratic estimators						
BW = full, continuous	–	0.561	0.782	-0.593	0.561	0.133
BW = FG, continuous	7495	0.569	0.809	-0.593	0.569	0.130
BW = CCT, continuous	1665	0.994	0.880	-0.586	0.994	0.070
BW = full, discontinuous	–	0.697	0.726	-0.893	0.697	0.041
BW = FG, discontinuous	7495	0.712	0.755	-0.891	0.712	0.038
BW = CCT, discontinuous	1665	1.248	0.851	-0.892	1.248	0.039

Table summarizes the performance of 12 RKD estimators, which differ in the degree of the underlying polynomial (linear or quadratic), bandwidth selector (full range of data, largest symmetric band, Fan-Gjbels, or Calonico et al. (2014)), and whether a discontinuity is imposed. The data are generated using a true kink of $\tau = 8.75 \times 10^{-5}$, corresponding to a semi-elasticity of 0.5 at the mandate kink point. The coverage rate is the fraction of confidence intervals containing this kink.

Table B.3: Summary of performance of RKD estimators in Monte Carlo Simulation, 2016, assuming low semi-elasticity

Estimator	Median Bandwidth	$\frac{RMSE}{\tau}$	Coverage Rate	$\frac{Bias}{\tau}$	$\frac{Variance}{\tau^2}$	Rejection Rate
	(1)	(2)	(3)	(4)	(5)	(6)
A. Linear estimators						
BW = full, continuous	–	0.360	0.897	0.170	0.360	0.913
BW = FG, continuous	2357	0.494	0.924	0.185	0.494	0.711
BW = CCT, continuous	1297	0.662	0.942	0.177	0.662	0.476
BW = full, discontinuous	–	0.454	0.923	0.080	0.454	0.713
BW = FG, discontinuous	2357	0.619	0.936	0.083	0.619	0.464
BW = CCT, discontinuous	1297	0.845	0.944	0.057	0.845	0.301
B. Quadratic estimators						
BW = full, continuous	–	1.370	0.790	-1.492	1.370	0.014
BW = FG, continuous	7636	1.406	0.814	-1.496	1.406	0.014
BW = CCT, continuous	1738	2.313	0.895	-1.396	2.313	0.017
BW = full, discontinuous	–	1.704	0.722	-2.241	1.704	0.006
BW = FG, discontinuous	7636	1.754	0.753	-2.247	1.754	0.005
BW = CCT, discontinuous	1738	2.914	0.865	-2.306	2.914	0.010

Table summarizes the performance of 12 RKD estimators, which differ in the degree of the underlying polynomial (linear or quadratic), bandwidth selector (full range of data, largest symmetric band, Fan-Gjbels, or Calonico et al. (2014)), and whether a discontinuity is imposed. The data are generated using a true kink of $\tau = 3.5 \times 10^{-5}$, corresponding to a semi-elasticity of 0.2 at the mandate kink point. The coverage rate is the fraction of confidence intervals containing this kink, and the rejection rate is the fraction of confidence intervals that exclude zero.

Table B.4: Summary of performance of RKD estimators in Monte Carlo Simulation, 2016, assuming very low semi-elasticity

Estimator	Median Bandwidth (1)	$\frac{RMSE}{\tau}$ (2)	Coverage Rate (3)	$\frac{Bias}{\tau}$ (4)	$\frac{Variance}{\tau^2}$ (5)	Rejection Rate (6)
A. Linear estimators						
BW = full, continuous	–	0.482	0.926	0.225	0.482	0.734
BW = FG, continuous	2337	0.671	0.934	0.249	0.671	0.477
BW = CCT, continuous	1304	0.963	0.931	0.258	0.963	0.314
BW = full, discontinuous	–	0.606	0.944	0.080	0.606	0.460
BW = FG, discontinuous	2337	0.841	0.954	0.102	0.841	0.273
BW = CCT, discontinuous	1304	1.212	0.942	0.128	1.212	0.187
B. Quadratic estimators						
BW = full, continuous	–	1.882	0.789	-2.175	1.882	0.006
BW = FG, continuous	7493	1.925	0.826	-2.161	1.925	0.004
BW = CCT, continuous	1695	3.385	0.895	-2.173	3.385	0.016
BW = full, discontinuous	–	2.363	0.720	-3.346	2.363	0.001
BW = FG, discontinuous	7493	2.412	0.759	-3.337	2.412	0.000
BW = CCT, discontinuous	1695	4.300	0.861	-3.299	4.300	0.005

Table summarizes the performance of 12 RKD estimators, which differ in the degree of the underlying polynomial (linear or quadratic), bandwidth selector (full range of data, largest symmetric band, Fan-Gjbels, or Calonico et al. (2014)), and whether a discontinuity is imposed. The data are generated using a true kink of $\tau = 3.5 \times 10^{-5}$, corresponding to a semi-elasticity of 0.2 at the mandate kink point. The coverage rate is the fraction of confidence intervals containing this kink, and the rejection rate is the fraction of confidence intervals that exclude zero.

C Additional figures

Figure C.1: Discontinuities in months insured, additional coverage types

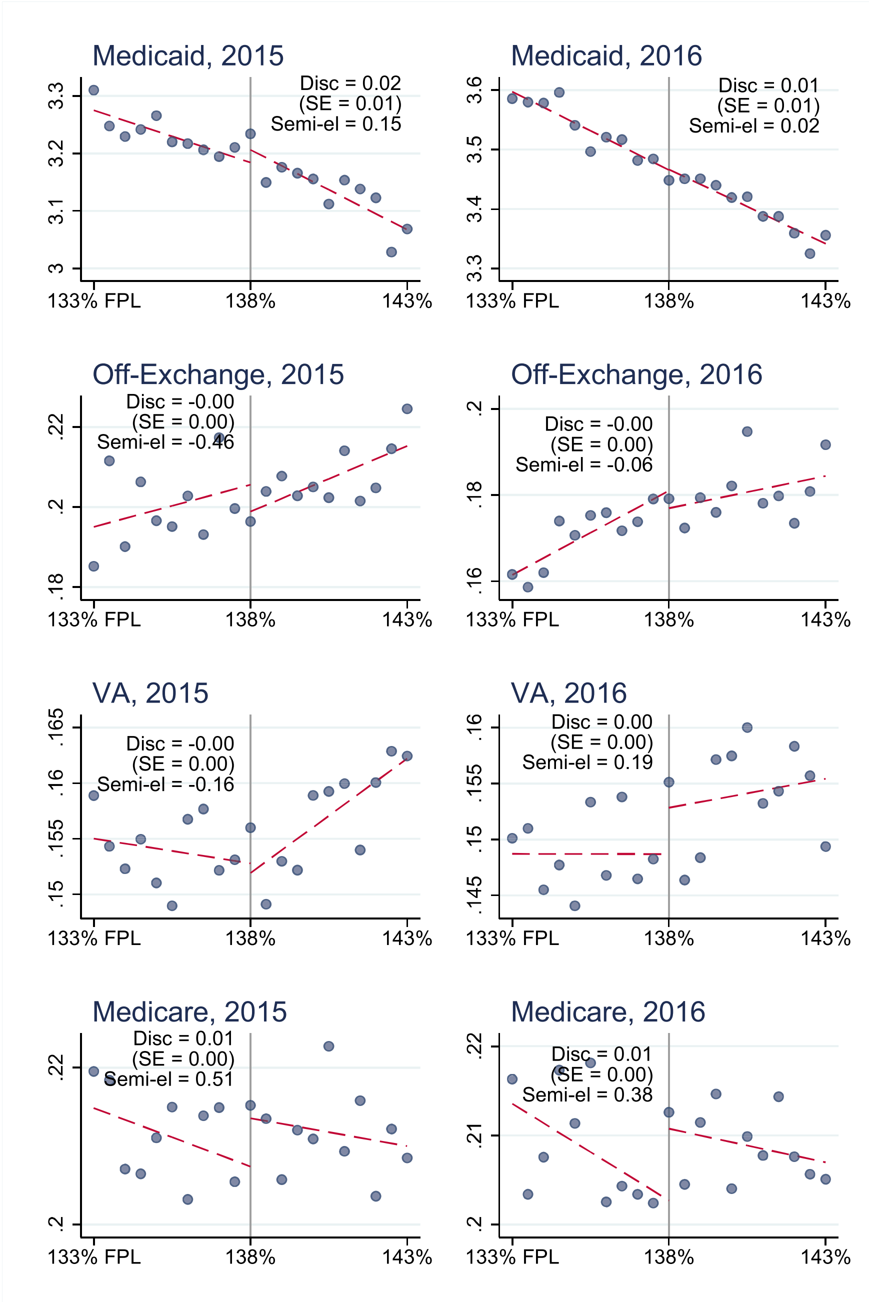

Notes: Figure shows the average number of months of insurance of the indicated type, in each 0.5 FPL point bin. Sample consists of people aged 0-64 in the indicated year, living in non-expansion states. Figure reports the estimated discontinuity, its standard error, and the implied semi-elasticity with respect to the penalty paid.

Figure C.2: Kinks in months insured, additional coverage types

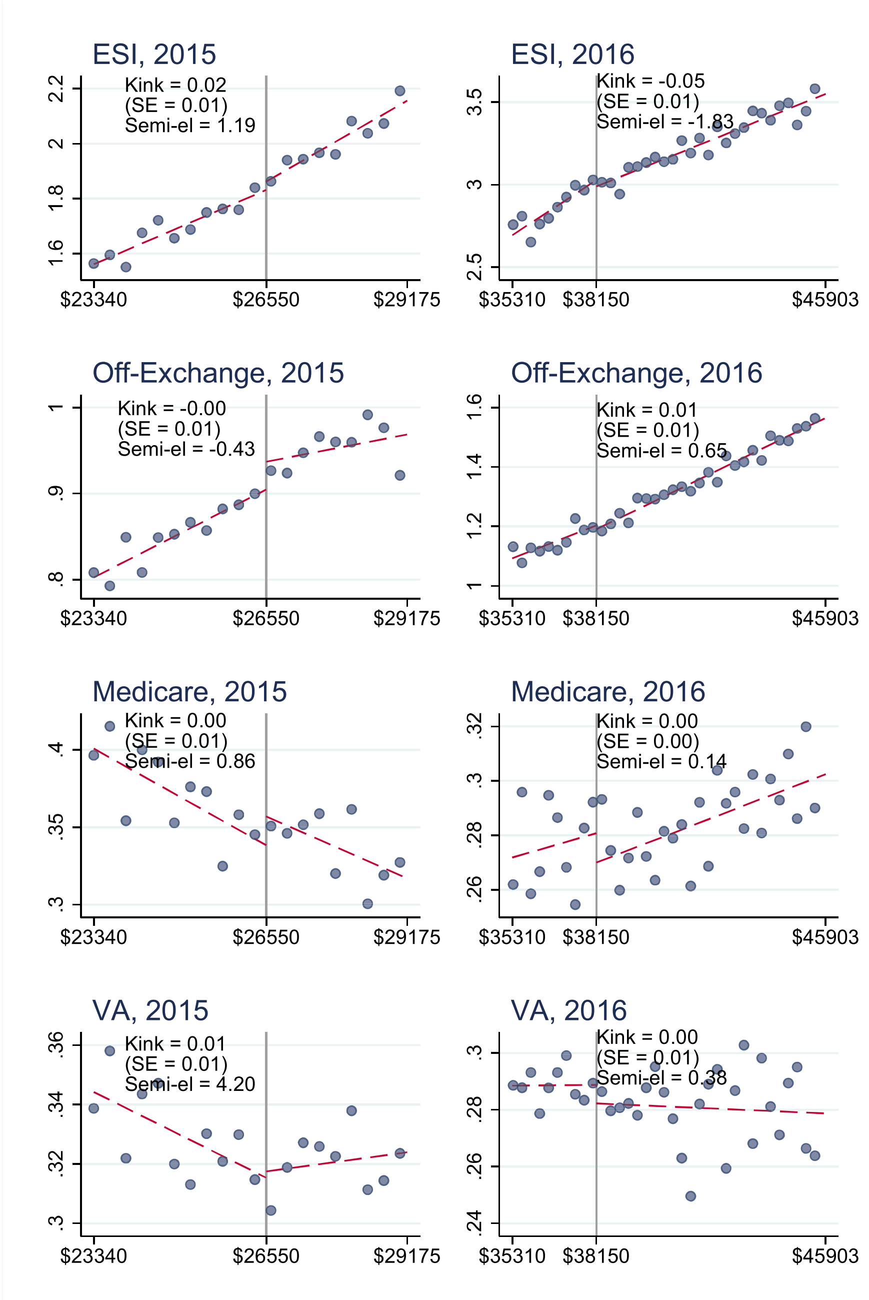

Notes: Figure shows the average number of months of insurance of the indicated type, in each $300 bin. Sample consists of people aged 27-64 in the indicated year, without signs of ESI offer, with single person tax returns and no dependents. Figure reports the estimated kink, its standard error, and the implied semi-elasticity with respect to the penalty paid.

D Further evidence on missing density

In our main discontinuity samples, we find statistically significant discontinuities in the density of the running variable. The densities do not indicate manipulation of the running variable, because there is no evidence for excess mass above the threshold. However they do indicate missing observations above the threshold. If this missing mass is driven by strategic non-filing of tax returns by uninsured people trying to avoid paying the penalty, then our RD estimates will be biased. Such strategic non-filing is unlikely to drive our results, however, because most households have a strong incentive to file a tax return, even if it means paying the penalty. Specifically, households can only claim their refundable tax credits if they file a tax return. The two most important of these credits are the earned income tax credit (EITC) and the child tax credit. The size of these credits depends on income and family structure. For households with income near 138 percent of the poverty line, we report in Appendix Table D.1 the sum of these credits, along with the maximum mandate penalty the household could face in 2016.[30] The table shows that, for many household structures, the refundable tax credits exceed the mandate penalty even in the worst case, sometimes by a large amount. Therefore there is no incentive to avoid filing, at least for some households.

This logic suggests that we should see lower filing discontinuities in households with a stronger incentive to file. In column (2) of the table, we report the discontinuity in the distribution of the running variable (in percents, so that differences are comparable across groups) in 2016. (We focus on 2016 because we have the largest first stage here.) Households with children have small distribution discontinuities, especially single-parent households. Often these discontinuities are statistically insignificant and in some cases they are positive. If differential selection explained our results, we would expect to see no coverage discontinuity where there is no density discontinuity. But as the table shows, nearly every group exhibits a coverage discontinuity, and in fact the group with the largest coverage discontinuity shows an upward discontinuity in the density of the running variable.

[30] These calculations assume that all income in MAGI is earned income and that all children in the household are 17 or younger. The credits vary with household composition for two reasons. First, the child tax credit pays $1000 per child. Second, the EITC depends on absolute income, not income relative to FPL. As the family size increases but we fix income at 138 percent of FPL, absolute income increases, because the poverty line increases. This reduces the EITC.

Table D.1: Tax credits, density discontinuities, and coverage discontinuities, 2016, by family type

Family composition		Credits & Penalty		Density discontinuity		Coverage discontinuity		
# Adults (1)	# kids (2)	Credits (3)	Penalty (4)	Estimate (5)	(SE) (6)	Estimate (7)	(SE) (8)	Semi-elasticity (9)
1	0	0	695	-0.028	(0.017)	0.308	(0.029)	0.371
1	1	3789	1042	-0.019	(0.013)	0.090	(0.027)	0.068
1	2	5622	1390	-0.031	(0.020)	0.010	(0.025)	0.007
1	3	6139	1738	0.014	(0.027)	0.143	(0.049)	0.182
1	4	5262	2085	0.047	(0.050)	0.283	(0.113)	0.895
2	0	1894	1390	-0.037	(0.015)	0.265	(0.038)	0.316
2	1	3442	1738	-0.049	(0.016)	0.249	(0.036)	0.184
2	2	3959	2085	-0.003	(0.026)	0.186	(0.032)	0.161
2	3	3779	2085	0.023	(0.024)	0.324	(0.039)	0.434
2	4	4000	2085	-0.057	(0.029)	0.132	(0.048)	0.237

Notes: Table shows the maximum tax credits (EITC plus CTC) that a family of the indicated type would be eligible for in 2016, if they had income at 138 percent of FPL, along with the estimated discontinuity in the density of the running variable (in percent) and coverage. Robust standard errors in parentheses. The sample consists of people with income between 133 and 143 percent of FPL, aged 0-64, living in non-expansion states.

E Robustness tests

Placebo test based on expansion states: To show the validity of our RD design, we begin with a placebo test looking at expansion states. In these states, there is no discontinuity in the mandate penalty at 138 percent of FPL, so if our model is well-specified, we should also see no coverage discontinuity.[31] We present RD plots for the placebo sample in Figure E.1. The top panels show the penalty paid per uninsured month. There is a discontinuity of \$0.35 in 2015 and \$0.18 in 2016, about a twentieth of the penalty discontinuity in the non-expansion states. Consistent with this small first stage, we find in the remaining panels very small discontinuities in coverage: about 0.01 months for both any coverage and verified coverage, in 2015 and in 2016. This discontinuity is statistically insignificant, often wrong-signed, and again about a twentieth of the estimate in the expansion states. Overall therefore this placebo tests shows no effect in the expansion states.

Alternative bandwidths: We consider robustness to bandwidth and to alternative specification choices. Figure E.2 shows the estimated discontinuity and its 95% confidence interval, as a function of the bandwidth. We indicate the optimal bandwidth (computed using the procedure of Calonico et al. (2014)) with a vertical line. In general the estimates are not too sensitive to the bandwidth, although with a very small bandwidth we would find insignificant estimates in 2015.

We plot the analogous figure for the 2015 kink estimate in Figure E.3. The point estimate is stable over a wide range of bandwidths. At small enough bandwidths, the point estimate fluctuates and its confidence interval becomes quite large. The mean-squared error optimal bandwidth of Calonico et al. (2014) is about \$900. At this bandwidth, the point estimate is 0.17, thrice as large as our main estimate, but the confidence interval is larger still, and so the estimate is marginally significant ($p = 0.08$). Looking across the different coverage types, the point estimates are fairly stable until the bandwidth becomes small, at which point the estimates become less stable and much less precise.[32]

Alternative specifications: We consider robustness to alternative specification choices: allowing for nonlinearities in income (i.e. a quadratic or cubic), controlling for demographics (female dummy and a quadratic in age, plus filing status dummies and dummies for number of exemptions in the RD sample), imposing continuity at the kink point, and excluding people with ESI in the RK sample. The results are in Appendix Tables E.1 to E.4. The RD estimates are generally robust to alternative functional forms. In 2015, the point estimates are similar for the linear and quadratic specifications, but fall a bit with the cubic. In 2016 however the cubic specification yields slightly higher point estimates. The RD results are also unchanged when we include controls.

For the kink sample, the nonlinear income terms generally produce larger estimates,

[31] Of course, in these states, eligibility for Medicaid for childless adults ends at 138 percent of FPL, so we might expect to see a downward discontinuity. However, as we emphasize elsewhere, Medicaid eligibility is not determined by annual taxable income; it is determined by recent income as reported to state Medicaid authorities. We do not expect to see a downward discontinuity in Medicaid coverage at this threshold, and indeed we do not find one.

[32] We do not plot the 2016 kink estimates against the bandwidth because that figure only confirms that the 2016 estimates are not statistically significant.

and sometimes substantially larger ones. For example the kink in any coverage in 2015 increases by 60 percent when we include a quadratic term, and it more than doubles when we include a cubic. The standard errors also rise, consistent with the findings from our Monte Carlo analysis that these nonlinear terms substantially increase the MSE of our estimates. We focus on the linear specifications because of these large standard errors, although we find it reassuring that allowing for higher order terms would, if anything, strengthen the conclusion that the mandate penalty increases coverage. In general the results are not sensitive to other choices of controls. They change little when we impose continuity or add demographic controls. The estimates typically rise when we drop people with ESI coverage.

Extensive margin responses: As a final robustness check, we show that our results are not sensitive to our focus on months of coverage. We re-estimate our models but looking at the extensive margin—the probability of having at least one month of coverage of a given type. Generally we find very similar results. In RD samples, the estimated semi-elasticities are about 0.17 (for any coverage) and 0.13 (for verified coverage), which are similar to but slightly smaller than the estimates using months of coverage. The regression kink estimates show a similar pattern; the semi-elasticities are roughly similar in magnitude whether we examine months of coverage or the extensive margin of any coverage.

Permutation test for the RKD We have found clear kinks in months of any insurance coverage and months of verified coverage in 2015. One concern with the regression kink approach, however, is that it may detect spurious kinks, simply due to curvature in the relationship between the outcome and running variable (Ganong and Jäger, 2018). We assess this concern by re-estimating our RKD models, but varying the kink point across a fine grid of placebo locations. If the kink is spurious, then we expect that our estimate is unexceptional in the distribution of placebo estimates.

Figure E.4 shows the distribution of placebo kink estimates, for any coverage and for verified coverage. We consider permutation kinks every \$25, starting from \$500 above 200 percent of FPL, and ending at \$500 below 250 percent of FPL. We look in this range because we do not believe looking elsewhere in the income distribution would be informative about the possibility of a false positive at our income level. There are likely to be other policy-induced kinks elsewhere in the income distribution (for example, because of the PTC). We exclude kink points near the boundaries because estimating a kink near the boundary produces very large, very noisy estimates, because there is very little data with which to estimate a slope on one side of the kink.

The histograms show, first, a long left tail of placebo kink points. This is generated by the fact that placebo kink locations near the boundaries tend to produce large, negative placebo estimates. Second, the estimated kink, shown with the vertical line, is larger than all but a handful of the placebo kinks. The implied p-value—the fraction of placebo point estimates that exceed the true point estimate—is 0.078 for any coverage and 0.098 for verified coverage. The reader may worry that these p-values are small in part because of the inclusion of the many very negative placebo kink points estimated near the boundary. If we instead estimate p-values, but excluding placebo kink points within \$1000 of the boundary, we obtain p-values of 0.058 for any coverage and 0.118 for verified coverage.

Figure E.1: Placebo test in expansion states

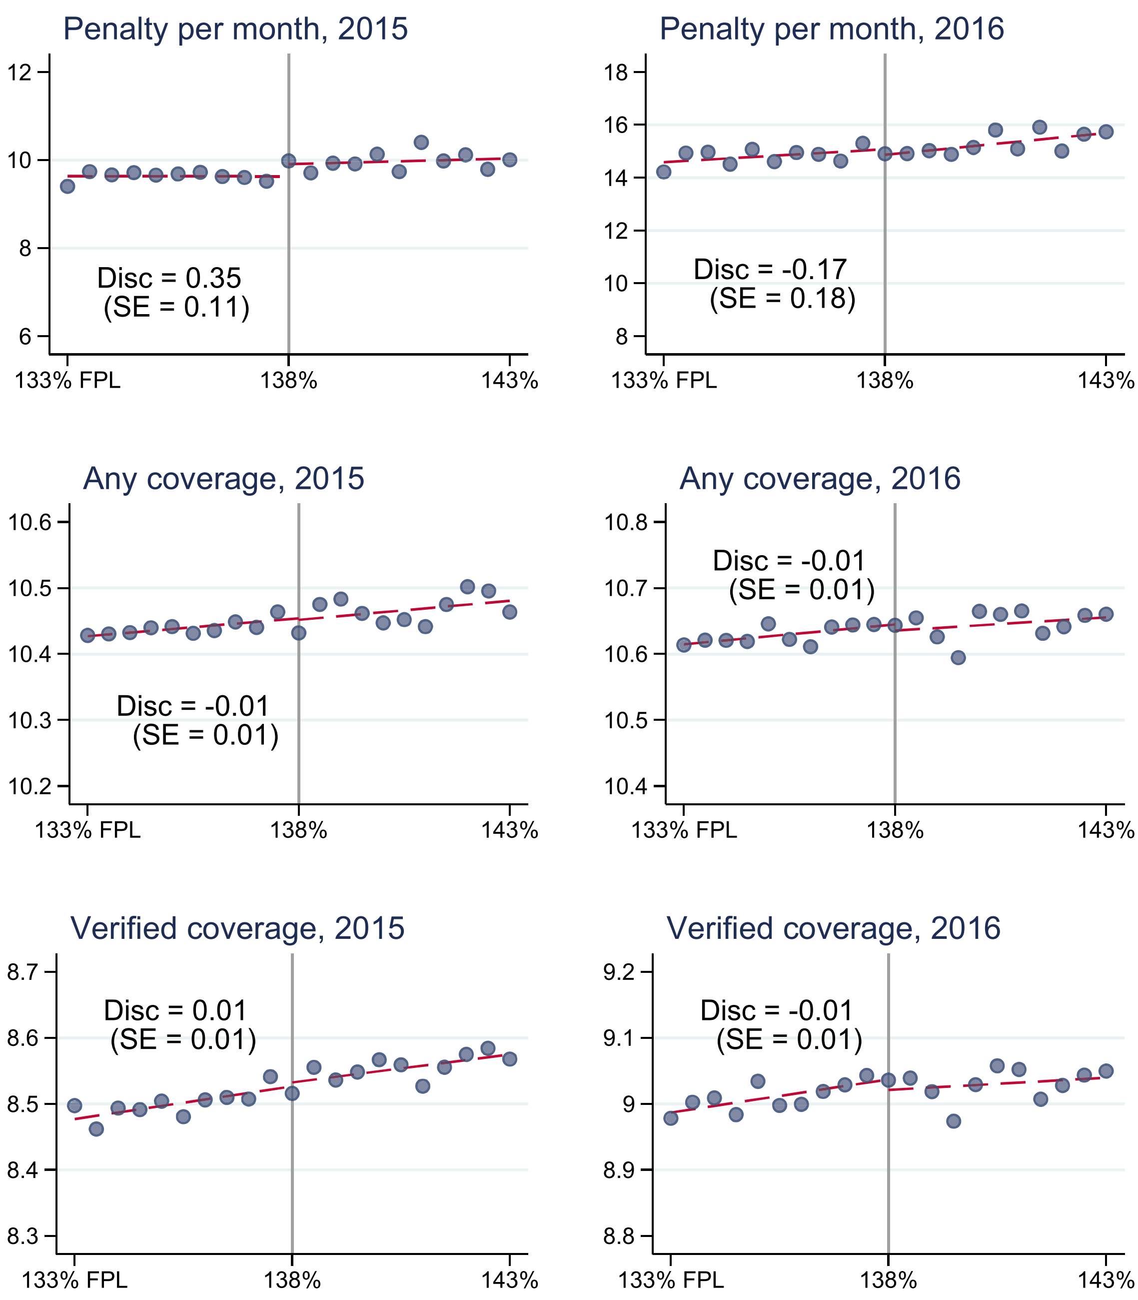

Notes: Figure shows the average of the indicated outcome in each 0.5 FPL point bin. This is a placebo test because, in expansion states, people on both sides of 138 percent of FPL are subject to the mandate. Sample consists of people aged 0-64 in the indicated year, living in Medicaid expansion states. Figure reports the estimated discontinuity and its standard error.

Figure E.2: Estimated discontinuity as a function of bandwidth

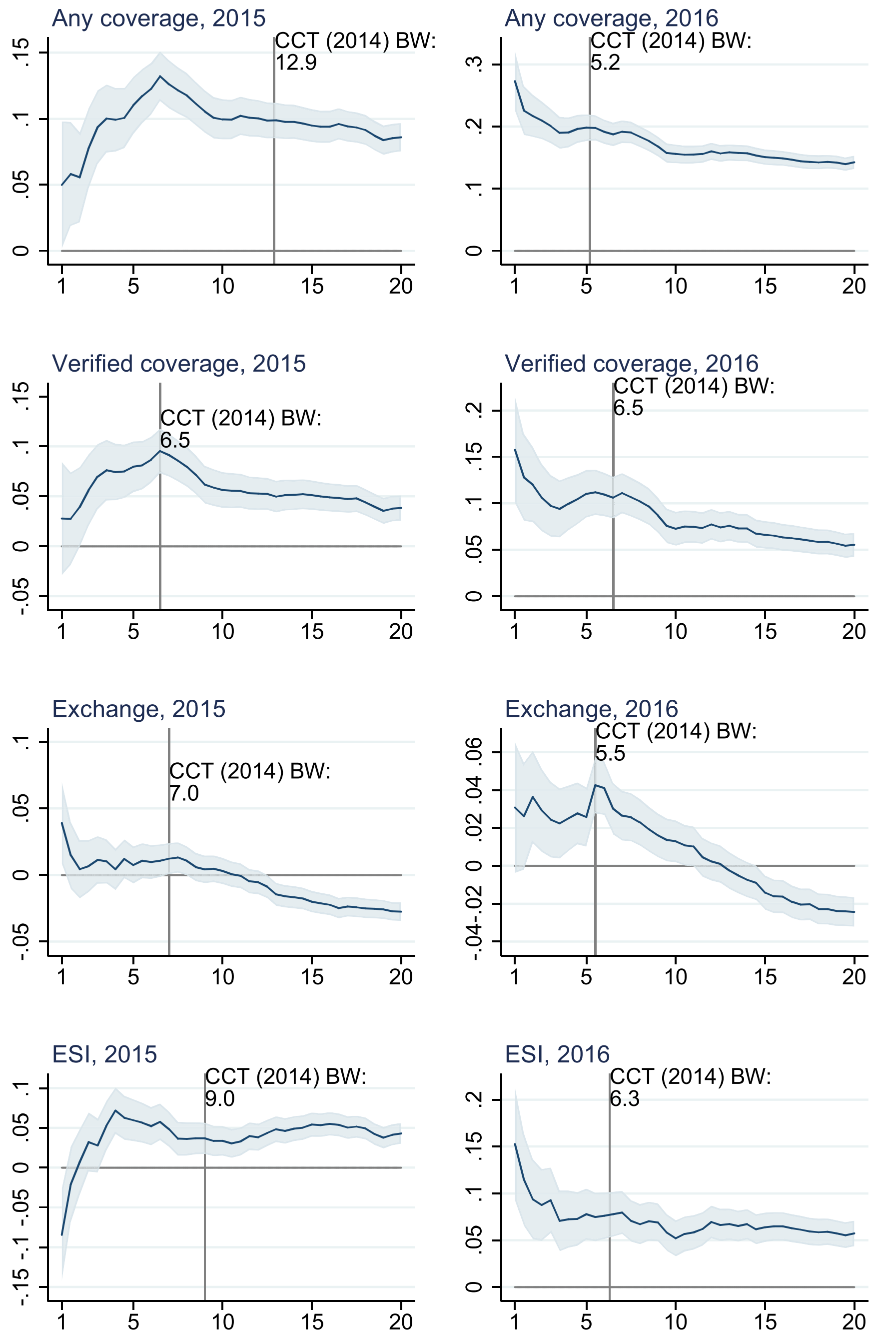

Notes: Figure shows estimated discontinuity (and 95% confidence interval) in months of coverage of the indicated type, as a function of bandwidth, for the indicated types of coverage. The vertical line is the MSE-optimal bandwidth of Calonico et al. (2014). The estimates in the paper use a bandwidth of 5, the largest symmetric bandwidth that avoids looking across the 133 percent of FPL discontinuity in the premium tax credit.

Figure E.3: Estimated kink as a function of bandwidth, 2015

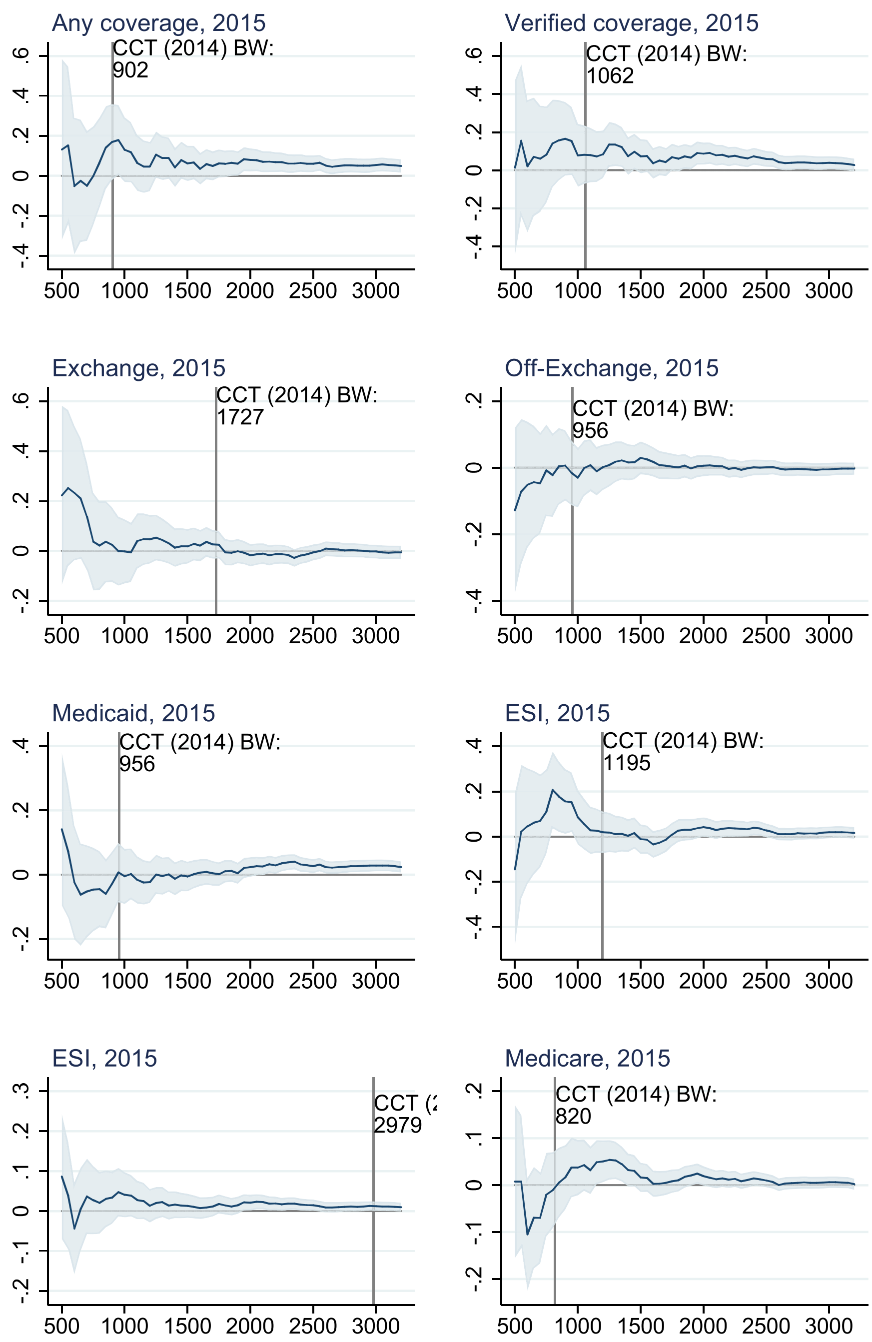

Notes: Figure shows estimated kink (and 95% confidence interval) in 2015 months covered insured as a function of bandwidth, for the indicated types of coverage. The CCT (2014) bandwidth is the MSE-optimal bandwidth of Calonico et al. (2014).

Figure E.4: Estimated kink in months insured at placebo kink points, 2015

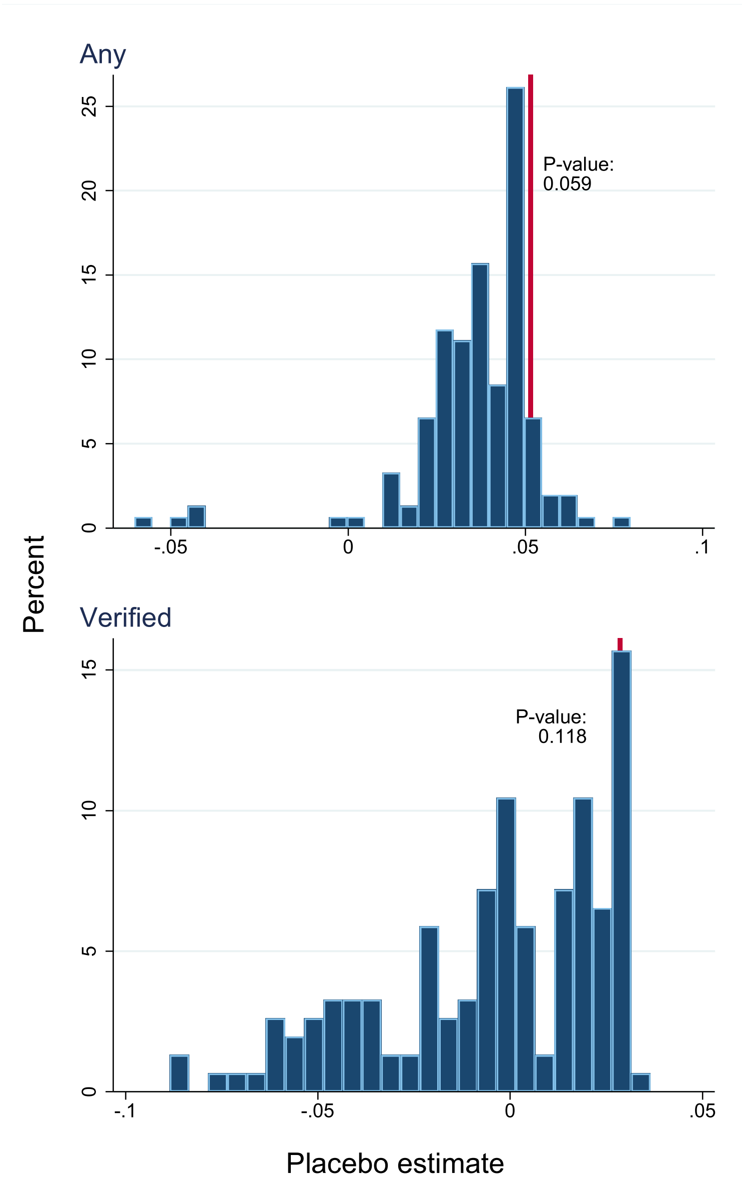

Notes: Figure shows estimated kink in 2015 months insured at placebo kink points. The p-value is the fraction of placebo kinks that exceed the true estimate.

Table E.1: Robustness of 2015 RD estimates to specification choices

	(1)	(2)	(3)	(4)
A. Y = Months any coverage				
Discontinuity	0.110	0.072	0.058	0.105
	(0.011)	(0.017)	(0.022)	(0.011)
Semi-elasticity	0.23	0.15	0.12	0.20
B. Y = Months verified coverage				
Discontinuity	0.080	0.055	0.028	0.074
	(0.013)	(0.019)	(0.026)	(0.013)
Semi-elasticity	0.20	0.14	0.07	0.16
D. Y = Months Exchange				
Discontinuity	0.008	0.012	0.009	0.012
	(0.007)	(0.011)	(0.014)	(0.007)
Semi-elasticity	0.19	0.29	0.22	0.55
C. Y = Months ESI				
Discontinuity	0.060	0.034	-0.040	0.065
	(0.013)	(0.020)	(0.027)	(0.013)
Semi-elasticity	0.31	0.17	-0.20	0.68
F. Y = Months Medicaid				
Discontinuity	0.024	-0.002	0.041	0.006
	(0.013)	(0.019)	(0.025)	(0.011)
Semi-elasticity	0.15	-0.01	0.26	0.02
E. Y = Months off-Exchange				
Discontinuity	-0.005	-0.007	-0.012	-0.004
	(0.004)	(0.006)	(0.007)	(0.004)
Semi-elasticity	-0.46	-0.65	-1.17	-0.27
Degree	Linear	Quadratic	Cubic	Linear
Controls	No	No	No	Yes

Notes: Table shows robustness of the regression discontinuity coverage estimates to alternative specifications (polynomial degree) or controls. Column (1) is the base estimates. The degree specification controls for polynomials of the indicated degree, allowed to vary on either side of the discontinuity. The controls in column (4) are a female dummy, a quadratic in age, and dummies for filing status and number of exemptions. The sample consists of people with income between 133 and 143 percent of FPL, aged 0-64, living in non-expansion states.

Table E.2: Robustness of 2016 RD estimates to specification choices

	(1)	(2)	(3)	(4)
A. Y = Months any coverage				
Discontinuity	0.198	0.203	0.243	0.194
	(0.011)	(0.016)	(0.022)	(0.011)
Semi-elasticity	0.22	0.24	0.29	0.19
B. Y = Months verified coverage				
Discontinuity	0.110	0.098	0.128	0.106
	(0.013)	(0.020)	(0.026)	(0.013)
Semi-elasticity	0.14	0.14	0.18	0.12
D. Y = Months Exchange				
Discontinuity	0.026	0.029	0.029	0.031
	(0.008)	(0.012)	(0.016)	(0.008)
Semi-elasticity	0.30	0.36	0.36	0.64
C. Y = Months ESI				
Discontinuity	0.078	0.088	0.118	0.085
	(0.014)	(0.021)	(0.028)	(0.014)
Semi-elasticity	0.21	0.25	0.35	0.55
F. Y = Months Medicaid				
Discontinuity	0.006	-0.020	-0.034	-0.015
	(0.013)	(0.020)	(0.027)	(0.011)
Semi-elasticity	0.02	-0.07	-0.11	-0.02
E. Y = Months off-Exchange				
Discontinuity	-0.001	-0.002	0.002	-0.000
	(0.004)	(0.005)	(0.007)	(0.004)
Semi-elasticity	-0.06	-0.16	0.15	-0.00
Degree	Linear	Quadratic	Cubic	Linear
Controls	No	No	No	Yes

Notes: Table shows robustness of the regression discontinuity coverage estimates to alternative specifications (polynomial degree) or controls. Column (1) is the base estimates. The degree specification controls for polynomials of the indicated degree, allowed to vary on either side of the discontinuity. The controls in column (4) are a female dummy, a quadratic in age, and dummies for filing status and number of exemptions. The sample consists of people with income between 133 and 143 percent of FPL, aged 0-64, living in non-expansion states.

Table E.3: Robustness of 2015 RK estimates to specification choices

	(1)	(2)	(3)	(4)	(5)	(6)
A. Y = Months any coverage						
Kink	0.052	0.112	0.073	0.050	0.047	0.056
	(0.017)	(0.069)	(0.171)	(0.017)	(0.017)	(0.019)
Semi-elasticity	0.92	1.97	1.27	0.87	0.59	1.10
B. Y = Months verified coverage						
Kink	0.029	0.210	0.061	0.030	0.023	0.027
	(0.017)	(0.070)	(0.174)	(0.017)	(0.017)	(0.019)
Semi-elasticity	0.64	4.58	1.33	0.64	0.30	0.68
C. Y = Months Exchange						
Kink	-0.004	0.021	0.067	-0.007	-0.008	-0.001
	(0.013)	(0.054)	(0.135)	(0.013)	(0.013)	(0.016)
Semi-elasticity	-0.24	1.23	3.97	-0.42	-0.51	-0.03
D. Y = Months ESI						
Kink	0.017	0.054	0.103	0.020	0.016	0.007
	(0.013)	(0.051)	(0.129)	(0.013)	(0.013)	(0.005)
Semi-elasticity	1.21	3.94	7.49	1.43	0.43	5.17
E. Y = Months Medicaid						
Kink	0.023	0.027	-0.169	0.021	0.023	0.020
	(0.009)	(0.036)	(0.089)	(0.009)	(0.009)	(0.010)
Semi-elasticity	3.45	4.06	-24.81	3.23	3.78	2.67
F. Y = Months off-Exchange						
Kink	-0.003	0.026	0.004	-0.001	-0.003	0.003
	(0.010)	(0.038)	(0.095)	(0.009)	(0.010)	(0.011)
Semi-elasticity	-0.40	3.80	0.53	-0.16	-0.44	0.38
Degree	Linear	Quadratic	Cubic	Linear	Linear	Linear
Controls	No	No	No	No	Yes	No
Discontinuity	Yes	Yes	Yes	No	Yes	Yes
Include ESI?	Yes	Yes	Yes	Yes	Yes	No

Notes: Table shows robustness of the regression kink coverage estimates to alternative specifications (polynomial degree), controls, and samples. Column (1) is the base estimates. The degree specification controls for polynomials of the indicated degree, allowed to vary on either side of the kink. In column (4) we impose continuity. The controls in column (5) are a female dummy and a quadratic in age. The sample consists of people with income between 200 and 250 percent of FPL, aged 27-64, without signs of ESI offers, single filing status, and one exemption.

Table E.4: Robustness of 2016 RK estimates to specification choices

	(1)	(2)	(3)	(4)	(5)	(6)
A. Y = Months any coverage						
Kink	0.020	-0.079	-0.161	0.026	0.016	0.050
	(0.015)	(0.060)	(0.150)	(0.012)	(0.014)	(0.018)
Semi-elasticity	0.29	-1.14	-2.31	0.37	0.16	0.84
B. Y = Months verified coverage						
Kink	-0.002	-0.164	-0.029	0.008	-0.007	0.030
	(0.016)	(0.063)	(0.158)	(0.013)	(0.015)	(0.018)
Semi-elasticity	-0.04	-2.84	-0.51	0.14	-0.08	0.67
C. Y = Months Exchange						
Kink	0.021	0.056	-0.126	0.014	0.019	0.013
	(0.012)	(0.048)	(0.120)	(0.009)	(0.012)	(0.015)
Semi-elasticity	1.31	3.54	-7.78	0.88	1.41	0.61
D. Y = Months ESI						
Kink	-0.046	-0.193	0.005	-0.036	-0.049	-0.004
	(0.014)	(0.056)	(0.140)	(0.011)	(0.014)	(0.005)
Semi-elasticity	-1.87	-7.59	0.18	-1.47	-0.85	-1.94
E. Y = Months Medicaid						
Kink	0.012	-0.024	0.044	0.014	0.012	0.015
	(0.007)	(0.027)	(0.067)	(0.005)	(0.007)	(0.009)
Semi-elasticity	2.63	-5.03	9.47	3.14	2.12	2.61
F. Y = Months off-Exchange						
Kink	0.006	-0.063	-0.062	0.012	0.006	0.008
	(0.010)	(0.039)	(0.098)	(0.008)	(0.010)	(0.013)
Semi-elasticity	0.64	-6.20	-6.12	1.23	0.74	0.61
Degree	Linear	Quadratic	Cubic	Linear	Linear	Linear
Controls	No	No	No	No	Yes	No
Discontinuity	Yes	Yes	Yes	No	Yes	Yes
Include ESI?	Yes	Yes	Yes	Yes	Yes	No

Notes: Table shows robustness of the regression kink coverage estimates to alternative specifications (polynomial degree), controls, and samples. Column (1) is the base estimates. The degree specification controls for polynomials of the indicated degree, allowed to vary on either side of the kink. In column (4) we impose continuity. The controls in column (5) are a female dummy and a quadratic in age. The sample consists of people with income between 300 and 390 percent of FPL, aged 27-64, without signs of ESI offers, single filing status, and one exemption.

Table E.5: Estimated discontinuities and kinks in extensive margin of coverage

Coverage type	Any (1)	Verified (2)	Exchange (3)	Off-Exchange (4)	ESI (5)	Medicaid (6)	Medicare (7)	VA (8)
				A. 2015 Discontinuity Sample				
Discontinuity	0.007	0.005	0.000	-0.001	0.005	0.002	0.001	-0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)
Semi-elasticity	0.17	0.13	0.01	-0.58	0.27	0.12	0.60	-0.17
				B. 2016 Discontinuity Sample				
Discontinuity	0.014	0.010	0.002	-0.000	0.007	0.001	0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)
Semi-elasticity	0.17	0.13	0.25	-0.07	0.20	0.03	0.45	0.19
				C. 2015 Kink Sample				
Kink	0.0048	0.0033	-0.0008	-0.0000	0.0012	0.0022	0.0009	0.0002
	(0.0014)	(0.0015)	(0.0013)	(0.0009)	(0.0012)	(0.0009)	(0.0005)	(0.0005)
Semi-elasticity	0.96	0.77	-0.46	-0.07	0.93	2.97	4.25	0.87
				D. 2016 Kink Sample				
Kink	0.0014	0.0003	0.0016	0.0010	-0.0039	0.0014	0.0001	-0.0000
	(0.0012)	(0.0013)	(0.0011)	(0.0009)	(0.0012)	(0.0007)	(0.0004)	(0.0004)
Semi-elasticity	0.24	0.06	1.00	1.13	-1.69	2.60	0.69	-0.04

Notes: The dependent variable is an indicator for at least one month of coverage of the indicated type. The table reports RD and RK estimates. The RD sample consists of people aged 0-64 with income between 133 and 143 percent of FPL, living in non-expansion states. The RK sample consists of people with single tax returns, one exemption, aged 27-64, without sings of an ESI offer, and income between 200 and 250 percent of FPL (in 2015) or 300 and 390 percent of FPL (in 2016).

F Exploring the ESI offer sample

Our main kink sample excludes people with signs of an ESI offer. Hence, our main estimates do not reflect two channels through which a greater mandate penalty could affect insurance coverage: by changing ESI offer rates, or by changing take up among people offered ESI. Here we present evidence that neither of these channels is quantitatively important, at least in the context of the kink. To do so, we expand our kink samples to include people with signs of ESI offers.

We begin by showing that there is no kink in the probably of having a sign of ESI offer at the mandate kink point. Appendix Figure F.1 plots the fraction of people who have a sign of an ESI offer, as a function of income. In both 2015 and 2016 the offer rate is increasing in income but essentially smooth through the mandate kink point. In 2015 the estimated kink is -0.06, with a standard error of 0.18, meaning that each thousand dollars of income above the kink point reduces the offer rate by 0.06 percentage points—an economically small, statistically insignificant and wrong-signed amount. In 2016, the estimated kink is -0.21 (standard error=0.13)—larger in absolute value, but still small and statistically insignificant. This evidence shows that a kink in the ESI offer rate does not bias our kink sample results (through endogenously changing our sample in the neighborhood of the kink). We emphasize that this evidence is not particularly informative about the effect of the individual mandate on ESI offers, as it is unlikely that employers could tailor ESI offer to individual employees, and so we should not expect to see a sharp kink in the ESI offer rate.

Next we show the effect of a greater mandate penalty among people with signs of an ESI offer. Appendix Figure F.2 shows months of any coverage among people with signs of an ESI offer. The figure reveals several important patterns. First, and unsurprisingly, coverage is much higher in the offer sample than in the no-offer sample. Second, there is some evidence curvature. Looking above the kink point only, for example, the slope seems to decline as income rises. Third, looking locally to the kink point, the slope in 2015 appears flatter below the kink than above it, although in 2016 the pattern is ambiguous. Taken together, these patterns indicate that perhaps there is a slight kink in 2015, but any kink will be difficult to detect and may be sensitive to the assumed polynomial and bandwidth.

In Appendix Table F.1, we report estimated RK models for the offer sample. We begin in column (1) with the linear specification using the full range of data, which is our main specification for the no-offer sample. The estimates are negative and statistically significant in both years. However, the linear/full range of data specification may not be appropriate for the ESI offer sample, with the evident curvature and larger sample size. When we use the asymptotically optimal bandwidth, the 2015 kink becomes positive and statistically insignificant (but not small, with a semi-elasticity of 0.33); the 2016 kink remains similar (because the optimal bandwidth uses nearly all the data).

One potential concern with the specifications in columns (1) and (2) is that they do not adequately control for the concavity evident in Figure F.2. This concavity could bias us toward finding a negative kink. In the remaining columns we therefore control for quadratic and cubic functions of the running variable. The quadratic specification produces kinks and positive point estimates in 2015 and negative estimates in 2016, both insignificant. However the 2016 first stage disappears when we use control for a quadratic in income. Likewise we see in columns (5) and (6) that the first stage is gone in both 2015 and in 2016 when we control for a cubic in income. It is therefore impossible to interpret the reduced form kink in these specifications.

Overall, therefore, the ESI results are sensitive to specification choices. Based on the concavity in Figure F.2, we would prefer specifications that control for higher order terms. These specifications yield insignificant first stages and coverage kinks that are very noisy. Given this ambiguous evidence, we remain agnostic about the effect of the penalty among people with ESI offers at higher incomes. At lower incomes, we pool the offer and non-offer sample, and we find significant and positive effects on ESI coverage, suggesting that the mandate raises coverage even in ESI, at least at lower income levels.

Figure F.1: Probability of sign of ESI offer as a function of income

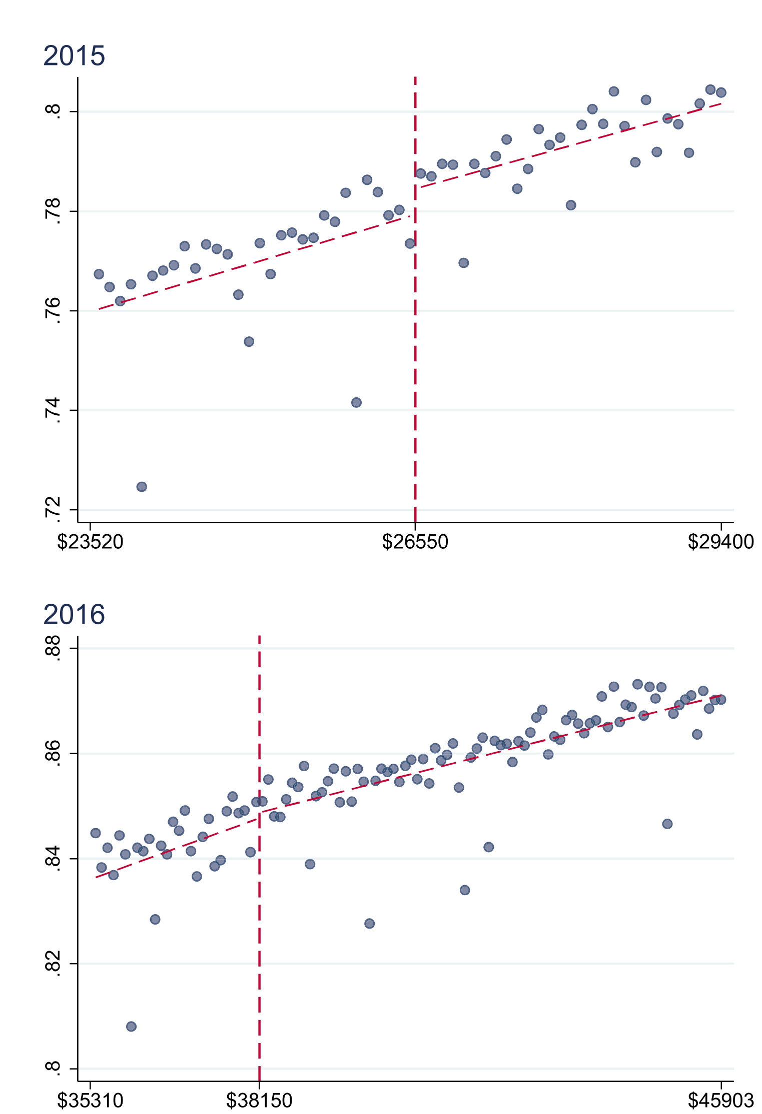

Notes: Figure shows the fraction of people with a sign of ESI offer, in each $100 bin of modified adjusted gross income. Sample consists of people aged 27-64 in the indicated year, who filed single tax returns with no dependents. The hollow circles indicate round number incomes ($1000 multiples); such incomes are much more common among the self-employed, who lack signs of an ESI offer. We include dummies for "round number incomes" when estimate RKD models for signs of an ESI offer.

Figure F.2: Months insured as a function of income, signs-of-ESI-offer sample

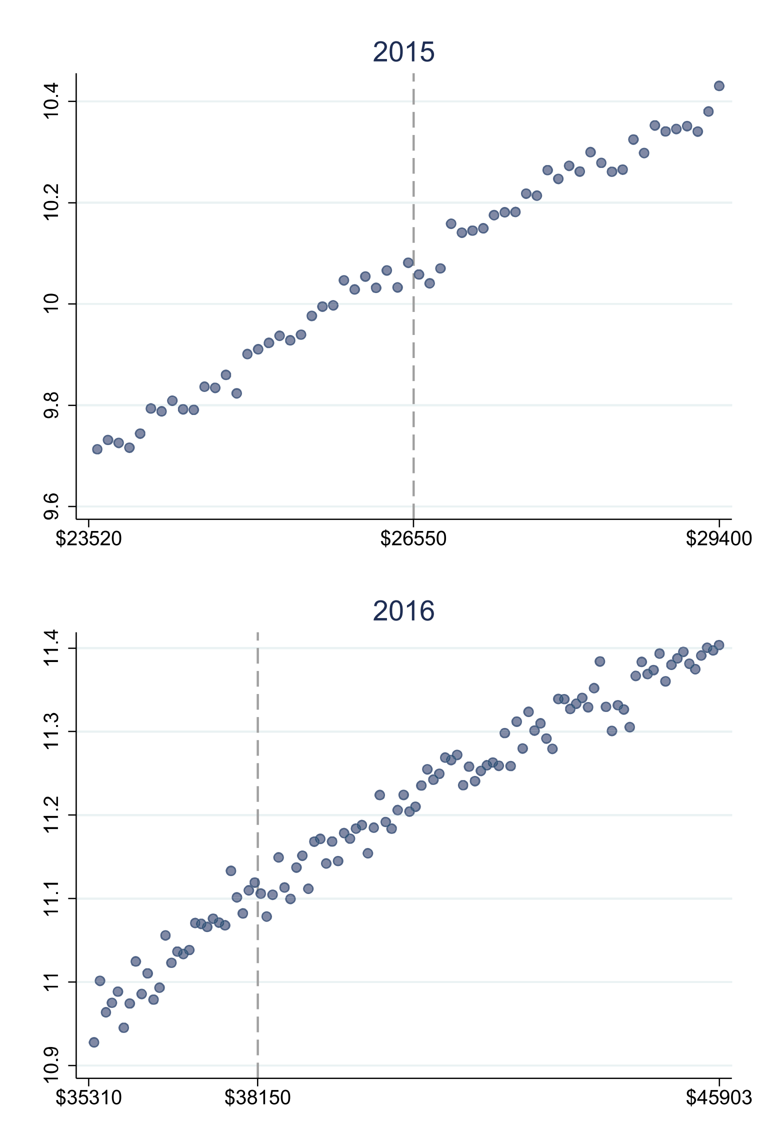

Notes: Figure shows the average number of months of any insurance, in each $100 bin of modified adjusted gross income. Sample consists of people aged 27-64 in the indicated year, who filed single tax returns with no dependents, but with signs of ESI offers.

Table F.1: RK estimates in the ESI offer sample

Degree	Linear		Quadratic		Cubic	
Bandwidth	Full (1)	CCT (2)	Full (3)	CCT (4)	Full (5)	CCT (6)
A. 2015, Y= penalty per month						
Kink	0.750	0.750	1.221	1.221	0.762	0.762
	(0.085)	(0.085)	(0.332)	(0.332)	(0.827)	(0.827)
BW	2918	9020	2918	4192	2918	8744
B. 2015, Y= Months any coverage						
Kink	-0.020	-0.018	0.045	0.121	0.118	0.161
	(0.007)	(0.009)	(0.026)	(0.046)	(0.065)	(0.068)
Semi-elasticity	-0.258	-0.244	0.369	1.731	1.537	1.115
BW	2918	2260	2918	1959	2918	2994
C. 2015, Y= Months verified coverage						
Kink	-0.028	-0.027	0.030	0.026	0.085	0.085
	(0.008)	(0.009)	(0.031)	(0.037)	(0.077)	(0.077)
Semi-elasticity	-0.427	-0.488	0.284	0.204	1.284	1.284
BW	2918	2596	2918	2523	2918	3790
D. 2016, Y= penalty per month						
Kink	0.780	0.296	0.020	0.020	-0.963	-0.963
	(0.161)	(0.210)	(0.629)	(0.629)	(1.442)	(1.442)
BW	5296	3029	5296	8514	5296	13931
E. 2016, Y= Months any coverage						
Kink	-0.022	-0.018	-0.023	-0.021	-0.004	-0.004
	(0.004)	(0.004)	(0.014)	(0.015)	(0.035)	(0.035)
Semi-elasticity	-0.258	-0.250	-10.702	1.025	0.041	0.035
BW	5296	5334	5296	4892	5296	7367
F. 2016, Y= Months verified coverage						
Kink	-0.036	-0.035	-0.051	-0.056	-0.057	-0.062
	(0.005)	(0.010)	(0.019)	(0.022)	(0.047)	(0.048)
Semi-elasticity	-0.452	-1.105	-25.428	1.507	0.580	0.546
BW	5296	2072	5296	3860	5296	6168

Table reports the estimated kink for the indicated year and outcome. For each outcome, we report linear, quadratic, and cubic specifications (meaning we control for separate linear, quadratic, or cubic specifications), and we report estimates using the full range of data and the bandwidth of Calonico et al. (2014) ("CCT"). When calculate coverage semi-elasticities, we use the same bandwidth for calculating the first stage as we use for estimating the kink. The sample consists of single tax returns in 2015 or 2016 with one exemption claimed and a sign of an ESI offer, aged 27-64, with income between 200 and 250 percent of FPL (in 2015), or 300 and 390 percent of FPL (in 2016). Robust standard errors in parentheses. The bandwidth we report in the full section is half the full range, but it is not symmetric.

G Digging into the Medicaid response

We dig further into the large Medicaid response and small individual market response by re-estimating our RK models, stratifying on Medicaid expansion status. We expect larger Medicaid responses in states that expanded Medicaid, and larger individual market responses in non-expansion states. We report the estimates in Appendix Table G.1. The Medicaid response occurs almost entirely in Medicaid expansion states, in both 2015 and in 2016. Likewise, in 2016, we find the Exchange and off-Exchange responses are concentrated in non-expansion states. In 2015, however, we find no individual market response even in non-Expansion states. The point estimates are all statistically insignificant, fairly small, and some are wrong signed. These estimates suggest that people respond to a greater mandate penalty by obtaining Medicaid coverage if at all possible. Only at fairly high income levels and in non-expansion states do we see an individual market response.

The substantial Medicaid response raise an important question: how is it that people with income above 200 percent of FPL obtain Medicaid coverage? Medicaid eligibility is assessed based on rolling income, with infrequent recertification, rather than on realized annual income. It is likely that people in our sample obtain Medicaid coverage because they found or lost a job during the year, and were temporarily eligible for Medicaid. We expect to see the biggest increases in partial year Medicaid coverage—people with a few more months of coverage, rather than an increase from 0 to 12 months of coverage. To test this hypothesis, we estimate regressions of the form

$$Pr(\text{Medicaid Months}_i <= m) = \beta_0^m + \beta_1^m v_i + \beta_2^m 1\{v_i \geq 0\} + \beta_3^m v_i 1\{v_i \geq 0\} + \varepsilon_i^m. \quad (4)$$

This is an RKD where the dependent variable is an indicator for having at most m months of Medicaid coverage. We expect to find larger effects on the probability having an intermediate number of months of coverage (1-11). This implies that we should find less negative kinks as m grows larger. We present the estimates graphically in Appendix Figure G.1, and we report the estimated kinks in Appendix Table G.2. The effect is largest for 0-5 months of coverage. Specifically we show the baseline CDF at the 2015 kink point, and the new CDF induced by a \$10 per month increase in the mandate penalty, along with the new CDF's 95% confidence interval. The baseline CDF is given by the estimates of β_m^0 from Equation 4. We obtain the new CDF by adding the implied effect of a \$10 penalty increase to the baseline CDF. The new CDF is lower everywhere than the old CDF, implying that the penalty shifts people towards more months of Medicaid. However the distance between the CDFs is greatest for relatively low months of coverage. The mandate penalty increases months of Medicaid coverage primarily at the bottom end of the coverage spectrum, pulling people up from zero months of coverage to 1-6 months coverage, with a relatively smaller effect higher up.

Figure G.1: CDF of months of Medicaid coverage, 2015

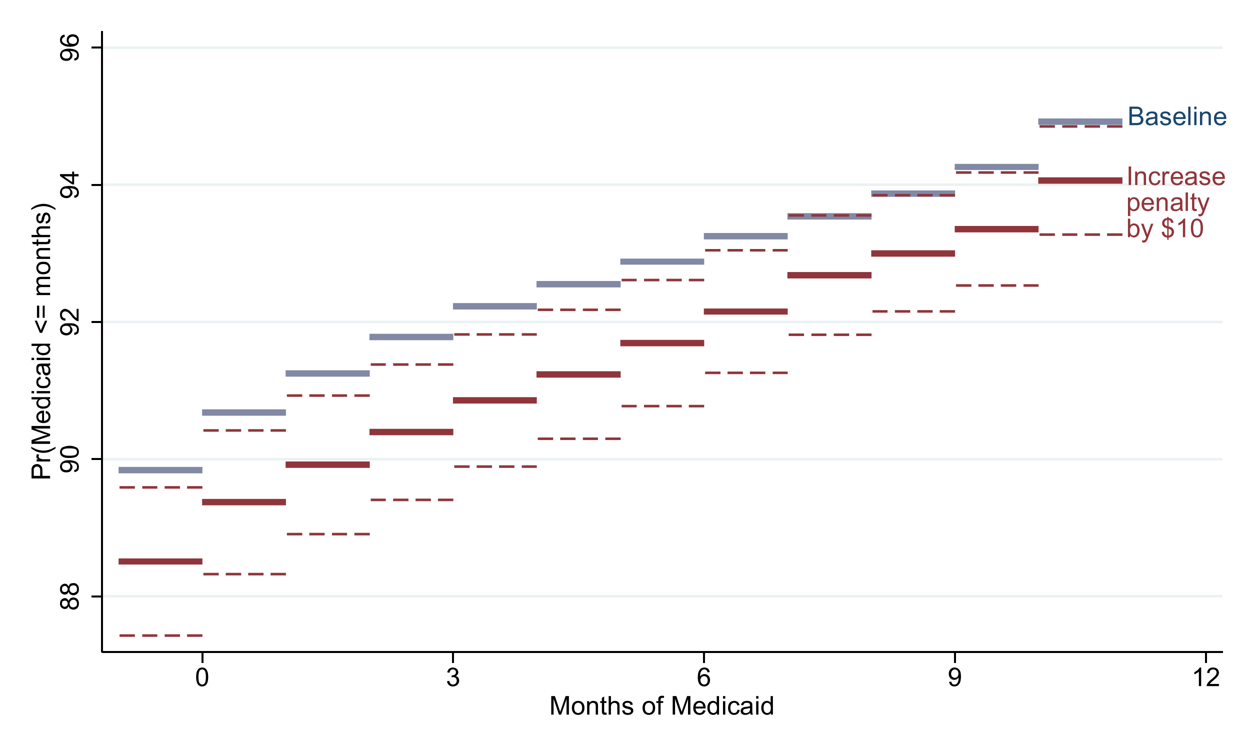

Notes: Source: Figure shows the CDF of months of Medicaid coverage at the 2015 mandate kink point ("baseline") and the counterfactual CDF induced by a $10 increase in the monthly mandate penalty, along with the 95% confidence interval.

Table G.1: Kinks in months of insurance coverage, all categories, by Medicaid expansion status

Coverage type	Any		Verified		Medicaid		Exchange		Off-Exchange	
Expanded Medicaid?	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	A. Coverage year 2015									
Kink	0.056	0.044	0.022	0.029	0.003	0.033	-0.022	0.008	0.018	-0.018
	(0.028)	(0.022)	(0.027)	(0.022)	(0.006)	(0.014)	(0.021)	(0.017)	(0.015)	(0.012)
P-value (difference)	0.736		0.858		0.051		0.270		0.063	
	B. Coverage year 2016									
Kink	0.015	0.025	-0.009	0.005	0.007	0.017	0.045	0.006	0.019	-0.001
	(0.025)	(0.018)	(0.026)	(0.020)	(0.006)	(0.010)	(0.019)	(0.015)	(0.015)	(0.012)
P-value (difference)	0.754		0.672		0.378		0.113		0.301	

Table reports the estimated kink obtained from a regression of the indicated coverage type on income, allowing for a kink and discontinuity at the mandate kink point, estimated separately by whether the state of residence has expanded Medicaid by the indicated coverage year. The sample consists of single tax returns in 2015 (Panel A) or 2016 (Panel B) with one exemption claimed and no signs of ESI offers, aged 27-64, with income between 200 and 250 percent of FPL (in 2015), or 300 and 390 percent of FPL (in 2016). Robust standard errors in parentheses. The reported p-value is the p-value of the hypothesis that the kinks are the same for the expansion and non-expansion states.

Table G.2: RKD estimates for CDF of Medicaid months, 2015

	Kink (1)	Standard error (2)	Penalty effect (3)
Pr(Medicaid months ≤ 0)	-0.222	(0.092)	-0.133
Pr(Medicaid months ≤ 1)	-0.218	(0.089)	-0.131
Pr(Medicaid months ≤ 2)	-0.222	(0.086)	-0.133
Pr(Medicaid months ≤ 3)	-0.231	(0.084)	-0.138
Pr(Medicaid months ≤ 4)	-0.229	(0.082)	-0.138
Pr(Medicaid months ≤ 5)	-0.219	(0.080)	-0.131
Pr(Medicaid months ≤ 6)	-0.198	(0.078)	-0.119
Pr(Medicaid months ≤ 7)	-0.183	(0.076)	-0.110
Pr(Medicaid months ≤ 8)	-0.143	(0.074)	-0.086
Pr(Medicaid months ≤ 9)	-0.145	(0.072)	-0.087
Pr(Medicaid months ≤ 10)	-0.151	(0.070)	-0.091
Pr(Medicaid months ≤ 11)	-0.143	(0.067)	-0.086

The sample consists of single tax returns with 2015 income between 200 and 250 percent of FPL, with one exemption claimed, no signs of ESI offers, aged 27-64. Each row is a separate regression; the outcome is an indicator for having at most the indicated number of months of Medicaid coverage on income (multiplied by 100). The independent variable is 2015 income (in thousands), allowing for a kink and discontinuity at the 2015 kink point. Column (1) shows the estimated kink, column (2) shows the standard error, and column (3) shows the implied effect of an extra dollar of penalty per month, which is $kink/20 * 12$.

H Imperfect expectations and ignorance of the penalty do not explain our low sensitivity

H.1 Rational but imperfect expectations

Our model of accounting for imperfect expectations allows for only two possibilities: perfect expectations for some, and non-rational expectations for others. It is possible that some people have imperfect but rational expectations, but these people would not contribute to the discontinuity when we use lagged income as the running variable, because they would not expect a first stage discontinuity. However, we do not find evidence that some people have imperfect but rational expectations. We test for such people by stratifying our sample according to measures of income uncertainty: the number of W2s and the presence of unemployment insurance income. The idea here is that income is much easier to predict for people with a single job (i.e. a single W2) and no job loss (i.e. no unemployment insurance income). We begin by defining "predictable" earners, people in households with at most one W2 per earner, and no unemployment insurance income. About three quarters of the sample meets this definition. If people make imperfect but rational forecasts of the penalty, then we would expect to find much larger penalty sensitivities among predictable earners, because for them realized income, which we observe, is likely close to expected income, the unobserved but theoretically desirable running variable.

We provide two pieces of evidence to suggest that our measure of predictability is valid. First, the change in income is more tightly concentrated around zero for predictable earners than for non-predictable earners, as can be seen in Appendix Figure H.2, which shows the distribution of income growth rates for predictable and non-predictable earners. Second, income is indeed much more predictable for this group than for non-predictable earners. For example, when we regress income in year t on a degree 7 polynomial in lagged income, interacted with year and filing status dummies, we obtain an R^2 of 60 percent for predictable earners but only 37 percent for non-predictable earners.

Of course, both these facts indicate that income and family structure alone do not perfectly predict next year's income, even among predictable earners. Nonetheless we think that looking at predictable earners is useful for two reasons. First, if expectations error in the running variable are a serious source of bias, we should still see larger discontinuities for predictable earners, for whom there is less measurement error. Second, predictable earners as we define them may experience idiosyncratic increases in income, for example coming from raises or promotions. If these idiosyncratic income changes are predictable to the people in our data (given their private information), then realized income is still the appropriate running variable, even if we cannot predict these income changes given our limited information.

To test the importance of income expectations for our low semi-elasticity, we re-estimate our main RD models, separately for predictable and non-predictable earners. The results are in Table H.1, and we show the RD plots in Figure H.1. We see large and statistically significant differences between predictable and non-predictable earners but, surprisingly, it is the predictable earners who are less responsive to the penalty. Their semi elasticity is 50 to 100 percent lower than non-predictable earner's, depending on the specification. This low responsiveness does not arise from high baseline coverage rates; predictable earners' coverage rates are only 3 percent higher. Nor is it an artifact of self-reported coverage or access to ESI; we find similar differences across coverage categories.

These results suggest that income unpredictabiltiy does not necessarily explain our low semi-elasticity. This conclusion, however, is potentially sensitive to our (admittedly arbitrary) classification of "predictable" and "non-predictable." As an alternative approach, we divide our data up into cells defined by unemployment insurance receipt (yes/no), number of W2s (zero, one, or two or more), filing status (single, married filing jointly, head of household, other), and year. Note that our "predictable" measure is a coarsening of these cells. For each cell, we estimate the predictability of income (wit the R^2 from a regression of income on a degree 7 polynomial in its lag), and we estimate our main RD model for coverage (continuing to use a 5-FPL point bandwidth). In Appendix Figure H.3, we plot the estimated semi-elasticities against income predictability for each cell. The figure shows no apparent relationship between the semi-elasticity and income predictability, although a weak relationship would be hard to see because some of the estimates are noisy. To more precisely measure the relationship between penalty responsiveness and income predictability, we regress cell-level semi-elasticity against cell-level R^2. We obtain a constant of 0.19 and a slope of -0.59; that is, semi-elasticities and predictability are negatively correlated.[33]

These results show that people with more predictable income do not exhibit a substantially greater coverage discontinuity. Of course, an important caveat is that there may be unobserved factors correlated with income predictability that lead to low responsiveness; such factors would muddy the interpretation of these results. Nonetheless, this finding is consistent with the view that imperfect but rational forecast errors are not a major factor in explaining our low semi-elasticity estimates.

H.2 Imperfect knowledge of the penalty

Another possible explanation for our relatively low semi-elasticity is that people are simply unaware of the mandate's existence. If this ignorance is permanent, then our estimates would still reflect the long-run effect of the mandate. If, however, people have learned about the mandate in recent years, then our estimates might understate the long-run effect of penalty. To provide some evidence on the role of knowledge of the penalty, we consider whether people who paid the penalty in the past—and were therefore likely aware of the penalty—respond more to penalty in a given year. Of course, people who paid the penalty in the past could not have been insured for the whole year. So to make the comparison clean, we look only at people who were not insured for the whole year, comparing responsiveness to the mandate among those who did and did not pay the penalty in the past. Specifically, we merge in to the 2015 and 2016 data information on penalty paying in 2014, along with an indicator for self-reported whole-year coverage, i.e. for checking the coverage box on Form 1040.

[33]We weight by the inverse of the standard error, to down weight the cells where we have less statistical precision.

Table H.2 shows the RD estimates for 2015 and 2016 coverage for three groups: people who "checked the box" reporting full year coverage, people who did not check the box but did not pay a penalty, and people who did not check the box and paid the penalty. We present the RD plots in Figure H.4. The results show that, among the uninsured in 2014, people who did and did not pay the penalty had fairly similar coverage rates in 2015 and 2016, but people who paid the penalty were if anything less responsive than people who did not pay the penalty; they have lower discontinuities and, especially, lower semi-elasticities.

Under the assumption that people who previously paid the penalty are fully informed of the penalty, these results do not support the view that widespread ignorance of the penalty explains our low semi-elasticities. That is, if people who paid the penalty are fully aware of it, then people who didn't pay the penalty should also be fully aware (because they are more responsive to the penalty) so ignorance of the penalty cannot be widespread. Of course, an alternative interpretation of these results is that people who paid the penalty are not particularly aware of it. We cannot rule out this possibility—it could be that people simply answer questions in their tax software, and pay the final amount. Nonetheless, the evidence here is consistent with the view that people are reasonably aware of the penalty.

Figure H.1: Months of coverage by type, stratifying on income predictability

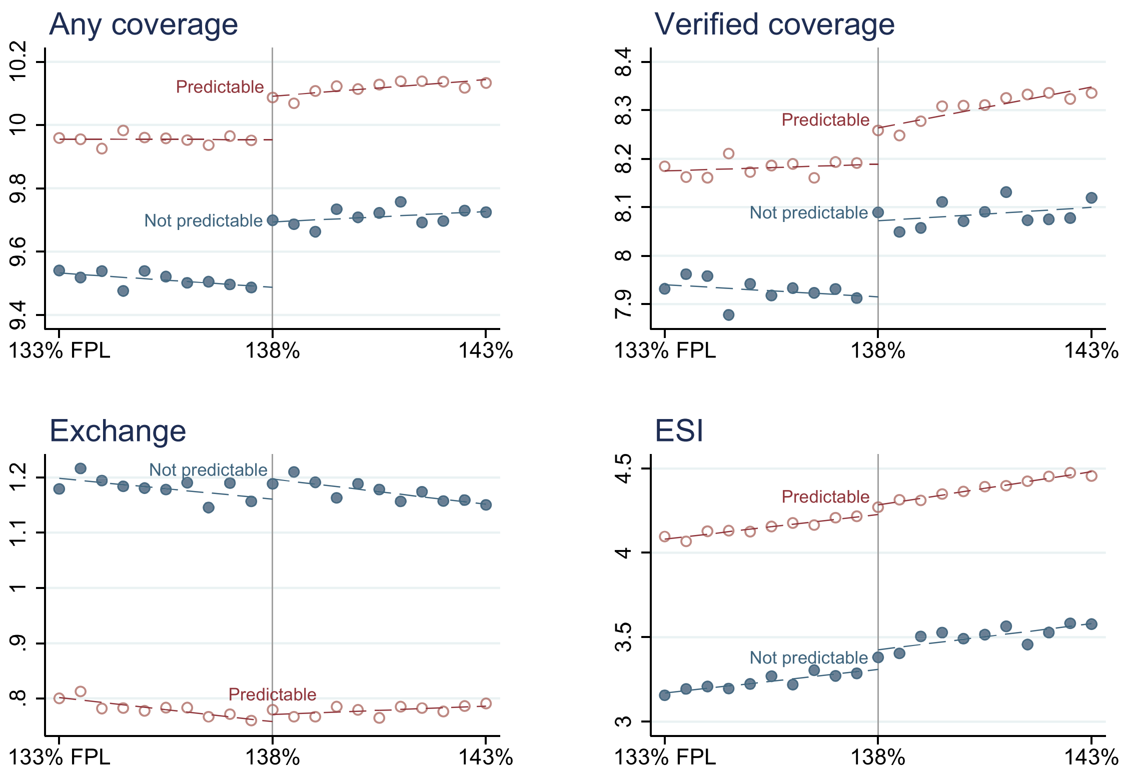

Notes: Figure shows the average number of months of insurance of the indicated type, in each 0.5 percent of FPL bin, separately for predictable earners and not predictable earners. Predictable earners are defined as people living in households with no unemployment insurance income and at most one W2 per earner. Sample consists of people aged 0-64, pooling 2015-2016.

Figure H.2: Change in log income, predictable and non-predictable earners

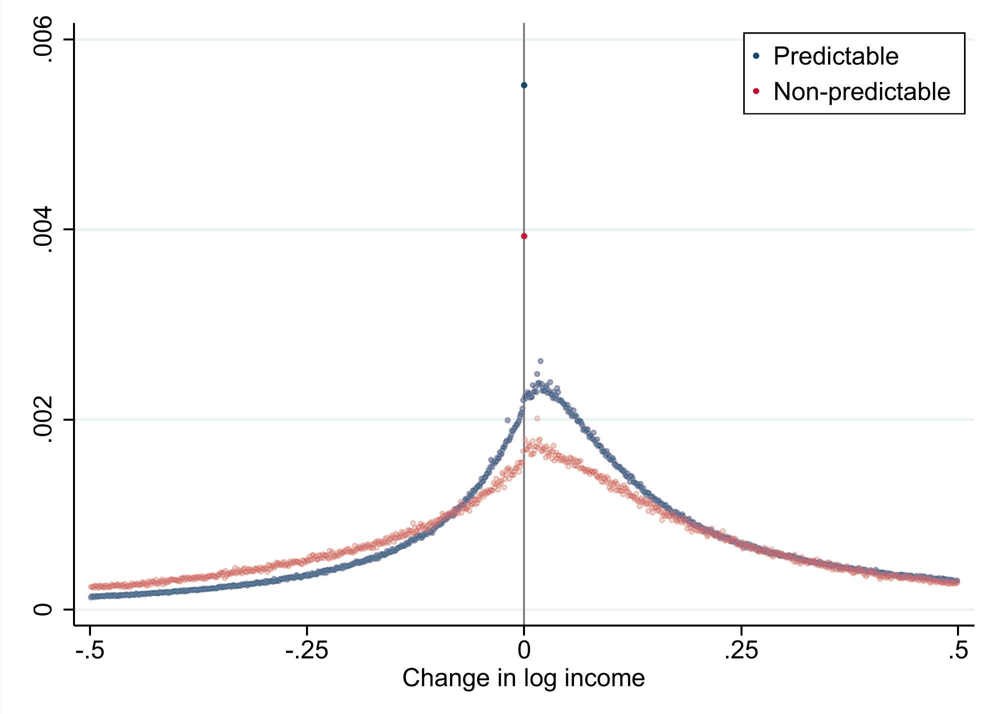

Notes: Figure shows the distribution of the change in log income for predictable and non-predictable earners. Each bin is 0.001 points wide. Not shown are changes greater than 50 percent in absolute value. About 18 percent of the non-predictable are in this category, and 17 percent of predictable earners. The sample consists of people aged 0-64 in 2015 or 2016, with income between 110 and 160 percent of FPL, living in non-expansion states. Predictable earners are defined as having at most one W2 per earner and no unemployment insurance income.

Figure H.3: Income predictability and responsiveness to the mandate penalty

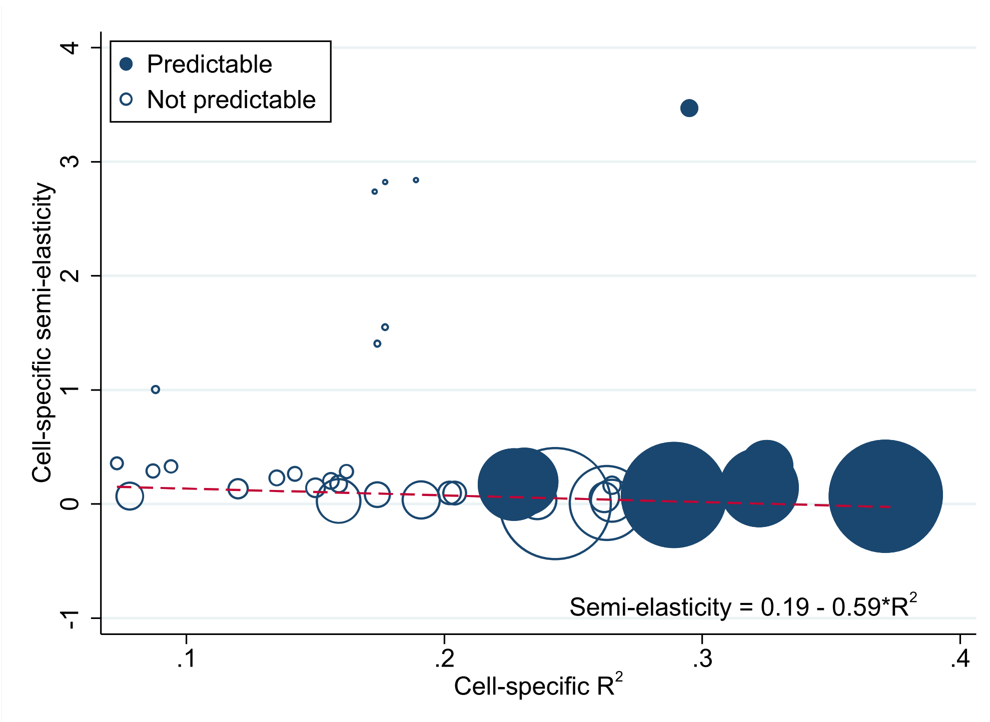

Notes: Each point is a cell defined by uninsurance receipt (yes/no), number of w2s (zero, one, or two or more), filing status (single, married filing jointly, head of household, other), and year. For each cell, we estimate semi-elasticities using our RD model and a 5-FPL point bandwidth, and we plot the estimated semi-elasticity on the y-axis. On the x-axis, we plot income predictability, measured as the R^2 from a regression of income on a degree seven polynomial in its lag. The size of each point is proportional to the inverse of the standard error of the semi-elasticity estimate. We also report the estimate from a weighted least squares regression, where the weights are the inverse of the standard error.

Figure H.4: Penalty paid and months of coverage, stratifying on past experience with the penalty

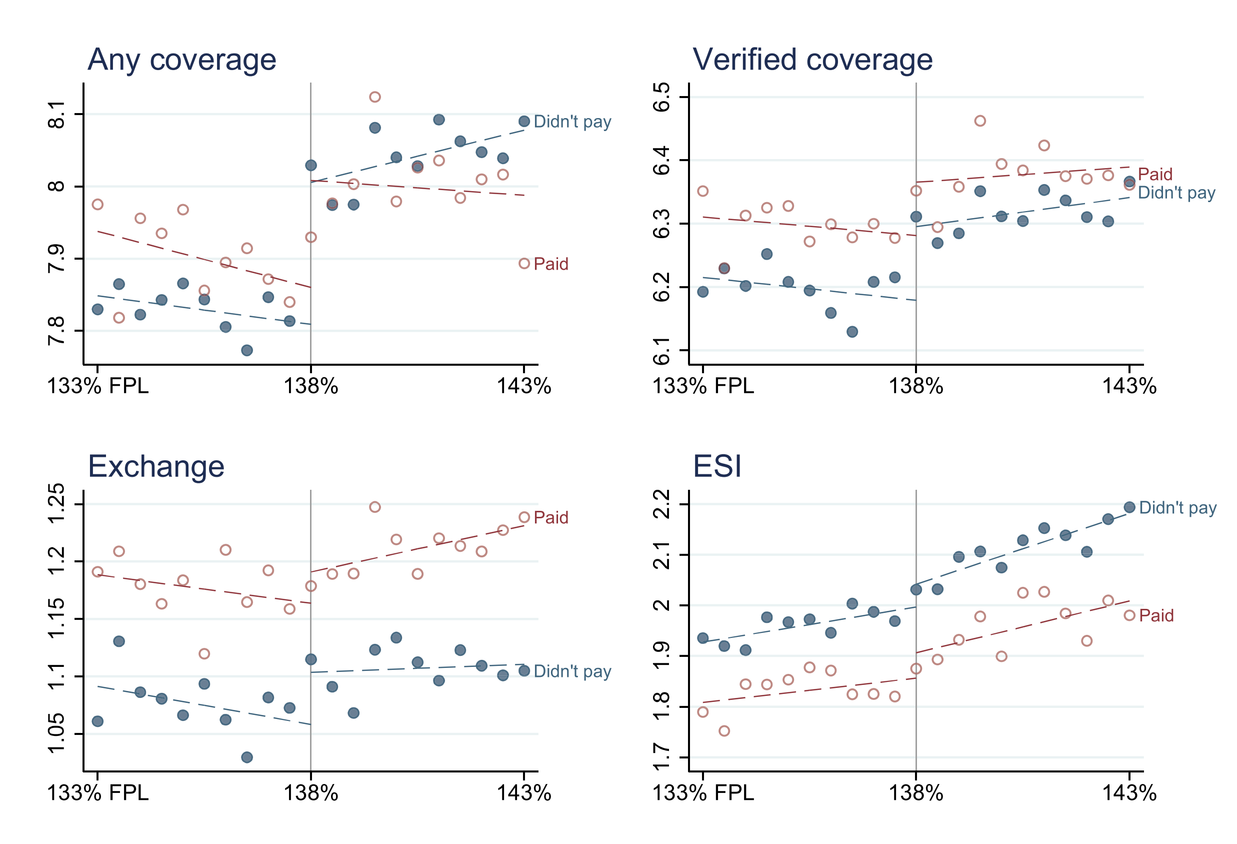

Notes: The figures plot the indicated outcome against income, separately for people who did and did not pay the mandate penalty in 2014. Sample is limited to people who indicated in 2014 that they did not have full-year coverage. The sample pools 2015 and 2016, and consists of people aged 0-64.

Table H.1: Coverage discontinuities by income predictability

Group	# in group (1)	Mean months (2)	Discontinuity (3)	Standard error (4)	Semi-elasticity (5)
			A. Any coverage		
Not stable	1,352,328	9.61	0.201	(0.015)	0.303
Stable	3,584,356	10.03	0.132	(0.009)	0.189
p-value of equality			0.000		
			B. Verified coverage		
Not stable	1,352,328	8.00	0.157	(0.018)	0.284
Stable	3,584,356	8.24	0.069	(0.011)	0.121
p-value of equality			0.000		
			C. Exchange		
Not stable	1,352,328	1.18	0.037	(0.012)	0.453
Stable	3,584,356	0.78	0.011	(0.006)	0.201
p-value of equality			0.044		
			D. ESI		
Not stable	1,352,328	3.36	0.112	(0.017)	0.482
Stable	3,584,356	4.26	0.047	(0.012)	0.158
p-value of equality			0.002		

Notes: Table reports the estimated discontinuity in coverage, as well as the semi-elasticity, estimated separately for predictable earners and not predictable earners, and the p-value for the test that discontinuities in each group are equal. Sample consists of people aged 0-64 living in non-expansion states with income between 133 and 143 percent of FPL. To maximize power, we pool 2015 and 2016 and include a year dummy variable.

Table H.2: Coverage discontinuities, stratifying on past experience with the mandate

Group	# in group (1)	Mean months (2)	Discontinuity (3)	Standard error (4)	Semi-elasticity (5)
			A. Any coverage		
Checked box	2,874,287	11.20	0.092	(0.007)	0.133
No box, no penalty	1,058,551	7.93	0.191	(0.021)	0.353
No box, penalty	605,539	7.95	0.131	(0.028)	0.199
p-value (box vs. no box, no pen)			0.000		
p-value (box vs. no box, pen)			0.165		
p-value (no box, no pen vs. no box, pen)			0.082		
			B. Verified coverage		
Checked box	2,874,287	9.60	0.053	(0.010)	0.089
No box, no penalty	1,058,551	6.25	0.117	(0.022)	0.273
No box, penalty	605,539	6.34	0.083	(0.028)	0.157
p-value (box vs. no box, no pen)			0.008		
p-value (box vs. no box, pen)			0.321		
p-value (no box, no pen vs. no box, pen)			0.334		
			C. Exchange		
Checked box	2,874,287	0.83	0.009	(0.007)	0.181
No box, no penalty	1,058,551	1.09	0.043	(0.013)	0.575
No box, penalty	605,539	1.19	0.026	(0.017)	0.267
p-value (box vs. no box, no pen)			0.019		
p-value (box vs. no box, pen)			0.350		
p-value (no box, no pen vs. no box, pen)			0.437		
			D. ESI		
Checked box	2,874,287	5.46	0.045	(0.013)	0.132
No box, no penalty	1,058,551	2.03	0.045	(0.016)	0.326
No box, penalty	605,539	1.89	0.046	(0.021)	0.296
p-value (box vs. no box, no pen)			0.979		
p-value (box vs. no box, pen)			0.941		
p-value (no box, no pen vs. no box, pen)			0.961		

Notes: Table reports the estimated discontinuity in coverage in 2015 and 2016, as well as the semi-elasticity, estimated separately for three groups: people who checked the box in 2014 to indicate full-year coverage; people who did not check the box and did not pay a mandate penalty; and people who did not check the box and paid a penalty. Table also reports the p-value for the test that discontinuities in each group are equal Sample consists of people aged 0-64 living in non-expansion states with income between 133 and 143 percent of FPL. To maximize power, we pool 2015 and 2016 and include a year dummy variable.