

Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts[†]

By AMY FINKELSTEIN, NATHANIEL HENDREN, AND MARK SHEPARD*

How much are low-income individuals willing to pay for health insurance, and what are the implications for insurance markets? Using administrative data from Massachusetts' subsidized insurance exchange, we exploit discontinuities in the subsidy schedule to estimate willingness to pay and costs of insurance among low-income adults. As subsidies decline, insurance take-up falls rapidly, dropping about 25 percent for each \$40 increase in monthly enrollee premiums. Marginal enrollees tend to be lower-cost, indicating adverse selection into insurance. But across the entire distribution we can observe (approximately the bottom 70 percent of the willingness to pay distribution) enrollees' willingness to pay is always less than half of their own expected costs that they impose on the insurer. As a result, we estimate that take-up will be highly incomplete even with generous subsidies. If enrollee premiums were 25 percent of insurers' average costs, at most half of potential enrollees would buy insurance; even premiums subsidized to 10 percent of average costs would still leave at least 20 percent uninsured. We briefly consider potential explanations for these findings and their normative implications. (JEL G22, H51, H75, I13, I18)

Governments spend an enormous amount of money on health insurance for low-income individuals. For instance, the US Medicaid program (at \$550 billion in 2015) dwarfs the size of the next largest means-tested programs, food stamps

* Finkelstein: Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, Building E52, Cambridge, MA 02142, and NBER (email: afink@mit.edu); Hendren: Department of Economics, Harvard, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138, and NBER (email: nhendren@fas.harvard.edu); Shepard: Harvard Kennedy School, 79 JFK Street, Mailbox 114, Cambridge, MA 02138, and NBER (email: Mark_Shepard@hks.harvard.edu). Thomas Lemieux was the coeditor for this article. We thank Melanie Rucinski for excellent research assistance and Lizi Chen, Ray Kluender, and Martina Uccioli for helping with several calculations. We thank the Massachusetts Health Connector (and particularly Marissa Woltmann and Michael Norton) for assistance in providing and interpreting the CommCare data. We thank our discussants Jeff Clemens and Marika Cabral for thoughtful and constructive comments. We also thank Abhijit Banerjee, Amitabh Chandra, Keith Ericson, Josh Goodman, Jon Gruber, Jon Kolstad, Tim Layton, Jeff Liebman, Brigitte Madrian, Neale Mahoney, Rebecca Sachs, and seminar participants at Harvard Kennedy School, the AEA Annual Meetings, University of Texas, MIT, and NBER Summer Institute and Public Economics meeting for helpful comments and discussions. We gratefully acknowledge data funding from Harvard's Lab for Economic Applications and Policy. Shepard acknowledges funding from National Institute on Aging grant T32-AG000186 (via the National Bureau of Economic Research). Hendren acknowledges funding from the National Science Foundation.

[†]Go to <https://doi.org/10.1257/aer.20171455> to visit the article page for additional materials and author disclosure statements.

and the Earned Income Tax Credit (\$70 billion each).¹ Perhaps because of these high and rising costs, public programs increasingly offer *partial* subsidies for health insurance, requiring enrollees to pay premiums to help cover costs. Partial subsidies are a key feature of market-based programs such as Medicare Part D and the Affordable Care Act (ACA) exchanges, and even traditional low-income programs like Medicaid and the Children's Health Insurance Program (CHIP) increasingly require premiums for some enrollees (Smith et al. 2015, Brooks et al. 2017). Partial subsidies are also a textbook policy response to adverse selection if a full coverage mandate may not be efficient (Einav and Finkelstein 2011). In such settings, measuring willingness to pay and costs is important for analyzing the impact and desirability of alternative subsidies.

In this paper, we estimate low-income individuals' willingness to pay (WTP) for health insurance, assess how it compares to the cost they impose on the insurer, and discuss the positive and normative implications for subsidized health insurance programs. We do so in the context of Massachusetts' pioneer health insurance exchange for low-income individuals, known as Commonwealth Care or CommCare. Established in the state's 2006 health care reform, CommCare offered heavily-subsidized private plans to non-elderly adults below 300 percent of poverty who did not have access to insurance through an employer or another public program. Public subsidies were substantial: on average for our study population, enrollee premiums are only about \$70 per month, or less than one-fifth of insurer-paid medical claims (\$359 per month) or insurer prices (\$422 per month). There was also a health insurance mandate backed by financial penalties.

We use a regression discontinuity design, together with administrative data on enrollment and medical claims, to estimate demand and cost for CommCare plans. Our main analysis focuses on fiscal year 2011, when the insurance options had a convenient vertical structure. We also present some complementary results for the full 2009–2013 period over which we have data.

The analysis leverages discrete changes in subsidies at several income thresholds. Subsidies were designed to make enrollee premiums for the cheapest insurer's plan "affordable." In practice, the subsidy amount changes discretely at 150 percent, 200 percent, and 250 percent of the federal poverty line (FPL). These discontinuities in program rules provide identifying variation in enrollee premiums. The cheapest plan's (post-subsidy) monthly enrollee premium increases by about \$40 at each of the discontinuities, and more generous plans experience a \$40 to \$50 increase in (post-subsidy) monthly enrollee premiums.

We first document two main descriptive patterns. First, enrollee demand is highly sensitive to premiums. With each discrete increase in enrollee premiums, enrollment in CommCare falls by about 25 percent, or a 20–24 percentage point fall in the take-up rate. Second, we find that despite the presence of a coverage mandate, the market is characterized by adverse selection. As enrollee premiums rise, lower-cost enrollees disproportionately drop out, raising the average cost of the remaining insured population. We estimate that average medical claims rise by \$10–\$50 per month (or 3–14 percent) with each premium increase.

¹ See US Department of Health and Human Services (2015), US Department of Agriculture (2016), and US Internal Revenue Service (2015).

We use a simple model to analyze the implications of these descriptive patterns. The nature of the individual choice problem lends itself naturally to a vertical model of demand in which individuals choose among a “high-coverage” (H) option, a “low-coverage” (L) option, and a third option of uninsurance (U); the vast majority of enrollees who buy insurance choose the high-coverage option, H . We use the model, which is a slight extension of Einav, Finkelstein, and Cullen (2010), along with the premium variation to map out curves for willingness to pay, average costs of insurance, and costs of marginal enrollees.

The model allows us to translate the descriptive patterns into two main results. First, even large insurance subsidies are insufficient to generate near-complete take-up of insurance by low-income adults. For example, at the median of the willingness to pay distribution, willingness to pay for H is about \$100 per month, less than one-quarter of average costs of \$420 per month if all those with above-median willingness to pay enrolled in H . Even with a subsidy that makes enrollee premiums for the H plan equal to 25 percent of insurers’ average costs, at most one-half of the population would purchase insurance if offered H . Subsidies making enrollee premiums 10 percent of insurers’ average costs still leave at least 20 percent uninsured.

These findings suggest that even modest enrollee premiums can be a major deterrent to universal coverage among low-income people. This deterrent is likely to be even larger in the ACA exchanges, in which income-specific premiums are significantly higher than in CommCare: 2–10 percent of income for the benchmark plan in the ACA versus 0–5 percent of income in CommCare.² Our results may thus help explain coverage outcomes in the ACA exchanges, where early evidence suggests highly incomplete take-up (Tebaldi 2017, Kaiser Family Foundation 2016). The price responsiveness that we document is also useful for generating predictions of coverage rates under alternate reform proposals or subsidies.

Second, although adverse selection exists, it is not the primary driver of low take-up. The cost of marginal consumers who enroll when premiums decline is less than the average costs of those already enrolled, indicating that plans are adversely selected (Einav, Finkelstein, and Cullen 2010). But across the entire in-sample distribution, which spans the sixth to the seventieth percentile of the willingness to pay distribution, the willingness to pay of marginal enrollees still lies far below their own expected costs imposed on insurers for either the H or L plans. For example, for the median willingness to pay individual, the gap between the costs of the marginal enrollee and average costs of enrollees explains only one-third of the \$300 gap between willingness to pay and average costs. In other words, most individuals would not enroll even if prices were subsidized to reflect marginal enrollees’ own expected insurer costs.

This finding contrasts with textbook models of insurance markets in which demand is assumed to exceed own cost, and adverse selection is widely viewed as the major barrier to insurance coverage. In our setting, enrollment is low not simply because of adverse selection, but because people are not willing to pay their own cost imposed on the insurer.

²These are based on authors’ calculations using ACA and CommCare policy parameters. The ACA premiums are for the second-cheapest silver plan; the CommCare premiums are for the cheapest plan.

In the final section of the paper, we briefly explore potential explanations for our findings and analyze their normative implications. One explanation is that because of uncompensated care, the costs individuals impose on the insurer differs from the costs they would pay if they were uninsured. Back-of-the-envelope calculations using other estimates of the prevalence of uncompensated care suggest this could explain the low WTP. Additionally, a range of behavioral explanations, such as optimistic beliefs that underestimate expected costs, could explain low WTP, and also have important normative implications. We briefly discuss potential normative rationales for subsidies based on behavioral biases, as an offset to the externalities resulting from uncompensated care (i.e., the Samaritan's dilemma, Buchanan 1975, Coate 1995), or as a means of redistribution to low-income households.

Related Literature.—While a substantial literature estimates demand and costs for health insurance, there is relatively little work providing such estimates for low-income adults on whom much public policy attention is focused.³ This is likely because, until recently, most of the low-income uninsured either were not offered health insurance or faced prices that were difficult to measure. This precluded standard demand and welfare analysis based on choices, as has been widely used in the study of private (often employer-provided) health insurance markets (see Einav, Finkelstein, and Levin 2010 for an overview). One effort to surmount this obstacle is Krueger and Kuziemko (2013), who conducted a survey experiment designed to elicit willingness to pay for hypothetical plan offerings among a broad sample of the uninsured from the full spectrum of the income distribution. In another attempt to circumvent the lack of direct estimates of willingness to pay, Finkelstein, Hendren, and Luttmer (2015) assume a normative utility function over estimates of the reduced-form impact of Medicaid in order to infer willingness to pay for Medicaid by a low-income population.

The 2010 passage of the ACA has given researchers an opportunity to directly study how low-income insurance demand responds to subsidies (e.g., Tebaldi 2017; Frean, Gruber, and Sommers 2017), although the ACA's subsidy schedule lacks the sharp discontinuities present in Massachusetts, which we exploit for our research design.⁴ Nonetheless, our estimates of insurance demand in the low-income adult population in Massachusetts are roughly similar to Tebaldi's (2017) estimates for a largely low-income population in the California ACA exchange.⁵ Such findings are also consistent with substantially incomplete take-up of subsidized insurance under the ACA (e.g., Kaiser Family Foundation 2016).

³There is more work on demand for employer-sponsored insurance, although this literature does not typically go so far as to estimate a demand curve. However, our results are qualitatively consistent with incomplete take-up of employer coverage, despite the large subsidies of employee premiums (Cooper and Schone 1997, Farber and Levy 2000).

⁴In other related papers, Dague (2014) examines how enrollment duration in Wisconsin Medicaid responds to increases in monthly premiums, and Chan and Gruber (2010) study the intensive margin of low-income individuals' price sensitivity in their choice among health plans in Massachusetts. Ericson and Starc (2015) estimate demand in the high-income (>300 percent of poverty) Massachusetts exchange using age discontinuities in insurer prices, but their estimates are for demand among plans *conditional* on buying insurance. None of these studies estimate willingness to pay for insurance.

⁵Tebaldi (2017) estimates a 2–4 percent decline in enrollment for an across-the-board \$100 annual premium increase for subsidized enrollees without children. Proportionally scaling down our central estimate (25 percent decline for \$39/month = \$468/year) would imply a decline of 5.3 percent for a \$100/year premium increase.

Our finding that low-income enrollees in Massachusetts value formal health insurance products at substantially below their average cost is consistent with other estimates for other low-income populations (e.g., Finkelstein, Hendren, and Luttmer 2015; Tebaldi 2017) but contrasts with findings for higher-income populations. In particular, Hackmann, Kolstad, and Kowalski (2015) study the *unsubsidized* Massachusetts health insurance exchange for individuals *above* 300 percent of poverty. They also find evidence of adverse selection but estimate that willingness to pay *exceeds* own costs over the entire population of potential consumers, in contrast to our estimates for a low-income population. One natural explanation for these divergent findings is that low-income individuals likely have much greater access to uncompensated care. Indeed, a growing empirical literature documents the large role of uncompensated care for the (predominantly low-income) uninsured and the impact of insurance in decreasing unpaid bills (see, e.g., Garthwaite, Gross, and Notowidigdo 2015; Finkelstein et al. 2012; Mahoney 2015; Dobkin et al. 2016; Hu et al. 2016).⁶ Another potential explanation is differential behavioral biases among lower and higher income individuals (e.g., Mullainathan and Shafir 2014).

Finally, our results have implications for the broader literature on adverse selection in health insurance markets. The empirical literature has extensively documented the presence of adverse selection in health insurance markets but concluded that the welfare cost of the resultant mispricing of contracts is relatively small. This literature however has “looked under the lamppost,” primarily focusing on selection across contracts that vary in limited ways, rather than selection that causes a market to unravel, leaving open the possibility of larger welfare costs on this margin (Einav, Finkelstein, and Levin 2010). Our work, however, finds evidence of significant adverse selection on the extensive margin of purchasing insurance versus remaining uninsured, a finding consistent with past work on the Massachusetts reform (Chandra, Gruber, and McKnight 2011; Hackmann, Kolstad, and Kowalski 2015; Jaffe and Shepard 2017). But it also finds that adverse selection is not the primary driver of limited demand for formal insurance among low-income adults.

The rest of the paper proceeds as follows. Section I presents the setting and data. Section II presents the basic descriptive empirical evidence, documenting the level and responsiveness to price of both insurance demand and average insurer costs. Section III uses a simple model of insurance demand to translate the empirical results from Section II into estimates of willingness to pay and costs for insurance. Section IV briefly considers potential explanations for low willingness to pay and normative implications. The final section concludes.

⁶Differences in the availability of uncompensated care may also reconcile our findings with results from a calibrated life-cycle model suggesting that the low-income elderly's willingness to pay for Medicaid is above their costs (De Nardi, French, and Jones 2016). Unlike low-income adults, low-income elderly do not have access to substantial uncompensated nursing home care (the primary health care covered by Medicaid), either in the De Nardi, French, and Jones (2016) model or in practice.

I. Setting and Data

A. Setting: Massachusetts Subsidized Health Insurance Exchange

CommCare.—We study Commonwealth Care (CommCare), a subsidized insurance exchange created under Massachusetts' 2006 "Romneycare" health insurance reform that laid the foundation for many of the health insurance exchanges created in other states under the Affordable Care Act (ACA). CommCare operated from 2006–2013 before shifting form in 2014 to comply with the ACA. We focus on the market in fiscal year 2011 (July 2010 to June 2011) but also present descriptive results for fiscal years 2009–2013, the full period over which we have data. The market rules described below apply to 2011; the rules for other years are similar except in some details.

CommCare covered low-income adults with family income below 300 percent of the federal poverty level (FPL) and without access to insurance from another source, including an employer or another public program (i.e., Medicare or Medicaid). This population is similar to those eligible for subsidies on the ACA exchanges. Given Medicaid eligibility rules in Massachusetts,⁷ the CommCare-eligible population consisted of adults aged 19–64 without access to employer coverage and who were either (i) childless and below 300 percent of FPL, (ii) non-pregnant parents between 133 percent and 300 percent of FPL, or (iii) pregnant women between 200 percent and 300 percent of FPL.

CommCare specified a detailed benefit structure (i.e., covered services and a schedule of cost sharing rules) and then solicited competing private insurers to provide these benefits. Each insurer offered a single plan that had the standardized set of benefits but could differ in its network of hospitals and doctors. Between four and five insurers participated in the market each year. Benefit design and participating insurers were very similar to the Massachusetts Medicaid program. In particular, CommCare enrollees faced very modest copays.⁸

Eligible individuals could enroll during the annual open enrollment period at the start of the fiscal year, or at any time if they experienced a qualifying event (e.g., job loss or income change). To enroll, individuals filled out an application form with information on age, income over the last 12 months, family size, and access to other health insurance. The state used this form to determine whether an applicant was eligible for Medicaid, CommCare, or neither. The form was also used to calculate income as a share of FPL for determining an enrollee's premiums. However, as discussed below, the translation from income and other information on the form into FPL units was not readily transparent to applicants on the form.

⁷During our study period, Medicaid covered all relevant children (up to 300 percent FPL) and disabled adults, as well as parents up to 133 percent FPL and pregnant women up to 200 percent FPL. Medicaid also covered long-term unemployed individuals earning up to 100 percent FPL and HIV-positive individuals up to 200 percent FPL, both relatively small groups.

⁸Enrollees below 100 percent of FPL received benefits equivalent to Medicaid. Enrollees between 100–200 percent FPL received a plan that we estimate (based on claims data) has a 97 percent actuarial value, while those between 200–300 percent FPL received a 95 percent actuarial value plan. The slight change in generosity at 200 percent FPL is a potential threat to the regression discontinuity (RD) analysis of demand and costs at 200 percent FPL; we show below that our main results are not sensitive to excluding this discontinuity.

Enroll Now! Select and Enroll in a Commonwealth Care health plan

Below are the Commonwealth Care health plans you can choose from. The dollar amount next to each health plan is what you must pay each month to stay enrolled in that plan. If you select a health plan with \$0.00 next to it, you will not be charged a monthly premium. The premiums listed below are based on your plan type, which depends on your income and your family size. Based on the information you provided, you are eligible for **Plan Type X**.

1. Choose your health plan and premium. Choose only one.

These plans are available to you. Read each Health Plan Information description to learn about the Commonwealth Care health plans.

<BMC HealthNet Plan	\$0.00	web address	Phone number>
<CeltiCare Health Plan	\$0.00	web address	Phone number>
<Fallon Community Health Plan	\$0.00	web address	Phone number>
<Neighborhood Health Plan	\$0.00	web address	Phone number>
<Network Health	\$0.00	web address	Phone number>

FIGURE 1. SNAPSHOT OF COMM-CARE PLAN CHOICE FORM

Notes: The figure shows a snapshot of the key section of the plan choice form sent to accepted CommCare applicants. As noted in the text, enrolling in CommCare involves two steps: (i) an application form, which collects information on income, family size, and access to other insurance, which lets the state determine eligibility, and (ii) a plan choice form, which applicants must return to choose a plan. More extensive snapshots of these forms are included in online Appendix A.

If approved for CommCare, individuals were notified (by mail and/or email) and provided information on the premiums for CommCare plans. They then had to complete a second form (or contact CommCare by phone/online) to choose a plan and pay the first month's premium. Individuals who did not make a plan choice and the associated payment did not receive coverage. Coverage commenced at the start of the month following receipt of payment. Once enrolled, individuals stayed enrolled as long as they remained eligible and continued paying premiums. Income and eligibility status changes were supposed to be self-reported and were also verified through an annual "redetermination" process that included comparisons to tax data and lists of people enrolled in employer insurance.

Figure 1 shows a snapshot of the key section of the plan choice form displaying an enrollee's plan options and premiums. Online Appendix A shows the entire plan choice form and snapshots of the initial application form. We take away two conclusions from these documents. First, enrolling in subsidized insurance may involve nontrivial hassles. Our willingness to pay measure will implicitly incorporate these hassle costs (see Section IIIA). Second, the plan choice form displays enrollee premiums prominently, while referring enrollees online for information about provider networks; employee premium information thus appears to be quite salient, which may help explain our findings that potential enrollee demand responds strongly to premiums.

Subsidy Structure.—Insurers in CommCare set a base plan price that applied to all individuals, regardless of income (or age, region, or other factors). The actual payment the insurer received from CommCare equaled their base price times a risk score intended to capture predictable differences in health status.

Panel A. Premiums for cheapest plan, 2009–2013 Panel B. Prices, subsidies, and premiums in 2011

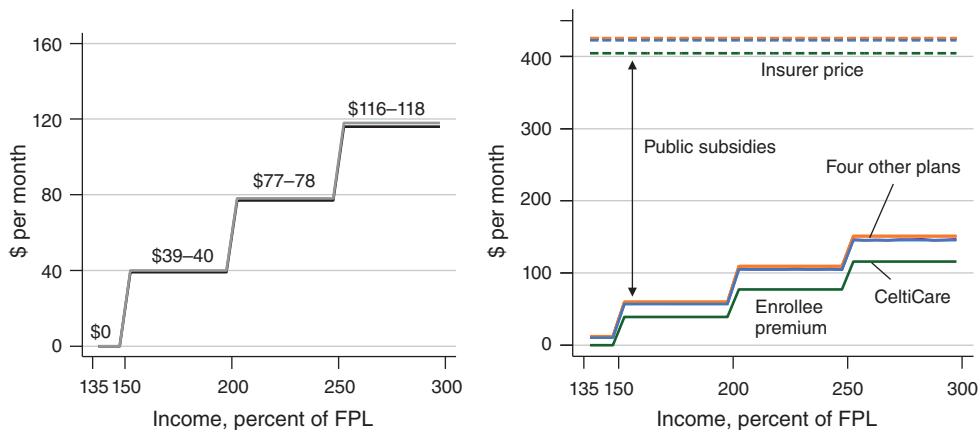


FIGURE 2. INSURER PRICES AND ENROLLEE PREMIUMS IN COMM-CARE MARKET

Notes: Panel A plots enrollee premiums for the cheapest plan by income as a percent of FPL, noting the thresholds (150 percent, 200 percent, and 250 percent of FPL) where the amount increases discretely. The black lines show the values that applied in 2009–2012; the gray lines show the (slightly higher) values for 2013. Panel B shows insurer prices (dotted lines) and enrollee premiums (solid lines) for the five plans in 2011. In this year, four insurers set prices within \$3 of a \$426/month price cap, while CeltiCare set a lower price (\$405) and therefore had lower enrollee premiums.

Enrollees paid premiums equal to their insurer's base price minus an income-varying subsidy paid by the state.⁹ Subsidies were set so that enrollee premiums for the lowest-price plan equaled a target "affordable amount." This target amount was set separately for several bins of income, with discrete changes at 150 percent, 200 percent, and 250 percent of FPL. Panel A of Figure 2 shows the result: enrollee premiums for the cheapest plan vary discretely at these thresholds. For the years 2009–2012 (shown in black), the cheapest plan is free for individuals below 150 percent of FPL and increases to \$39 per month above 150 percent FPL, \$77 per month above 200 percent FPL, and \$116 per month above 250 percent of FPL. In 2013 (shown in gray), these amounts increase slightly to \$0/\$40/\$78/\$118. Consistent with the goal of affordability, these premiums were a small share of income. For instance, for a single individual in 2011 (whose FPL equaled \$908 per month), these premiums ranged from 0–5 percent of income (specifically, 2.9 percent of income just above 150 percent FPL, 4.2 percent just above 200 percent FPL, and 5.1 percent just above 250 percent FPL).

2011 Plan Options.—We analyze the market in 2009–2013 but focus especially on fiscal year 2011 when the market had a useful vertical structure with plans falling into two groups. In 2011 CommCare imposed a binding cap on insurer prices of \$426 per month. Four insurers (BMC HealthNet, Fallon, Neighborhood Health Plan, and Network Health) all set prices within \$3 of this cap. The exception was CeltiCare, which set a price of \$405 per month. Panel B of Figure 2 shows these

⁹We will use "price" to refer to the pre-subsidy price set by insurers and "premium" to refer to the post-subsidy amount owed by enrollees.

insurer prices and the resulting post-subsidy enrollee premiums by income. The prices and premiums of the four high-price plans are nearly indistinguishable, while CultiCare's premium is noticeably lower.

Along with its lower price, CultiCare also had a more limited network than other plans. We estimate that CultiCare covered 42 percent of Massachusetts hospitals (weighted by bed size), compared to 79 percent or higher for the other three plans offered statewide.¹⁰ Both because of this limited network and because of its lack of long-term reputation with consumers (it had entered the state insurance market only in 2010), CultiCare was perceived by enrollees as less desirable, aside from its lower price.¹¹

As a result, in much of our analysis that follows we pool the 2011 plans into two groups: CultiCare as a low coverage (L) option and the other four plans with extremely similar prices pooled together as a high coverage (H) option. We interpret H as a composite contract that gives enrollees a choice among the four component insurers, with its utility equal to the max over these four insurers. When we specify and estimate a model of insurance demand in Section III, we will further assume that H is perceived as higher quality than L . We also show in an extension in Section IIID that we can generate fairly tight bounds on willingness to pay in a more general model that does not assume this vertical structure.

Figure 3 zooms in on enrollee premiums for the H plan and the L plan in 2011 by income. We define the enrollee premium for H as the share-weighted average of the component plans; online Appendix Table 5 reports these separately for each component plan. As previously discussed, enrollee premiums for the cheapest plan L (p_L), subsidized to equal a target affordable amount, jump at 150 percent, 200 percent, and 250 percent of FPL. The premium of the H plan (p_H) also jumps at these thresholds. Notably, p_H jumps by *more than* p_L at each of these thresholds. This occurs because CommCare chose to apply non-constant subsidies across plans with the goal of narrowing premium differences across plans for lower-income groups. Importantly for our demand estimation, this subsidy structure creates variation in both premium levels and differences between H and L . Specifically, the difference $p_H - p_L$ grows from \$11 below 150 percent FPL, to \$19 from 150–200 percent FPL, to \$29 from 200–250 percent FPL, and to \$31 above 250 percent FPL.

The final relevant option for people eligible for CommCare was to remain uninsured and pay a penalty for being uninsured, the so-called “mandate penalty” which increased the cost of remaining uninsured. The dotted gray line in Figure 3 shows the statutory mandate penalty at each income, which the state set to be half of the lowest CommCare premium (p_L). In practice, however, the actual penalty an individual would owe likely diverges from the gray line for two reasons. First, the mandate is assessed based on total annual income (reported in end-of-year tax filings), whereas the measure used to determine enrollee premiums is self-reported on the program application and measures income over the last 12 months (e.g., the prior July to June for someone enrolling during open enrollment). Thus, the actual *expected*

¹⁰One plan (Fallon Community Health Plan) was only active in central Massachusetts, so its network is difficult to compare to the other insurers.

¹¹Consistent with this perception, when all plans were available for free, which was the case for enrollees below 100 percent of FPL, 96 percent of enrollees chose a plan other than CultiCare.

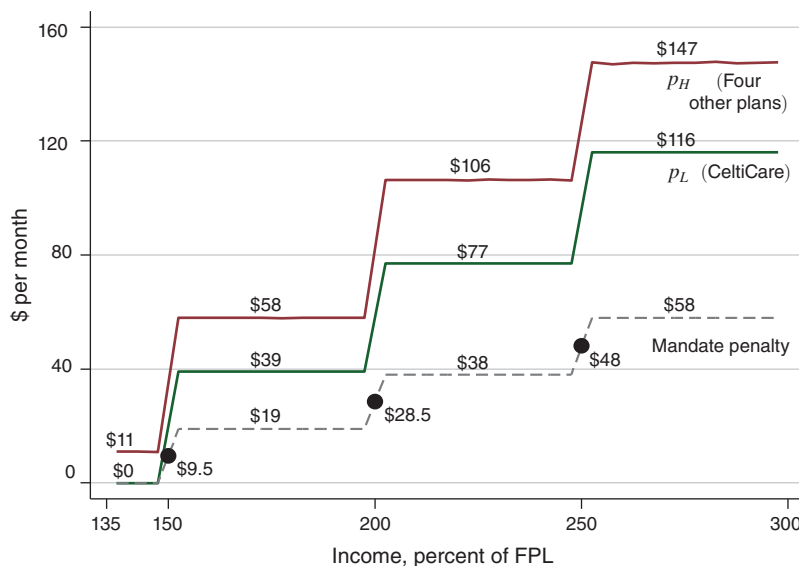


FIGURE 3. PREMIUM AND MANDATE PENALTY VARIATION, 2011

Notes: The figure shows how 2011 enrollee premiums and the mandate penalty vary across incomes (as a percent of the federal poverty level, FPL). p_L denotes the enrollee premium for the L plan (CeltiCare), while p_H is the share-weighted average of the enrollee premiums in the four H (non-CeltiCare) plans. *Mandate Penalty* (dashed gray line) is the statutory mandate penalty at each income level. The black dots show the expected mandate penalty for a person near the income discontinuities, which is the average of the two mandate penalties on either side of the discontinuity.

mandate penalty is unlikely to change discontinuously at the income thresholds, since someone just above a threshold is equally likely to have total annual income (relevant for the mandate) above or below the threshold. Figure 3 shows in black dots the expected mandate penalty for individuals near each threshold, which we assume is simply the average of the statutory penalty above and below the threshold. A second reason the actual mandate penalty may differ is that individuals may be able to avoid paying even if they are uninsured. For instance, the law does not impose a penalty if an individual is uninsured for three or fewer consecutive months during the year or if an individual qualified for a religious or hardship exemption.¹²

It is unclear how to use the mandate penalty when calculating revealed willingness to pay. For the reasons discussed above, an individual's actual mandate penalty is difficult to determine. Moreover, individuals may discount the mandate penalty because it is incurred in the following year's taxes, or even be unaware of it. In our baseline demand estimates, we will use the sticker premiums for insurance, effectively ignoring the saved penalty when an individual buys insurance. This will make our estimates a conservative upper bound on individuals' willingness to pay for insurance. In robustness analysis in Section IIID, we also report the lower

¹²The three-month exception is empirically important: based on a state report, almost 40 percent of the 183,000 uninsured people above 150 percent FPL in 2011 were uninsured for three or fewer months (Massachusetts Health Connector and Department of Revenue 2011).

willingness to pay estimates that we find when we normalize premiums by the expected mandate penalty values (shown in black dots).

B. Data

Administrative Data: Enrollee Plan Choices, Claims, and Demographics.—Our primary data are enrollee-level and claim-level administrative data from the CommCare program for fiscal years 2009–2013. We observe enrollee demographics and monthly plan enrollment linked to data on their claims and risk scores. All data are at the individual level because CommCare only offers individual (not family) coverage.¹³

We observe each enrollee's chosen plan at a monthly level. We define total enrollment as the annualized number of enrollee months in CommCare or a specific plan. In practice, most enrollees are in the same plan for the whole year. We also observe enrollees' choice sets, including the prices, subsidies, and enrollee premiums of each option (summarized in Figures 2 and 3). Enrollee-linked insurance claims data allow us to measure each person's monthly costs (both insurer-paid and out-of-pocket).

The most important demographic we observe is the individual's family income as a percent of FPL (rounded to the 0.1 percent level), which is the running variable for our RD analysis. This variable is calculated by the regulator from information on family income and composition that enrollees report in their initial CommCare application, and is used to determine premiums and subsidies. This variable is updated based on any subsequent known changes, which in principle, enrollees are required to self-report when they occur, and based on information from annual audits. We also observe information from CommCare's records on enrollee age, gender, zip code of residence, and risk score, a measure of predicted spending calculated by CommCare.

Throughout our analysis, we limit attention to individuals above 135 percent of FPL because of the significant eligibility change at 133 percent FPL. Above this threshold, parents cease to be eligible for Medicaid and become eligible for CommCare. Table 1 reports some summary statistics from the administrative data for CommCare enrollees in fiscal year 2011 between 135 and 300 percent FPL. Ninety percent of CommCare enrollees are in the *H* plan, despite higher enrollee premiums (see Figure 3). About 20 percent of enrollees are between 135 and 150 percent of the federal poverty line. CommCare's subsidies are quite large. Average enrollee premiums (\$70 per month) cover less than 20 percent of insurer-paid medical costs (\$359/month) or prices (\$422).

Survey Data: Eligible Population for CommCare.—We supplement the administrative data on CommCare enrollment with estimates of the size of the CommCare-eligible population from the 2010 and 2011 American Community Survey (ACS), an annual 1 percent random sample of US households. We use these data to estimate the number of people eligible for CommCare in each 1 percent FPL

¹³ These de-identified data were obtained via a data use agreement with the exchange regulator, the Massachusetts Health Connector. Our study protocol was approved by the IRBs of the Connector and our affiliated institutions (Harvard, MIT, and NBER).

TABLE 1—COMM CARE SUMMARY STATISTICS IN 2011: PREMIUMS, ENROLLMENT, AND COSTS

	Any plan (1)	H plan (2)	L plan (3)
<i>Enrollment</i>			
Observations (number of unique individuals)	107,158	96,391	14,828
Average monthly enrollment	62,096	55,599	6,497
Share of enrollees (%)	100	90	10
Average income (% of FPL)	193	193	189
Share below 150% FPL	20	18	29
Average age	44.4	44.9	40.2
Share male (%)	41	40	47
<i>Medical claims (monthly)</i>			
Insurer-paid	\$358.5	\$377.3	\$197.9
Total	\$377.3	\$396.4	\$213.4
<i>Prices, subsidies, and premiums (monthly)</i>			
Insurer price	\$422.2	\$424.2	\$404.9
Enrollee premium	\$70.0	\$72.8	\$46.0
Public subsidy	\$352.2	\$351.4	\$358.9

Note: Table shows summary statistics from CommCare administrative data for fiscal year 2011 for enrollees between 135 and 300 percent of FPL.

bin between 135 percent of FPL and 300 percent FPL. To be coded as eligible, an individual must live in Massachusetts and be: a US citizen, age 19–64, not enrolled in another form of health insurance (specifically, employer insurance, Medicare, or Tricare), and not eligible for Medicaid (based on income and demographics).¹⁴ Because the ACS is a 1 percent sample (and because of clustering in reported income at round numbers), our raw estimates of the size of the eligible population by 1 percent FPL bin are relatively noisy. We therefore smooth the estimates using a regression of raw counts by 1 percent FPL bin on a polynomial in income. Online Appendix B reports additional details on sample construction and shows the raw counts of eligibles by FPL, as well as the smoothing regression fit.

Rather than use the ACS estimates directly to estimate the size of the eligible population, we use it to estimate two statistics: the shape of the eligible income distribution and the average take-up rate for our study population. We do this because comparing the raw implied counts of the eligible population in the ACS to the number enrolled in CommCare from our administrative data would imply that only 37 percent of eligible individuals enroll in CommCare. This number seems low compared to the take-up estimate in the ACS data, where we find that 63 percent of eligible individuals report having insurance (see online Appendix B for details). We suspect the ACS take-up estimate is closer to the truth since it closely matches estimates from a Massachusetts health insurance survey in the fall of 2010 (Long, Stockley, and Dahlen 2012) and estimates from tax filing data.¹⁵ We conservatively

¹⁴The ACS does not distinguish Medicaid from CommCare coverage (both are coded as “Medicaid/other public insurance”), so we cannot directly exclude Medicaid enrollees.

¹⁵Specifically, Long, Stockley, and Dahlen (2012) find a 90 percent insurance coverage rate for non-elderly adults below 300 percent of FPL, nearly identical to the 89 percent rate we estimate using the ACS. Further, using statistics from state tax filings (used to determine who owes the mandate penalty), we estimate that about 107,000 tax filers earning more than 150 percent of the FPL were uninsured at a typical point in time during 2011

use the higher take-up estimate internal to the ACS and show in sensitivity analysis that if we instead use the ACS estimates of the eligible population directly, this produces a substantially lower take-up rate and in turn yields even lower estimates of the share insured under a given subsidy scheme.

Specifically, we take our estimates of the number of eligible individuals in the ACS by FPL bin (see panel A of online Appendix Figure 14) and scale the whole distribution down by a constant multiple (of 0.59) so that dividing the administrative count of enrollees by our adjusted eligible population size yields an average take-up rate of 63 percent (the rate calculated in the ACS).

Measuring the eligible population is difficult, and our approximation is, of course, imperfect. Fortunately, as we discuss in more detail below, the exact size of this population is not critical to estimate *changes* in enrollment and costs at the income discontinuities. Using this information (from administrative data alone), we can generate our key result: that willingness to pay is far below costs for marginal enrollees who drop coverage at each discontinuity. However, the ACS estimates are important for understanding *what share* of the eligible population these marginal enrollees comprise and *where* in the population distribution they lie. This is also necessary for translating our results into estimates of take-up shares under various subsidy policies. As discussed, in this sense, our baseline approach is a conservative one.

Figure 4 shows our (smoothed) estimate of the size of the eligible population by FPL bin. Note that the decline in the estimated number of eligible people by income does not reflect the shape of the overall income distribution in that range, but rather the shape of the *eligible population* income distribution. Eligibility requires, among other things, that the individual not have access to employer-sponsored insurance, which tends to increase with income (Janicki 2013). For comparison, Figure 4 also shows the raw counts of the number enrolled in any CommCare plan by FPL bin; we use the difference between the eligible population estimate and the number of CommCare enrollees as the number of people choosing uninsurance.

II. Descriptive Analysis

A. Regression Discontinuity Design

We use the discrete change in enrollee premiums at 150 percent, 200 percent, and 250 percent of FPL to estimate how demand and costs change with enrollee premiums. We estimate a simple linear RD in which we allow both the slope and the intercept to vary on each side of each threshold. Specifically, we run the following regression across income bins (b) collapsed at the 1 percent of FPL level:

$$(1) \quad Y_b = \alpha_{s(b)} + \beta_{s(b)} Inc_b + \epsilon_b,$$

(Massachusetts Health Connector and Department of Revenue 2011). This number is calculated from state-reported statistics on the number of full-year and part-year uninsured (separately for ≤ 3 months and > 3 months) and a mid-point assumption about the part-year groups' duration of uninsurance. From the ACS data, we estimate that there were 108,342 uninsured tax filers earning > 150 percent of FPL (treating each "health insurance unit" as a single tax-filer). These two estimates are remarkably close, suggesting that the ACS's uninsured estimates are accurate.

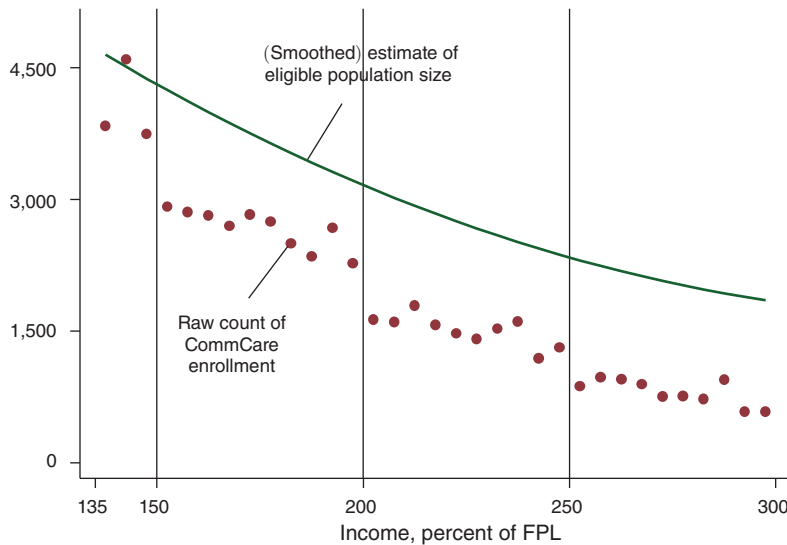


FIGURE 4. ELIGIBLE AND ENROLLED POPULATION, 2011

Notes: Figure shows our (smoothed) estimate of the CommCare-eligible population in 2011 (based on ACS data), and raw enrollment counts in CommCare in 2011 by bins of 5 percent of the FPL.

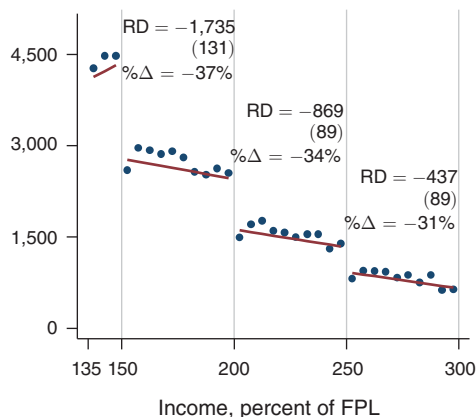
where Y_b is an outcome measure in that income bin b , Inc_b is income (as a percent of FPL) at the midpoint of the bin, and $s(b)$ is the income segment on which bin b lies (either 135–150 percent, 150–200 percent, 200–250 percent, or 250–300 percent FPL). Notice that the unit of observation is the income bin, while the slope and intercept coefficients vary flexibly at the segment level. Our outcomes are either measures of plan enrollment shares, or enrollee costs or characteristics. We run all regressions using bin-level data and report robust standard errors.

The key assumption is that the eligible population size is smooth through the income thresholds at which subsidies change (150 percent, 200 percent, and 250 percent FPL). This would be violated if people strategically adjust (or misreport)¹⁶ their income to get just below the thresholds and qualify for a larger subsidy.¹⁷ While in principle such manipulation would be possible, in our setting the process by which individuals' reported incomes were translated into the percent of FPL formula for determining subsidies were largely shrouded from the individuals during the application process. Perhaps as a result, we find minimal evidence of any such manipulation (see Section IIID). Moreover, because of the relatively linear patterns we find away from the discontinuity, alternative methods (such as constructing a donut-hole around the discontinuity) would lead to very similar estimates.

¹⁶Enrollees were required to show proof of income (e.g., via recent pay stubs) when applying but in theory could adjust hours or misreport self-employment income to get below subsidy thresholds.

¹⁷In addition, there are minor changes in eligibility just above 200 percent FPL (pregnant women and HIV-positive people lose Medicaid eligibility and become eligible for CommCare) that also technically violate the smoothness assumption. This will bias our RD estimate of demand responsiveness to price slightly toward zero, since the eligible population grows just above 200 percent FPL. In sensitivity analysis, we show that our main results are robust to excluding this discontinuity.

Panel A. Average monthly enrollment by income



Panel B. Average monthly insurer costs

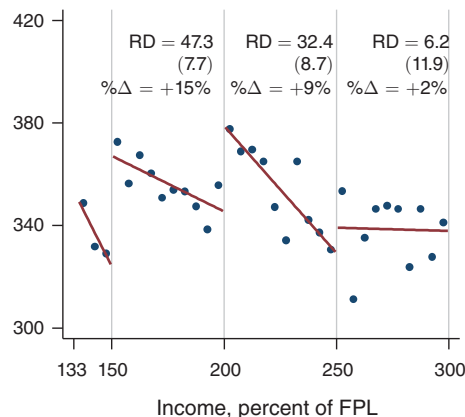


FIGURE 5. COMM-CARE ENROLLMENT AND AVERAGE INSURER COSTS, 2009–2013

Notes: The figure shows discontinuities in enrollment and average insurer costs at the income thresholds (150 percent, 200 percent, and 250 percent of FPL) at which enrollee premiums increase (see Figure 2). Panel A shows average enrollment in CommCare (total member-months, divided by number of months) by income over the 2009–2013 period our data span. Panel B shows average insurer medical costs per month across all CommCare plans over the same period. In each figure, the dots represent raw values for a 5 percent of FPL bin, and the lines are predicted lines from our linear RD specification in equation (1). RD estimates and robust standard errors (in parentheses) are labeled just to the right of each discontinuity; percent changes relative to the value just below the discontinuity are labeled as $\% \Delta =$.

B. Evidence from Pooled Years, 2009–2013

Insurance Demand.—Figure 2 showed that premiums increase discontinuously at 150, 200, and 250 percent of FPL. Panel A of Figure 5 shows that enrollment drops significantly at each of these income thresholds. Specifically, the figure plots average monthly enrollment in the CommCare market over the 2009–2013 period by income bin. It also superimposes estimates from the linear RD model in equation (1), using average monthly enrollment as the dependent variable. At each of the three discontinuities, enrollment falls by 30 percent to 40 percent. All three changes are statistically distinguishable from zero ($p < 0.01$).

Cost of Insurance and Adverse Selection.—Panel B plots average insurer costs by income bin, again superimposing estimates from the linear RD model in equation (1). Average insurer costs are defined as the average claims paid by the insurer for the set of people who are enrolled in that month.

The figure shows that average costs of the insured rise as the enrollee premium increases. For example, we estimate a discontinuous increase in costs of \$47.3 (standard error \$7.7) per enrollee per month at 150 percent FPL and of \$32.4 (standard error \$8.7) at 200 percent of FPL. We find a smaller but noisier increase of \$6.2 (standard error \$11.9) at 250 percent FPL; this imprecision may reflect the smaller size of the eligible and enrolled populations at 250 percent of FPL (see Figure 4).

These patterns indicate the presence of adverse selection: increases in average costs indicate that the marginal enrollees (who exit in response to the premium

increase) are less costly than the average enrollee who remains. An alternative way to test for adverse selection would be to examine whether characteristics of the enrollees that are associated with higher costs and not priced by insurers increase when premiums rise (Finkelstein and Poterba 2014). In online Appendix Figure 23 we also show that, consistent with adverse selection, the average age and risk score (i.e., predicted medical spending) of enrollees increase at these income thresholds. In other words, in response to higher premiums, younger and lower-risk enrollees are more likely to leave the market. These results are, not surprisingly, more precisely estimates than the analysis of realized insurer claims in panel B of Figure 5. The claims measure, however, captures all potential dimensions of selection (both observable factors that go into the risk score factors that do not).

A priori, it was unclear whether this market would suffer from adverse selection. On the one hand, insurers were not allowed to vary prices based on individuals' health characteristics (such as age, gender, or preexisting conditions); this would tend to generate adverse selection. On the other hand, in an effort to combat adverse selection, Massachusetts imposed a mandate on individuals to buy coverage, backed up by financial penalties. Our results suggest the coverage mandate and associated penalty were not sufficient to prevent adverse selection.

C. Evidence from 2011

In most of the rest of the paper, we study data from fiscal year 2011, which have the convenient vertical differentiation of plans discussed above. Here we present reduced-form evidence on demand and costs for 2011 alone, focusing on overall enrollment and enrollment in the *H* plan.

Insurance Demand.—Figure 6 shows statistically significant (at the 1 percent level) declines in overall CommCare and *H* plan enrollment at each enrollee premium threshold (see Figure 3). The drops in enrollment do not occur only when premiums rise from zero to a positive amount (150 percent FPL threshold): enrollment falls by 20–30 percent at all three thresholds.

Figure 7 transforms these raw enrollment counts into market shares, using our estimate of the eligible population (see Figure 4) as the denominator. Panel A shows that the share enrolled in any CommCare plan falls by a statistically significant 24–27 percent at each discontinuity at which enrollee premiums rise by \$38–39 per month. The size of these *percent* drops are identified directly from the fall in enrollment in the administrative data. But, we can also use our estimate of the size of the eligible population from the ACS to make inferences about the levels of take-up rates as a function of the enrollee premium. Take-up rates fall from 94 percent when insurance is free (below 150 percent FPL) to 70 percent where the cheapest premium increases to \$39 per month (above 150 percent FPL). Take-up rates continue to fall with premiums, declining to below 50 percent as the cheapest premium rises to \$116 per month (above 250 percent FPL).¹⁸

¹⁸The pattern of enrollment by income within a constant premium range should not be interpreted with caution; demand is estimated conditional on eligibility and, as can be seen in Figure 4, eligibility declines with income. Therefore, the sample is differentially selected by income, since the set of higher income people without access to

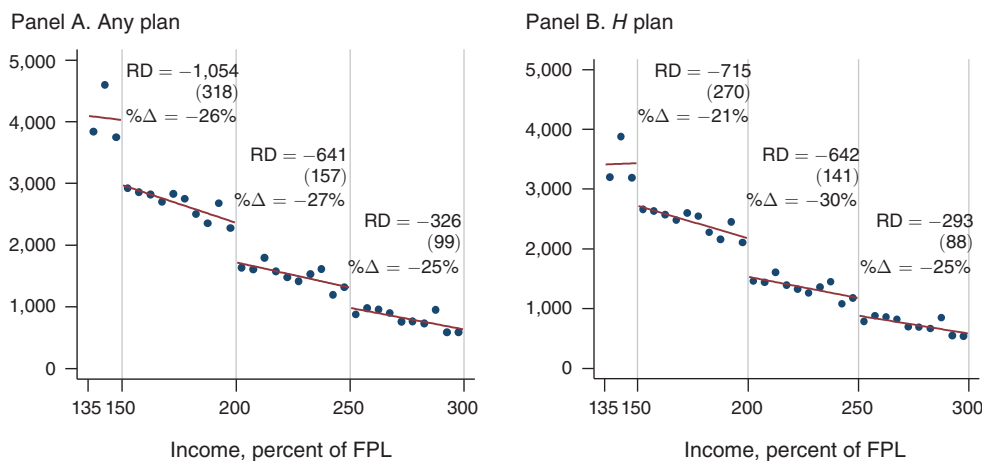


FIGURE 6. COMM-CARE ENROLLMENT, 2011

Notes: Figure shows average enrollment (defined as total member-months, divided by number of months) by income in 2011. Panel A shows enrollment in any CommCare plan, panel B shows enrollment in the *H* plan. In each figure, the dots represent raw averages for a 5 percent of FPL bin, and the lines (and labels) are predicted lines from our linear RD specification in equation (1). RD estimates and robust standard errors (in parentheses) are labeled just to the right of each discontinuity; percent changes relative to the value just below the discontinuity are labeled as %Δ =.

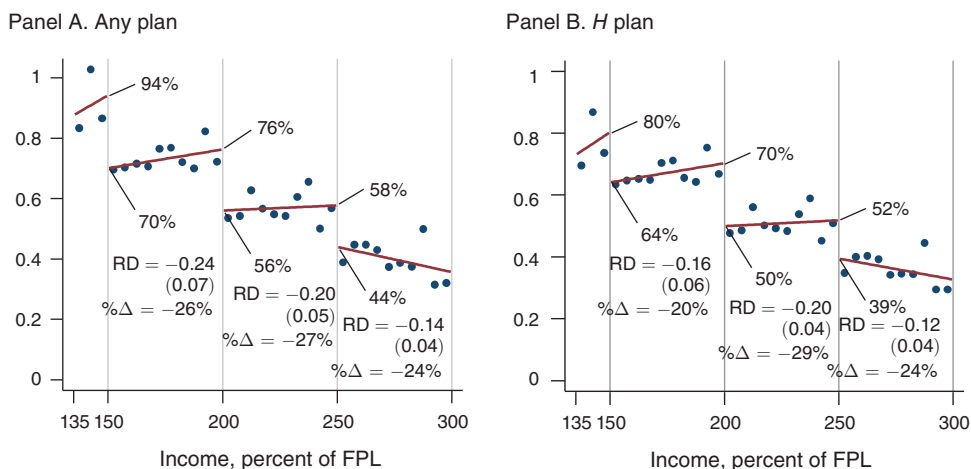


FIGURE 7. SHARE OF ELIGIBLE POPULATION ENROLLED, 2011

Notes: Figure shows share of eligible population enrolled (in bins of 5 percent of the federal poverty level) in any CommCare plan (panel A), and in the *H* plan (panel B). In each figure, the dots represent raw averages for a 5 percent of FPL bin, and the lines (and labels) are predicted lines from our linear RD specification in equation (1). RD estimates and robust standard errors (in parentheses) are labeled below each discontinuity; percent changes relative to the value just below the discontinuity are labeled as %Δ =.

Cost of Insurance.—Figure 8 shows that average monthly insurer costs rise as enrollee premiums increase at each income threshold. Panel A indicates that at 150 percent of FPL, average costs for CommCare enrollees increase by \$47 (about

employer-provided health insurance (an eligibility criteria) naturally is differentially selected than the set of lower income individuals without access.

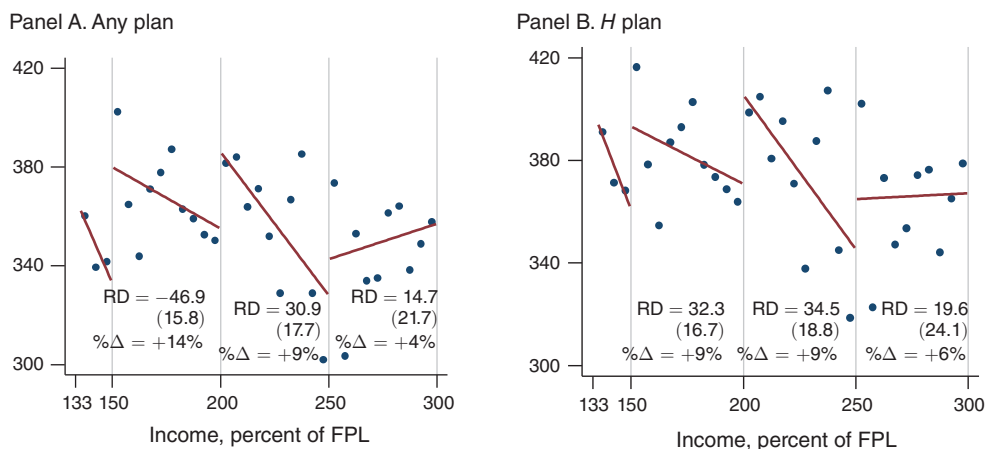


FIGURE 8. DISCONTINUITIES IN AVERAGE INSURER COSTS, 2011

Notes: Figure shows average monthly insurer medical costs for enrollees (in bins of 5 percent of the federal poverty level) for any CommCare plan (panel A) and the *H* plan (panel B). In each figure, the dots represent raw averages for a 5 percent of FPL bin, and the lines are predicted from our linear RD specification in equation (1). RD estimates and robust standard errors (in parentheses) are labeled below each discontinuity; percent changes relative to the value just below the discontinuity are labeled as $\% \Delta$.

14 percent); this is statistically distinguishable from zero at the 1 percent level. We also see increases in average costs at the 200 percent and 250 percent thresholds, but the increases are somewhat smaller (\$31 and \$15, or 9 percent and 4 percent) and less precisely estimated. These magnitudes are similar to the more precise estimates for the pooled 2009–2013 years shown above.

Panel B shows analogous estimates for the 2011 enrollees in the *H* plan. Again, we see increases in average costs at all three discontinuities. However, these are less precisely estimated.

III. Willingness to Pay and Cost Curves

We use a model of insurance demand and cost to map the 2011 descriptive results into estimates of willingness to pay and cost curves that we use for counterfactual analysis. The model follows Einav, Finkelstein, and Cullen (2010), but incorporates three plan options: the *H* plan, the *L* plan, and uninsured (*U*), as opposed to a binary model considered in Einav, Finkelstein, and Cullen (2010). Motivated by our institutional setting, we assume a vertical model of insurance demand. The vertical structure is helpful for tractability. We show in the sensitivity analysis of Section IIID that we can derive fairly tight bounds on willingness to pay that are similar to our point estimates below but do not rely on the vertical model assumptions.

A. Setup and Assumptions

Consider an insurance market where contracts j are defined by a generosity metric α . We assume there are two formal insurance contracts $j = H$ and L , with $\alpha_H > \alpha_L$. In addition, there is an outside option of being uninsured, U , which is

weakly less generous than L ($\alpha_U \leq \alpha_L$). Let $w(\alpha; i)$ be the (dollar) willingness to pay (WTP) of consumer i for an α -generosity contract. Let p_{ij} be the enrollee premium of contract j , and normalize $p_{iU} = 0$ so that premiums are defined relative to U . Finally, there is an (additively separable) “hassle cost” of the enrollment process for contract j , h_j . We normalize $h_U = 0$ and assume the formal insurance contracts H and L involve the same hassle cost $h_H = h_L \equiv h$. This hassle cost, h , will be positive if enrolling in formal insurance involves a greater hassle relative to remaining uninsured (e.g., due to the hassle of applying for insurance and making an active plan choice) and negative if staying uninsured involves greater hassle (or stigma).

With these assumptions, we write the utility of consumer i for plan j as

$$u_{ij} = w(\alpha_j; i) - h - p_{ij}, \quad \text{for } j \in \{L, H\},$$

$$u_{iU} = w(\alpha_U; i).$$

We denote the willingness to pay W_j for plan j relative to U as

$$W_j(i) = (w(\alpha_j; i) - w(\alpha_U; i)) - h, \quad \text{for } j \in \{L, H\},$$

which is the maximum price at which the consumer would choose plan j over U . We denote the willingness to pay $\Delta W_{HL}(i)$ for plan H relative to plan L as

$$(2) \quad \Delta W_{HL}(i) \equiv W_H(i) - W_L(i) = w(\alpha_H; i) - w(\alpha_L; i).$$

We impose a vertical demand model (see Tirole 1988) using the following two assumptions about preferences.

ASSUMPTION 1 (Vertical Preferences for Generosity): *Everyone prefers more generous contracts: $w(\alpha; i)$ is increasing in α .*

ASSUMPTION 2 (Single Dimension of Heterogeneity in Value for Generosity (Increasing Differences)): $w(\alpha; i) = w(\alpha; s)$, where $1 - s \in [0, 1]$ indexes increasing value for generosity, with $dw(\alpha; s)/d(1 - s) > 0$ and $d^2 w(\alpha; s)/d\alpha d(1 - s) > 0$.

Note that we use $1 - s$ as the index of WTP for generosity, i.e., $s = 0$ is the highest WTP type and $s = 1$ is the lowest. This ensures that s is the x -axis value on a standard demand curve (where highest-WTP types are on the left) and simplifies notation for our graphical analysis below.

Assumption 1 implies that $W_H(i) > W_L(i)$ for all i . It thus rules out cases in which people disagree about the quality of plans H and L (i.e., different people would prefer H or L at the same price). As noted in Section I, the data are consistent with this vertical assumption: when the price of H and L are the same, specifically CommCare enrollees below 100 percent of FPL for which all plans are free, 96 percent of enrollees choose H .

Assumption 2 imposes increasing differences in WTP for generosity. This means that both the value for H relative to L and value for L relative to U are increasing

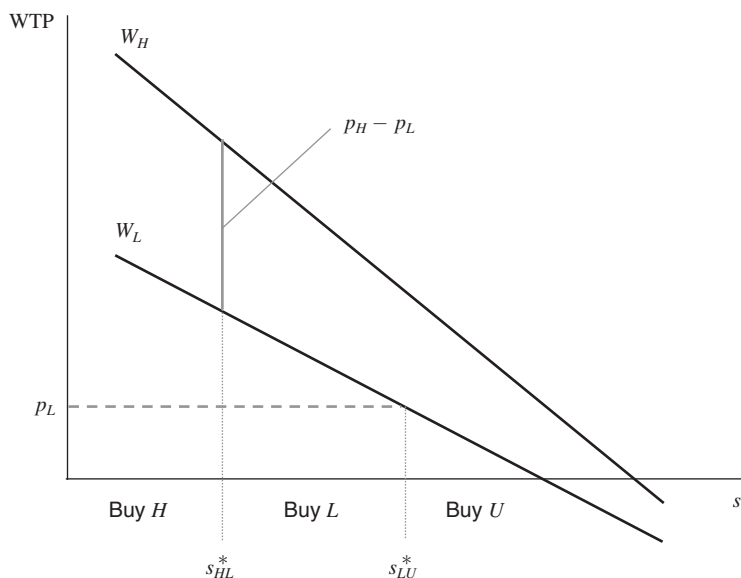


FIGURE 9. WILLINGNESS TO PAY CURVES: MODEL

Notes: Figure shows the theoretical implications of our vertical model for the willingness to pay (W_j) curves for the H and L plans. The model assumptions imply that both W_H and W_L are downward sloping (i.e., decreasing with s) and that the gap between $W_H - W_L$ is also narrowing as s increases. Under the positive demand condition for prices (which this graph assumes), the lowest- s types (furthest left on the x -axis) buy H , middle- s types (between s_{HL}^* and s_{LU}^*) buy L , and the highest- s types choose U .

in a single index of preferences, $1 - s$.¹⁹ This rules out cases in which the people who value L relative to U by more than average also care less than average about H relative to L , and vice versa.

Demand Curves.—We define the demand for product $j \in \{U, L, H\}$, $D_j(p_L, p_H)$, as the fraction of the population purchasing j at prices $\{p_L, p_H\}$. Assuming that prices are such that there is positive demand for all contracts, Assumptions 1 and 2 imply that those with the lowest s choose H , those with the highest s choose to remain uninsured, U , and those with interim values of s choose L . Moreover, with positive demand for all plans, the model has the tractable feature that demand depends only on *price differences* between adjacently ranked options. Specifically, D_H depends only on $p_H - p_L$, D_U depends only on p_L ²⁰ and D_L depends on both $p_H - p_L$ and p_L .

Figure 9 illustrates how the willingness to pay curves translate into the fraction enrolling in each plan. The figure plots W_L and W_H against a horizontal axis of s so that these curves are downward sloping. The vertical preference assumption (1) implies that $W_H > W_L$ at all points. The increasing differences assumption (2) implies that the gap $W_H - W_L$ widens as W_L increases (as one moves left).

¹⁹Note that Assumption 2 ensures that a single crossing property holds and generalizes the standard assumption in a vertical model of demand (Tirole 1988). The standard vertical model assumes that $v(\alpha_j; s) = \beta(1 - s) \cdot \alpha_j$ so that choice-specific utility equals $\beta(1 - s) \cdot \alpha_j - p_j$, where $\beta(1 - s)$ is the value of generosity for type s (with $\beta(1 - s) > 0$). This model satisfies our Assumption 2.

²⁰In general, demand for U would depend on $p_L - p_U$, but p_U is normalized to zero.

We denote the point s_{HL}^* to be the point of indifference between L and H , which occurs where the vertical distance between $W_H - W_L$ equals $p_H - p_L$. All types to the left of this enroll in H , and the demand for H equals s_{HL}^* . Likewise, the point s_{LU}^* at which p_L intersects the $W_L(s)$ curve determines the person who is indifferent between L and U . All types to the right of s_{LU}^* remain uninsured, and those just to the left enroll in the L plan. Mathematically, these points s_{HL}^* and s_{LU}^* are defined by the equations:

$$(3) \quad \begin{aligned} \Delta W_{HL}(s_{HL}^*) &\equiv W_H(s_{HL}^*) - W_L(s_{HL}^*) = p_H - p_L, \\ W_L(s_{LU}^*) &= p_L. \end{aligned}$$

Given these definitions, a necessary and sufficient condition for demand for all contracts to be positive is for

$$\text{Positive Demand Condition:} \quad 0 < s_{HL}^* < s_{LU}^* < 1.$$

Without loss of generality (since it is an arbitrary index), we assume a uniform distribution over s types. Because $\Delta W_{HL}(s)$ and $W_L(s)$ are monotonically decreasing functions (by Assumption 2), the equations in (3) implicitly define $s_{LU}^* = W_L^{-1}(p_L)$ and $s_{HL}^* = \Delta W_{HL}^{-1}(p_H - p_L)$. Define the demand for product $j \in \{U, L, H\}$ as the fraction of the population purchasing j at prices $\{p_L, p_H\}$. Assuming the positive demand condition holds,²¹ these are given by

$$(4) \quad \begin{aligned} D_H(p_H - p_L) &= s_{HL}^* = \Delta W_{HL}^{-1}(p_H - p_L), \\ D_L(p_L, p_H - p_L) &= s_{LU}^* - s_{HL}^* = W_L^{-1}(p_L) - \Delta W_{HL}^{-1}(p_H - p_L), \\ D_U(p_L) &= 1 - s_{LU}^* = 1 - W_L^{-1}(p_L), \end{aligned}$$

where the demand for H only depends on the price difference $p_H - p_L$, and the demand for L depends on both p_L and $p_H - p_L$. We will often analyze “demand for formal insurance” (i.e., pooled demand for H or L) which is calculated from the equations above as $1 - D_U$, which depends only on p_L , not p_H .

Insurer Costs.—We denote by $C_j(s)$ the expected costs to the insurer of enrolling type s in plan j .²² As is standard in the literature, we define insurer costs as medical claims paid and abstract from administrative costs. We also adopt the standard assumption that $C_j(s)$ is independent of the premium charged for the insurance plan.

²¹ Practically speaking, in our empirical setting, we observe positive demand for all products, so we will assume the positive demand condition holds (though it would be conceptually simple to generalize these curves to the more general case).

²² In a setting with a binary contract choice (as in Einav et al. 2010a), the variation in $C_j(s)$ with respect to s is referred to as the marginal cost curve for contract j ; with three contracts as we have here, there can be two different margins of selection into a contract and so the “marginal cost curve” language is less useful.

We define average costs $AC_j(s)$ as the average costs of all individuals with type $\tilde{s} \leq s$:

$$(5) \quad AC_j(s) = \frac{1}{s} \int_0^s C_j(\tilde{s}) d\tilde{s},$$

where recall that we have assumed $s \sim U[0, 1]$. If premiums are such that all types $\tilde{s} \leq s$ choose the j plan, then the cost imposed on the insurer would be given by $AC_j(s)$.

B. Constructing Willingness to Pay and Cost Curves

*Willingness to Pay (W_j).—*We combine the modeling assumptions above with the empirical patterns documented in Figure 7 to construct the empirical analogues of the $W_H(\cdot)$ and $W_L(\cdot)$ curves in Figure 9. Figures 10 and 11 walk through this exercise.

Panels A and B of Figure 10 plot the willingness to pay curves for the L contract ($W_L(\cdot)$) and for the H contract relative to L ($\Delta W_{HL}(\cdot)$), respectively. Each line segment represents points derived from our three income RDs at 150 percent FPL (in blue), 200 percent FPL (red), and 250 percent FPL (green). Equation (4) shows that $W_L^{-1}(p_L)$ equals $1 - D_U(p_L)$, the share of people who purchase formal insurance at enrollee premium p_L . Therefore, we obtain the W_L curve by plotting observations of $(1 - D_U, p_L)$ derived from market shares and premiums around each income discontinuity from our RD estimates (see Figures 3 and 7). Similarly, equation (4) shows that $\Delta W_{HL}^{-1}(p_H - p_L)$ equals $D_H(p_H - p_L)$, the share of people who purchase the H plan at premium difference $p_H - p_L$. We therefore obtain the ΔW_{HL} curve by plotting observations of $(D_H, p_H - p_L)$ from the same RD estimates.

In principle, we could identify part of a willingness to pay curve using only one premium discontinuity. In practice, we combine the data from all three discontinuities because this lets us observe demand over a wider range of premiums. As a result, at two of the enrollee premiums, we observe (and plot) two different market shares. This is because each pricing discontinuity identifies a demand curve for individuals at a given income level, and these demand curves need not be the same. For example, the premiums that apply between 150–200 percent FPL identify one point on the demand curve for 150 percent FPL (“from the right”) and one point on the demand curve for 200 percent FPL (“from the left”).

In practice, we observe that demand in fact varies little with income. In other words, market shares are relatively flat within an income range that has constant premiums, as was evident in Figure 7.²³ As a result, the demand line segments for the three income groups (shown in different colors in Figure 10) nearly intersect.

To adjust for remaining differences in demand across incomes, we extend our theoretical framework above to allow willingness to pay for insurance to vary with income, y . We define our index s conditional on a fixed income level, y , and we denote $w(\alpha; s, y)$ to be the willingness to pay of a type s for a single income group, y .

²³This does not necessarily imply that income effects of insurance demand are small; recall that the eligible population consists of people who, among other things, do not have access to employer-provided health insurance. The nature of the eligible population may therefore be changing with income as well.

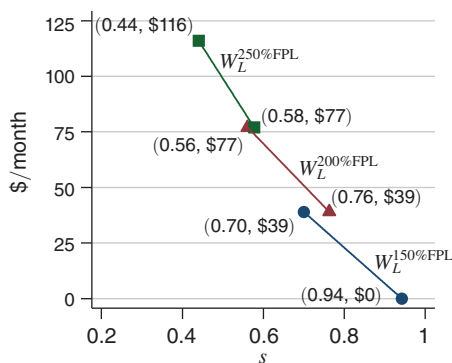
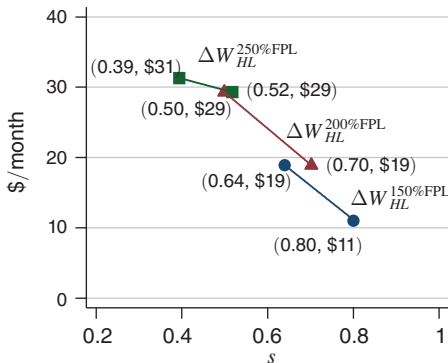
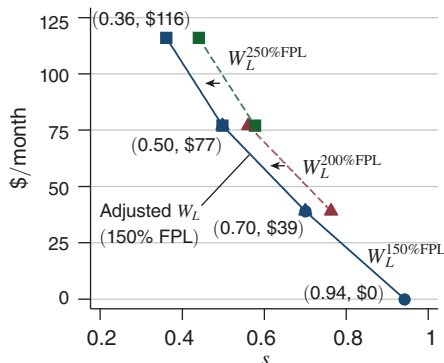
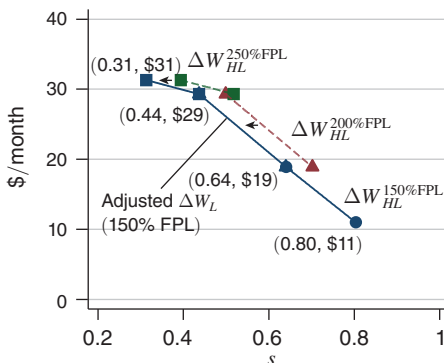
Panel A. W_L (based on $1 - D_U$)Panel B. ΔW_{HL} (based on D_H)Panel C. Adjusted W_L Panel D. Adjusted ΔW_{HL} 

FIGURE 10. WILLINGNESS TO PAY CURVES: EMPIRICAL

Notes: Figures show our construction of the willingness to pay curves (W_L and ΔW_{HL}) based on the demand points in our RD estimates in Figure 7 and the premium variation at each discontinuity from Figure 3. Panel A shows the W_L points, each of which represents an observation of $(1 - D_U, p_L)$ drawn from either side of our income discontinuities at 150 percent, 200 percent, and 250 percent FPL. Panel B shows the ΔW_{HL} points, each of which is an observation of $(D_H, p_H - p_L)$ from either side of the discontinuities. Panel C and panel D show how we adjust the W_L and ΔW_{HL} curves horizontally to line up with the 150 percent FPL line segment.

To allow us to combine demand information across income groups, we assume that income functions as a horizontal shifter of the willingness to pay curves:

$$w(\alpha; s, y) = w(\alpha; s + \lambda_y).$$

This assumption implies, for example, that $W_L^{150\%}(s) = W_L^{200\%}(s + \lambda_{200\%} - \lambda_{150\%})$.

Panels C and D of Figure 10 illustrate the implications of this assumption graphically. Specifically, we horizontally shift the panel A and B willingness to pay curves estimated at the discontinuities at 200 percent FPL and 250 percent FPL so that willingness to pay (i.e., demand) lines up with the curve estimated at 150 percent FPL. We chose to line up the curves at 150 percent of FPL since that is the threshold with the greatest share of the eligible population (see Figure 4); results would be qualitatively similar if we instead created a demand curve at the 200 percent or 250 percent FPL threshold. In practice, since demand is relatively flat with respect

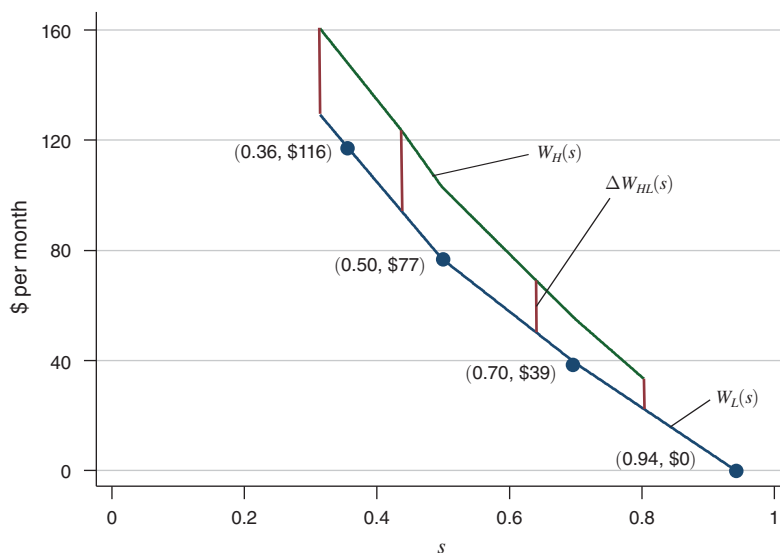


FIGURE 11. FINAL WILLINGNESS TO PAY CURVES

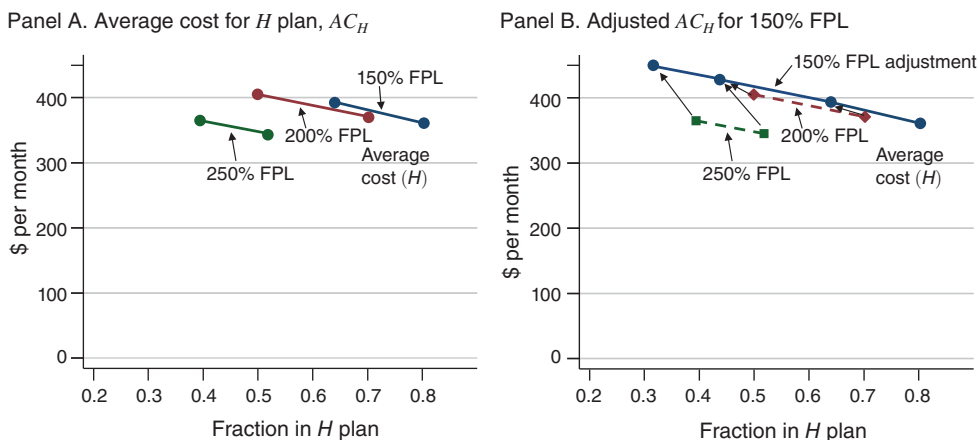
Notes: Figure shows our final estimated WTP curves. The blue curve is the adjusted W_L curve shown in panel C of Figure 10, with the large dots representing observed points. The W_H curve (in green) is constructed by vertically summing the W_L and ΔW_{HL} curve (as shown in panel D of Figure 10) at each x -axis value (of $1 - s$). The red vertical bars in the figure represent the observed points of the adjusted ΔW_{HL} curve. We extrapolate the W_L curve slightly out of sample (from $1 - s = 0.36$ down to $1 - s = 0.31$) to be able to add on the final point of the ΔW_{HL} curve.

to income, the shift is not very large: the 200 percent FPL curve is shifted leftward by 6 percentage points in s space, and the 250 percent FPL curve is shifted leftward by an additional 2 percentage points. The resultant willingness to pay curves consist of three piecewise linear segments.

Our horizontal shift approach assumes that we can infer the slope of the WTP curve for people at 150 percent of FPL at higher prices than we observe in the data from the slope of the WTP curve slope at these higher prices for people at 200 percent and 250 percent of FPL. While we cannot test this assumption, our sense is that it is likely to be conservative (in the sense of slightly overstating WTP at 150 percent of FPL); because the 150 percent FPL enrollees are poorer, we might expect them to drop out of the market more quickly at higher prices than do the 200 percent and 250 percent FPL enrollees.²⁴

Finally, in Figure 11 we use our estimates of $W_L(\cdot)$ and $\Delta W_{HL}(\cdot)$ from panels C and D of Figure 10 to construct $W_H(s)$ as $W_L(s) + \Delta W_{HL}(s)$ using equation (2). The resulting W_L and W_H curves allow us to infer willingness to pay for L and H for people earning 150 percent of FPL. Willingness to pay for H is nontrivially higher than for L . The additional value (ΔW_{HL}) ranges from \$11 to \$31 per month, or 11–30 percent of the median type's WTP for L . The median type has total WTP

²⁴Consistent with this, if we had simply linearly extrapolated the line segment for 150 percent FPL leftward, we would generate a slightly lower W_L curve. However, the difference would not be large: $W_L(0.50)$ would be \$71 (versus \$77 in our estimates) and $W_L(0.36)$ would be \$94 (versus \$116 in our estimates). The linearly extrapolated ΔW_{HL} curve would be even closer, never differing by more than \$3 from our version. The similarity of these estimates gives us additional confidence that our assumption is a reasonable approximation.

FIGURE 12. CONSTRUCTION OF AVERAGE COST CURVES (H)

Notes: Panel A shows the raw average cost points for the H plan, drawn from the RD estimates around each of our three income discontinuities (see panel B of Figure 8; for convenience, online Appendix Table 6 summarizes those estimates). Panel B shows how we generate our adjusted AC_H for 150 percent FPL by translating the 200 percent and 250 percent FPL line segments to line up with the 150 percent FPL segment.

for H (W_H) of \$103/month. Using our in-sample variation, we can infer W_L over the range $s \in [0.36, 0.94]$, i.e., all but the highest 36 percent and lowest 6 percent of the WTP distribution. Similarly, our variation lets us infer W_H over the range $s \in [0.31, 0.80]$, i.e., all but the highest 31 percent and lowest 20 percent of the distribution.

Cost Curves.—In Figure 12, we construct the average cost curve for the H plan, AC_H . In panel A we plot estimated average costs for enrollees in the H plan on each side of the premium discontinuities (from panel B of Figure 8) against the shares in the H plan at each discontinuity (from panel B of Figure 7). For instance, just below the 150 percent FPL discontinuity, 80 percent are in the H plan, and the average cost is \$361. Just above the discontinuity, 64 percent of people are in the H plan and average cost is \$393. Therefore, the average cost curve for 150 percent FPL flows through the points (64 percent, \$393) and (80 percent, \$361).

In panel B, we once again adjust the average cost curves to obtain a single curve applicable to individuals at 150 percent of FPL. To do so, we assume that the slopes of the average cost curves are stable with income so that one can vertically shift the average cost curves for the 200 percent FPL and 250 percent FPL thresholds to align with the 150 percent FPL average cost curve.²⁵ To be consistent with how we adjusted the W_j curves, we also shift the shares along the horizontal axis to align with the 150 percent FPL curve.

The cost of the marginal enrollee (C_H) can be straightforwardly derived from average costs (AC_H) and demand (D_H) shown in panel B of Figures 8 and 7,

²⁵Specifically, this assumes that the *slopes* of the average cost curves with respect to type, s , are stable with income even though the *levels* of costs are declining in income (see panel B of Figure 8).

respectively. The logic is identical to the two-plan case considered in past work (Einav, Finkelstein, and Cullen 2010). Online Appendix C provides more detail on the mechanics of constructing C_H .

Because we do not have variation in p_L and p_H that is orthogonal to $p_H - p_L$, we cannot use the same method to estimate AC_L and C_L . However, because the market share of L is relatively small (for example, just above 150 percent FPL, online Appendix Figure 24 shows that just 6 percent of the population enrolls in the L plan), the average L enrollee is similar to the marginal enrollee. We therefore use the average cost of the L plan for individuals just below and just above the 150 percent FPL threshold to approximate $C_L(s)$ for the range of s that these two sets of individuals span. Again, online Appendix C provides more detail.

C. Results and Implications for Take-Up

Figure 13 displays our key findings for individuals at 150 percent of FPL; online Appendix Table 7 summarizes the numbers in the figure at key points in the willingness to pay (s) distribution. Throughout the entire range of s spanning our data,²⁶ the $W_H(s)$ curve is substantially below both $AC_H(s)$ and $C_H(s)$. The gap between C_H and AC_H is sizable, indicating significant adverse selection, especially for lower-WTP types. For instance, if the highest-WTP half of the market ($s \leq 0.5$) enrolls in the H plan, the marginal enrollee (i.e., $s = 0.5$) costs $C_H(0.5) = \$333$ per month, about 20 percent less than the average enrollees' cost of $AC_H(0.5) = \$417$.

The fact that throughout the range of our data we find $C_H(s) > W_H(s)$ is particularly striking, since $W_H(s)$ and $C_H(s)$ represent enrollee WTP and insurer costs for the same people.²⁷ Throughout the observed distribution, $W_H(s)$ is less than half of $C_H(s)$. At the median of the WTP distribution, $W_H(0.5)$ is only \$103 per month, less than one-third of $C_H(0.5) = \$333$. Even at the highest in-sample point of the WTP distribution (the sixty-ninth percentile, or $s = 0.31$), $W_H(0.31)$ of \$162 per month is still substantially below average costs of insuring the top 31 percent of the WTP distribution ($AC_H(0.31) = \$448$), as well as costs for the sixty-ninth percentile WTP individuals ($C_H(0.31) = \$399$). Even if one could eliminate adverse selection and set premiums for marginal enrollees equal to their expected costs imposed on the insurer (i.e., $p_H(s) = C_H(s)$), at least 70 percent of individuals would not buy H .

WTP for L is also far below its costs, C_L . Indeed, the gap between W_L and C_L is larger than between W_H and C_H over the entire range that we can observe both sets of curves. Specifically, to the right of $s = 0.64$, C_L ranges from \$177 to \$241 per month whereas W_L is less than about \$50 per month. Indeed, the observed C_L points lie above our *maximum* in-sample W_L estimate of \$129 per month. Assuming adverse selection leads to a C_L curve that rises with W_L , C_L will also be above W_L for the 70 percent of the population for which we can measure demand for H or L .

²⁶ Our price variation spans roughly the seventieth to the sixth percentile of the WTP distribution (essentially all but the top 30 percent of the WTP distribution).

²⁷ Our willingness to pay and cost curves were adjusted to reflect those of the 150 percent FPL income group. If we had instead adjusted W_H and AC_H to line up with the estimates at 200 percent FPL or 250 percent FPL, we would still find C_H substantially above W_H .

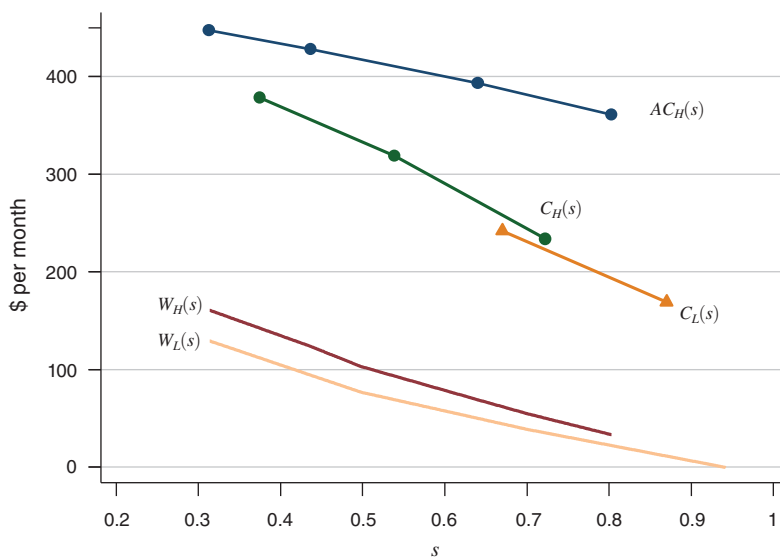


FIGURE 13. WILLINGNESS TO PAY AND COST

Notes: Figure reproduces AC_H , W_H , and W_L from panel B of Figures 11 and 12. The C_H curve is constructed using the formula in equation (6) in online Appendix C for each pair of AC_H points, with the x -value set as the midpoint between the two points x -values. The C_L curve is constructed using the approach described in the text.

The implied C_L curve is quite similar to the C_H curve over the regions of the s distribution where both are observed. This suggests that obtaining the more generous H contract instead of the L contract does not significantly increase costs. Therefore the much lower observed average cost in the L plan (see Table 1) is driven largely by favorable selection rather than by the causal impact of the plan on costs for the same type, s (i.e., moral hazard).

Take-Up under Counterfactual Subsidies.—These results imply low take-up of even heavily subsidized insurance for low-income adults. For example, at 150 percent of poverty, individuals in Massachusetts face a \$39 enrollee premium for the L plan, which is a 90 percent subsidy relative to the insurer price (see Figure 2, panel B). Our estimates of W_L and W_H indicate that, with a 90 percent price subsidy, only 69 percent of the market would enroll if offered L , and only 76 percent would enroll if offered H (with the corresponding enrollee monthly premium of \$42.40). These estimates have implications for understanding enrollment in the ACA subsidized exchanges, where enrollee premiums are significantly higher than in Massachusetts at a given income level. For instance, for an individual at 175 percent FPL in 2011, the ACA would make the second-cheapest silver plan cost 5.2 percent of income (or \$83 per month),²⁸ even though the ACA silver plan has an actuarial value below the

²⁸This calculation applies the ACA's subsidy rules. See <https://www.kff.org/health-reform/issue-brief/explaining-health-care-reform-questions-about-health/> (accessed February 11, 2019), which specify the premium of the second-cheapest silver plan as a percent of income, to the FPL for a single individual in 2011.

TABLE 2—IMPLICATIONS OF ALTERNATIVE SUBSIDIES FOR *H* PLAN

Subsidy		Enrollee premium	Share insured	Marginal enrollee		Average cost
Percent of AC	\$/month			WTP	Cost	
50%	\$246	\$246	3	\$246	\$488	\$492
	(\$25)	(\$25)	(9)	(\$25)	(\$41)	(\$50)
75%	\$314	\$105	49	\$105	\$335	\$418
	(\$18)	(\$6)	(3)	(\$6)	(\$34)	(\$25)
90%	\$328	\$36	79	\$37	\$208	\$364
	(\$16)	(\$2)	(1)	(\$2)	(\$113)	(\$18)
100%	\$329	\$0	96	\$0	\$139	\$329
	(\$34)	(\$0)	(1)	(\$0)	(\$131)	(\$34)

Notes: Table summarizes the implications of our estimates for enrollment and costs for *H* under alternative subsidies (shown in the rows), with bootstrapped standard errors shown in parentheses. We consider subsidies that correspond to 50 percent, 75 percent, 90 percent, and 100 percent of equilibrium average costs (column 1). These lead to enrollee premiums for *H* of \$246, \$105, \$36, and \$0 per month (column 3). The second column shows the corresponding dollar amount of the subsidy. The fourth column shows the share of the eligible population purchasing insurance (i.e., the offered *H* plan). The next two columns show WTP and costs of the marginal enrollee: the marginal WTP by definition equals the enrollee premium. The final column shows average costs of the insured population. Standard errors are calculated by a nonparametric bootstrap method that repeatedly resamples our original enrollee-level dataset, performs the RD regressions, and calculates the WTP and cost curves that generate these statistics. Note that while our analysis of 75 percent and 90 percent subsidies uses in-sample demand and cost curves, analysis of 100 percent and 50 percent subsidies requires extrapolating demand and costs outside of our in-sample range, which spans all but the top 31 percent and bottom 6 percent of the willingness to pay distribution. Online Appendix D shows the simple linear extrapolation we use to approximate willingness to pay and costs out of sample for these two estimates.

plans we study in Massachusetts.²⁹ Our estimates suggest that at an enrollee premium of \$83, only about half of the market would buy *L* and only about 60 percent would buy *H*.

The results also suggest that without subsidies that lower enrollee premiums substantially below average insurer cost of enrollees, relatively few low-income people would take up insurance. To illustrate this, Table 2 summarizes predicted take-up under potential subsidies for plan *H*. With enrollee premiums that are 75 percent below average costs (i.e., a subsidy in excess of \$300 per month) only 50 percent of the population would enroll. Premiums would need to be 90 percent below average costs in order to induce 80 percent enrollment. Interestingly, the per-enrollee subsidy cost increases by only \$2 as subsidies move from 90 percent to 100 percent of average cost, reflecting the fact that average enrollee costs are declining steeply as healthier individuals are brought into the market.

D. Sensitivity

Our key findings of low take-up and willingness to pay well below insurer costs are robust to a number of alternative implementation choices. Table 3 summarizes

²⁹We estimate that CommCare plans have an actuarial value of about 97 percent for enrollees between 100–200 percent of poverty. In the ACA, the baseline silver plan has an actuarial value of 87 percent for enrollees between 150–200 percent FPL.

TABLE 3—SENSITIVITY ANALYSIS: WTP AND COST ESTIMATES FOR *H* PLAN

Robustness specification	In-sample range of s	Median WTP ($s = 0.5$)			Share insured with subsidy (as percent of AC_H)	
		W_H	C_H	AC_H	75% subsidy	90% subsidy
Baseline estimates	[0.31, 0.94]	\$103	\$333	\$417	49	79
1. <i>Alternate RD specifications</i>						
Smaller bandwidth (25% FPL)	[0.29, 0.94]	\$100	\$318	\$418	48	78
Quadratic functional form	[0.28, 0.84]	\$98	\$351	\$403	49	73
Omit 200% FPL estimates	[0.30, 0.94]	\$97	\$343	\$412	46	79
2. <i>Alternate take-up estimates</i>						
Unscaled ACS eligible pop.	[0.19, 0.56]	\$24	\$186	\$354	29	47
3. <i>Accounting for mandate penalty</i>						
Use normalized premiums	[0.31, 0.94]	\$93	\$333	\$417	46	74

Notes: Top line (*Baseline estimates*) reproduces results from Figure 13 and Table 2. The remainder of the table shows analogous estimates with (row 1) alternate RD specifications, (row 2) alternate take-up estimates, and (row 3) premiums normalized by the expected mandate penalty. For each specification, we report the in-sample range of s values; estimates of W_H , AC_H , and C_H at the median of the WTP distribution ($s = 0.5$); and the share who purchase H under various subsidies as a percent of AC_H .

some of these results. Each row represents a single deviation from the baseline specification, as indicated. In all the alternative specifications we consider below, our main results continue to hold: W_H is substantially below AC_H and C_H , which implies limited take-up even with substantial subsidies. Indeed, the sensitivity analysis highlights the conservative nature of our baseline assumptions; under plausible alternative specifications, the share who enroll in H at a given subsidy level is always (weakly) lower, sometimes substantially so.

RD Specification.—Our baseline specification allowed a (linear) slope and intercept to vary on each side of the 150 percent, 200 percent, and 250 percent thresholds (see equation (1)). In practice, this meant a bandwidth of 50 percent of FPL everywhere but to the left of the 150 percent FPL threshold. The first two rows of Table 3 show results using a narrower (25 percent of FPL) bandwidth, and results with the baseline bandwidth but a quadratic (rather than linear) functional form for the running variable. The third row shows results from our baseline specification excluding one of our three thresholds, the 200 percent of FPL threshold.³⁰ As we discussed in Section I, this threshold is potentially problematic because of two other small changes at 200 percent FPL that could affect enrollment or costs independently of the change in enrollee premium: eligibility expands slightly at 200 percent FPL (to cover pregnant women and HIV-positive individuals for whom Medicaid eligibility ceases) and copays increase slightly at 200 percent FPL, resulting in a decline in plan actuarial value from 97 percent to 95 percent.

Examining Manipulation of the Running Variable.—A key threat to the validity of our empirical design is if individuals manipulate the running variable (the CommCare-specific income measure as a share of FPL) in order to qualify for

³⁰ We fill in the (now missing) space between the 150 percent and 250 percent FPL line segments by extrapolating the 150 percent FPL segment linearly.

higher subsidies. A standard way to look for such manipulation is to examine how the density of the population varies around the RD thresholds. Online Appendix Figure 14 shows the density of the eligible population in the ACS data for 2011, as well as for a similar ACS sample of CommCare-eligible individuals pooled over 2009–2013. The figure shows no evidence of bunching in the ACS income distribution at the thresholds. However, the data are somewhat noisy for 2011 alone, and more importantly, the income measure reported in the ACS may not be the same as what is reported to CommCare for purposes of determining subsidy eligibility. It is the CommCare-specific income measure for which we are concerned about potential manipulation, not what is reported to the ACS. Since our administrative data only include those who actually enroll in CommCare, we cannot look for bunching per se in the CommCare income measure as we cannot separately identify income manipulation from the take-up response to higher premiums.³¹

Turning to the administrative data on plan enrollment, we can look for other patterns consistent with strategic manipulation, including whether there is an upward spike in the number of enrollees just below the subsidy threshold or a decline just above the threshold (see Kleven 2016). Panel A of Figure 5 shows some slight evidence of lower enrollment (relative to the linear slope in income that we fit) to the right of the thresholds in the 2009–2013 pooled data. Online Appendix Figure 25 examines this further by showing enrollment by FPL separately for each year. The slightly lower-than-projected enrollment to the right of the subsidy threshold in the pooled 2009–2013 data appears to be entirely driven by the 2012 and (especially) 2013 data. There is no evidence of manipulation in the earlier years; as we explain in online Appendix J, there is an administrative rather than strategic explanation for the limited bunching that we see in 2012 and 2013. The lack of manipulation in 2011, our base year, suggests that any attempt to adjust for it using standard techniques (e.g., donut RDs as in Diamond and Persson 2016) would not substantively affect our baseline estimates.

Alternative Estimates of Take-Up.—As we discussed in Section I, the administrative data alone are sufficient to estimate willingness to pay, average cost, and own costs for the enrolled population and thus produce our key result that willingness to pay is substantially below average cost and own costs. However, an estimate of the eligible population size is essential for pinning down where in the distribution of willingness to pay for insurance our observed demand changes occur.

Our baseline estimates scaled the shape of the eligible population income distribution in the ACS to match the ACS estimate that on average, 63 percent of the eligible population enrolls in CommCare. As discussed, the ACS's coverage estimates match other survey estimates, as well as estimates based on tax filing data. However, if we instead divide the administrative counts of enrollment in CommCare

³¹ Recent work by Heim et al. (2016) and Kucko, Rinz, and Solow (2018) find evidence of modest income bunching in response to notches in the ACA subsidy schedule at, respectively, 400 percent of FPL (above which subsidies end) and 100 percent of FPL in non-Medicaid expansion states (below which people fall into a “coverage gap”). These notches are far larger than our \$40 per month amounts, e.g., the Kucko, Rinz, and Solow (2018) notch is about \$250 per month, and the responses are modest. For example, using the universe of IRS tax data, the Kucko, Rinz, and Solow (2018) paper suggests an excess mass of only about 20,000 tax filers and that is limited to the self-employed (i.e., no detectable response for wage earners).

by the raw ACS estimates of the size of the eligible population, we estimate only a 37 percent take-up rate. As shown in Table 3, this implies an even lower fraction of the population that will be insured at a given subsidy. For example, with this alternative take-up rate, we estimate that with a 75 percent subsidy of average costs, only 29 percent of eligible individuals would enroll in H , compared to 49 percent in our baseline analysis. Right below 150 percent FPL, using the raw ACS estimates for the denominator suggests 56 percent take-up, compared to our baseline estimate of 94 percent. The lower take-up estimates based on the ACS denominator may reflect the fact that income in the ACS is a noisy measure of the administratively recorded income in the CommCare data.

Accounting for the Mandate Penalty.—Our baseline analysis assumes that individuals do not take account of the expected mandate penalty for remaining uninsured when deciding whether to buy insurance. While we argue in Section I that this is a reasonable assumption in our institutional environment, the last row of Table 3 shows that accounting for the mandate penalty, which lowers the effective premiums, implies even lower willingness to pay than our baseline estimates.³² For example, our estimates now imply that with a 75 percent subsidy of average costs, only about 46 percent would enroll in H .

Inertia in Plan Demand and Robustness to New Enrollees.—Our estimates thus far have abstracted from inertia or switching costs, which have shown to be relevant for health insurance plan choices (Handel 2013, Ericson 2014, Polyakova 2016). If individuals do not make “active choices” each year once they are enrolled in CommCare this raises potential concerns with our estimates.

One concern is that enrollees’ income might change, leading them to move across the RD income threshold, but they might be unaware of the change or not re-optimize their choices.³³ This would suggest that our estimates understate the impact of higher premiums on insurance demand.

A second concern is that enrollees may not respond to changes in relative premiums for L compared to H , affecting our estimates of W_L relative to W_H . This seems potentially relevant, since the L plan (CeltiCare) was new to the market in 2010, so enrollees who entered the exchange prior to 2010 did not have it as an option when they first joined. However, our main findings are driven by a shift in demand from any formal insurance (H or L) to uninsurance at the RD income thresholds. Thus, switching between H and L is less likely to be empirically important for our main results. Further, because L was unavailable prior to 2010, lack of awareness of L would likely push upward our estimates of demand for H relative to L .

A third, and related, concern is that in years prior to 2011, the premiums for the different plans that make up the H composite plan varied.³⁴ This motivated our focus

³²Specifically, we normalize premiums by subtracting the expected mandate penalty values (shown by the black dots in Figure 3) from the “sticker premiums” shown in that same figure; note that effective premiums are everywhere lower, but the premium change at the FPL thresholds remains the same. Online Appendix Table 6 reports the normalized premiums by FPL.

³³Institutionally, lack of awareness seems less likely to be relevant, since the administrative income variable (used for calculating subsidies) changes only if an individual is audited (a salient event) or self-reports a change.

³⁴For instance, in 2010 for individuals in the 150–200 percent FPL group, enrollee premiums for the four H plans varied from \$39 to \$64 per month.

on 2011 when the premiums were quite similar, so these plans can be pooled into a single H option (defined as the preferred choice among the four component plans) with price p_H . However, if individuals made their choices in other years, this could be problematic.

To investigate the potential importance of such concerns for our estimates, we re-estimated demand on the subsample of new enrollees. We define new enrollees as those who enroll for the first time in 2011 (since the market opened in 2006); by definition, therefore, they must make an active choice in 2011. The results for new enrollees in 2011 are shown in online Appendix Figures 26 and 27 for enrollment in any plan and enrollment in the H plan, respectively. In each case, panel A shows results for all enrollees (for comparison) while panel B shows results for new enrollees. New enrollees comprise about one-sixth of all enrollee-months. The percent reductions in enrollment at the income discontinuities are similar for new enrollees and all enrollees; for example, enrollment in any plan declines by 22 to 28 percent for new enrollees, compared to 25 to 27 percent for all enrollees. This suggests that inertia is unlikely to be biasing downward our estimates of how much demand falls as premiums rise.

Relaxing Assumptions of Vertical Model.—Our analysis thus far has assumed a vertical model of demand with two CommCare options, H and L . In online Appendix E, we show that we can eliminate the vertical assumptions and still obtain *bounds* on WTP for CommCare (W^{Ins}), defined as an individual's WTP for her most preferred plan. Assuming only that consumers are optimizing when making their plan choices, we can use the enrollee premium for the cheapest plan (p^{min}) as a lower bound on W^{Ins} at a given point in its distribution, and the premium of the most expensive plan (p^{max}) as an upper bound. Our RD subsidy discontinuities then serve as exogenous variation in p^{min} and p^{max} that let us map out these lower and upper bounds on W^{Ins} across a range of the population distribution. We find that the resultant lower and upper bounds of W^{Ins} are, in fact, quite similar to the W_L and W_H estimates from the baseline vertical model. The lower bound on W^{Ins} is identical to W_L by construction: both are generated by plotting the share purchasing any insurance against the premium of the cheapest plan (L). The upper bound on W^{Ins} is also only slightly above W_H .

IV. Discussion and Normative Implications

Our finding that individuals are not willing to pay the costs they impose on the insurer suggests there are significant barriers to setting up private markets for low-income adults. Insurers cannot recoup their costs if individuals aren't willing to pay for them.

This raises a potential puzzle, since in the standard model of insurance, individuals are willing to pay their own expected costs *plus* a value of risk protection. One parsimonious explanation is that the relevant costs that drive an individual's demand for insurance is not the costs they impose on the insurer, which is what our cost curve estimates reflect, but rather the costs they would pay if they were uninsured. It is now well documented that uninsured (predominantly low-income) individuals do not pay their full medical costs when they receive medical care (see, e.g., Garthwaite,

Gross, and Notowidigdo 2015; Finkelstein et al. 2012; Mahoney 2015; Dobkin et al. 2016; Hu et al. 2016; Brevoort, Grodzicki, and Hackmann 2017). While there is considerable uncertainty in the exact prevalence of uncompensated care, national estimates suggest that the uninsured pay only about 20 percent to 35 percent of their cost of care (Coughlin et al. 2014; Hadley et al. 2008; Finkelstein, Hendren, and Luttmer 2015). This is remarkably similar to our estimated ratio of WTP to own costs imposed on the insurer for the *H* plan. Thus when we compare individual willingness to pay not to the cost of the insurer, but rather to the estimates of what they would pay if they were uninsured, we find that willingness to pay is quite similar to the costs they would incur if they were uninsured; in online Appendix F, we walk through this calibration exercise in more detail.

Of course, we do not directly observe the amount paid out of pocket by the uninsured in our sample. Thus, there is considerable uncertainty around whether our calibration of the individuals' own expected costs if uninsured is above or below WTP. Nor are we able to directly estimate the impact of access to uncompensated care on willingness to pay for insurance in our setting. This caveat is particularly important given the large literature suggesting that health insurance choices may be rife with behavioral biases.³⁵ This raises the possibility that consumers might not respond "rationally" to changes in uncompensated care; indeed, as suggested by Fadlon and Laibson (2017), the current uncompensated care system could be a response by a rational planner to anticipated behavioral biases that would lead individuals to not purchase insurance even in the absence of uncompensated care. For example, there is substantial evidence of individuals underestimating the probability of negative events (see, e.g., Moore and Healy 2008 for a review); Spinnewijn (2015) estimates that unemployed individuals underestimate the expected duration of unemployment by over 70 percent. In our context, such underestimation of expected costs could easily close the approximately 200 percent gap between willingness to pay and insurer costs, and would still depress demand even in the absence of uncompensated care.

Both behavioral biases and access to uncompensated care may play a larger role for lower income populations. Behavioral biases may be particularly acute among low-income populations who are may be making purchase decisions under greater constraints or stress (Mani et al. 2013; Mullainathan and Shafir 2014; Bhargava, Loewenstein, and Sydnor 2017). Lower-income individuals also have greater access to uncompensated care than higher income individuals (Mahoney 2015; Dranove, Garthwaite, and Ody 2015). In online Appendix G we show that greater access to uncompensated care (as measured by proximity to safety net providers) is associated with a greater gap between willingness to pay and insurer costs. We also show in online Appendix G that within our own sample, the gap between willingness to pay and insurer costs narrows as we move from the 150 percent FPL threshold to the 200 percent threshold to the 250 percent threshold. Out of sample, our

³⁵Ericson and Sydnor (2017) provide a recent overview. A partial list of examples of behavioral biases that have been documented for health insurance choices includes inertia (Handel 2013), confusion about contract dimensions (e.g., Bhargava, Loewenstein, and Sydnor 2017; Handel and Kolstad 2015), over-weighting certain financial plan features relative to others (e.g., Abaluck and Gruber 2011), and making suboptimal choices (e.g., Abaluck and Gruber 2011, Kling et al. 2012, Handel 2013). Our previous finding of similar (or if anything lower) demand responsiveness by new enrollees suggests that inertia or inattention are not primary drivers of low willingness to pay, but naturally we cannot rule out behavioral biases more broadly.

finding of willingness to pay below insurer costs for the low-income population in Massachusetts contrasts with Hackmann, Kolstad, and Kowalski's (2015) estimate that higher-income individuals in Massachusetts (above 300 percent of FPL) are willing to pay the cost they impose on the insurer.

We are unable to quantify the relative roles of behavioral biases and uncompensated care in reducing willingness to pay so far below insurer costs in our setting. However, we believe we can rule out several other potential explanations for our finding. One is moral hazard: some of the cost incurred may not be fully valued by the individuals because they only consume the care if they don't have to pay for it. But to close the gap in our setting, insurance would have to increase costs by a factor of at least 200 percent, which is an order of magnitude larger than most plausible estimates of the impact of moral hazard on utilization.³⁶ Another possibility, recently emphasized by Hendren (2017), is that observed demand may understate the ex ante value of insurance measured before the individual has learned something about her health type; WTP might be higher if the individual considered a purchase decision prior to learning their health type. In practice, however, as we show in online Appendix I, our estimates suggest that even though an ex ante willingness to pay measure may be significantly higher than our baseline willingness to pay estimates, it would still be below own cost.

Yet another potential factor creating a wedge between willingness to pay and costs that we suspect is not quantitatively important in our setting are liquidity constraints, i.e., the inability to borrow against future income at market rates of interest. As Casaburi and Willis (2016) observe, most insurance products require individuals to pay the premium up front, thus transferring income across time as well as states. We suspect, however, that liquidity constraints are unlikely the primary driver of low willingness to pay in our context. One reason is that the majority of marginal enrollees choose to pay for the *H* instead of the *L* plan, suggesting that although they might be liquidity constrained, they are not up against the corner of their budget constraint. In addition, the premiums we study represent only about 0–5 percent of family income. Of course, the fact that individuals are low income, and therefore high marginal utility of consumption, may reduce their willingness to pay, but that is a separate point from liquidity constraints and one we return to below when we consider normative implications.

Normative Implications.—Thus far, we have focused on positive analysis of demand for insurance among a low-income population, and its implications for how insurance take-up would vary under alternative subsidies. Normative analysis faces (at least) two additional hurdles. First, we must be willing to accept our estimates of willingness to pay as the welfare-relevant metric. The presence of behavioral biases, which, as noted, have been extensively described in health insurance choices, raises concerns with this assumption. So too does Hendren's (2017) point that observed

³⁶For example, in online Appendix H we translate the estimates from Chandra, Gruber, and McKnight (2014) on the impact on health care spending of more versus less generous plans in our Massachusetts CommCare population. This suggests that relative to being uninsured, insurance coverage would increase spending by roughly 15–25 percent. This is broadly consistent with the results from the Oregon Health Insurance Experiment, which found that Medicaid for low-income adults, a close analog to CommCare, increases health care spending by about 25 percent relative to being uninsured (Finkelstein et al. 2012).

demand may understate the *ex ante* value of insurance measured before the individual has learned something about her health type.

Second, even were we to accept estimated willingness to pay as the welfare-relevant metric, normative analysis needs to consider the fact that the subsidy recipients are a low-income population (perhaps with a high marginal utility of consumption), which presumably reduces their willingness to pay for any good, including health insurance. One approach would be to apply a social welfare function that takes this into account by using a parameterization of social marginal utilities of income that translate individual willingness to pay into social willingness to pay (Saez and Stantcheva 2016). In other words, even if recipient willingness to pay does not exceed costs, social willingness to pay may exceed cost. Another approach would be to follow Hendren (2016) and compute the marginal value of public funds (MVPF), i.e., the ratio of marginal benefit to marginal cost, for an incremental government subsidy for health insurance. This could then be compared to the MVPF of alternative redistributive programs to a low-income population, such as cash transfers through the Earned Income Tax Credit (EITC). For this normative approach, liquidity constraints are not relevant; if demand for health insurance is low in part because individuals have a high marginal value of current cash due to liquidity constraints, that will (and should) increase the MVPF of cash transfers relative to in-kind subsidies.

If the existence of substantial uncompensated care for the low-income uninsured is a primary driver of low willingness to pay for formal insurance, a crucial question concerns the ultimate economic incidence of the uncompensated care that is provided in the absence of formal insurance; is this, for example, on the government, the low-income uninsured themselves, or more affluent third parties? A large role for uncompensated care in explaining why willingness to pay is substantially below (gross) insurance costs also suggests a potential efficiency, rather than purely redistributive, rationale for subsidies, as an offset to the implicit tax that uncompensated care imposes on formal insurance (Coate 1995). There are also other possible distortions created by the current system of charity care provision, such as potential distortions in providers' decisions of whether to locate or expand certain services in low income areas (e.g., Dranove, Garthwaite, and Ody 2015).

V. Conclusion

This paper estimates willingness to pay and costs for health insurance among low-income adults using data from Massachusetts' pioneer subsidized insurance exchange. For at least 70 percent of the low-income eligible population, we find that willingness to pay for insurance is far below insurers' average costs. From a positive economics perspective, our results point to substantial challenges in getting to universal coverage via partially subsidized insurance programs like the ACA's exchanges. For example, we estimate that even subsidizing premiums down to 10 percent of insurer costs would generate only 80 percent coverage. This reality may underlie the incomplete take-up of insurance under the ACA, despite a coverage mandate and generous subsidies.

We find evidence of adverse selection, but show that, by itself, adverse selection cannot explain limited demand; we estimate that at least 70 percent of the eligible population would not enroll even if prices were subsidized to reflect own expected

costs. The magnitude of the gap between willingness to pay and own costs is also substantially larger than what could plausibly be explained by moral hazard effects of insurance. Of course, expected costs reflect costs to the insurer, not necessarily costs the individual would pay if uninsured. Adjusting our cost estimates for existing estimates of the magnitude of available uncompensated care for the uninsured, we find that WTP is roughly close to these costs. Other analyses using a calibrated utility model (rather than revealed preference) for welfare analysis of health insurance for low-income adults similarly finds willingness to pay that is substantially below gross costs imposed on the insurer and quite close to the net costs of insurance that account for the uncompensated care that would be provided if the individual remained uninsured (Finkelstein, Hendren, and Luttmer 2015).

The large literature on behavioral biases in health insurance purchase decisions suggests that such biases could also play a large role in reducing estimated of willingness to pay. Crucially, they also suggest caution in normative analysis of health insurance subsidies based on our demand estimates. While it is typical to use revealed preference in welfare analysis of demand for employer-provided health insurance (see, e.g., Einav, Finkelstein, and Levin 2010 for a review), there may be reluctance to do so in our context.

Normative analysis should also consider that the subsidy recipient population is poor, which provides a natural potential redistributive rationale for policy. This does not, however, speak directly to the relative merits of providing in-kind subsidies for health insurance as opposed to cash transfers, such as through an expansion of the Earned Income Tax Credit. Given the existence of substantial uncompensated care provision to the low-income uninsured, any redistributive analysis of subsidies to low-income individuals for health insurance must tackle the important, but challenging, question of the economic incidence of this uncompensated care provision. This is an important direction for future work.

REFERENCES

- Abaluck, Jason, and Jonathan Gruber.** 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101 (4): 1180–210.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor.** 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Options." *Quarterly Journal of Economics* 132 (3): 1319–72.
- Brevoort, Kenneth, Daniel Grodzicki, and Martin B. Hackmann.** 2017. "Medicaid and Financial Health." NBER Working Paper 24002.
- Brooks, Tricia, Karina Wagnerman, Samantha Artiga, Elizabeth Cornachione, and Petry Ubri.** 2017. *Medicaid and CHIP Eligibility, Enrollment, Renewal, and Cost Sharing Policies as of January 2017: Findings from a 50-State Survey*. San Francisco: Kaiser Family Foundation.
- Buchanan, James M.** 1975. "The Samaritan's Dilemma." In *Altruism, Morality, and Economic Theory*, edited by Edmund S. Phelps. New York: Russell Sage Foundation.
- Casaburi, Lorenzo, and Jack Willis.** 2016. "Time vs. State in Insurance: Experimental Evidence from Contract Farming in Kenya." Unpublished.
- Chan, David, and Jonathan Gruber.** 2010. "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review* 100 (2): 292–96.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight.** 2011. "The Importance of the Individual Mandate: Evidence from Massachusetts." *New England Journal of Medicine* 364 (4): 293–95.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight.** 2014. "The Impact of Patient Cost-Sharing on Low-Income Populations: Evidence from Massachusetts." *Journal of Health Economics* 33: 57–66.

- Coate, Stephen. 1995. "Altruism, the Samaritan's Dilemma, and Government Transfer Policy." *American Economic Review* 85 (1): 46–57.
- Cooper, Philip F., and Barbara Steinberg Schone. 1997. "More Offers, Fewer Takers for Employment-Based Health Insurance: 1987 and 1996." *Health Affairs* 16 (6): 142–49.
- Coughlin, Teresa A., John Holahan, Kyle Caswell, and Megan McGrath. 2014. *Uncompensated Care for the Uninsured in 2013: A Detailed Examination*. San Francisco: Kaiser Family Foundation.
- Dague, Laura. 2014. "The Effect of Medicaid Premiums on Enrollment: A Regression Discontinuity Approach." *Journal of Health Economics* 37: 1–12.
- De Nardi, Mariacristina, Eric French, and John Bailey Jones. 2016. "Medicaid Insurance in Old Age." *American Economic Review* 106 (11): 3480–520.
- Diamond, Rebecca, and Petra Persson. 2016. "The Long-Term Consequences of Teacher Discretion in Grading of High-Stakes Tests." NBER Working Paper 22207.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2016. "The Economic Consequences of Hospital Admissions." NBER Working Paper 22288.
- Dranove, David, Craig Garthwaite, and Christopher Ody. 2015. *A Floor-and-Trade Proposal to Improve the Delivery of Charity-Care Services by US Nonprofit Hospitals*. Washington, DC: Brookings Institution.
- Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives* 25 (1): 115–38.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877–921.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics* 2: 311–36.
- Ericson, Keith M. Marzilli. 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy* 6 (1): 38–64.
- Ericson, Keith M. Marzilli, and Amanda Starc. 2015. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics* 97 (3): 667–82.
- Ericson, Keith M. Marzilli, and Justin Sydnor. 2017. "The Questionable Value of Having a Choice of Levels of Health Insurance Coverage." *Journal of Economic Perspectives* 31 (4): 51–72.
- Fadlon, Iitzik, and David Laibson. 2017. "Paternalism and Pseudo-Rationality." NBER Working Paper 23620.
- Farber, Henry S., and Helen Levy. 2000. "Recent Trends in Employer-Sponsored Health Insurance Coverage: Are Bad Jobs Getting Worse?" *Journal of Health Economics* 19 (1): 93–119.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." NBER Working Paper 21308.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20171455>.
- Finkelstein, Amy, and James Poterba. 2014. "Testing for Asymmetric Information Using 'Unused Observables' in Insurance Markets: Evidence from the U.K. Annuity Market." *Journal of Risk and Insurance* 81 (4): 709–34.
- Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group. 2012. "The Oregon Health Insurance Experiment: Evidence from the First Year." *Quarterly Journal of Economics* 127 (3): 1057–106.
- Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers. 2017. "Premium Subsidies, the Mandate, and Medicaid Expansion: Coverage Effects of the Affordable Care Act." *Journal of Health Economics* 53: 72–86.
- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo. 2015. "Hospitals as Insurers of Last Resort." NBER Working Paper 21290.
- Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review* 105 (3): 1030–66.
- Hadley, Jack, John Holahan, Teresa A. Coughlin, and Dawn Miller. 2008. "Covering the Uninsured in 2008: Current Costs, Sources of Payment, and Incremental Costs." *Health Affairs* 27 (5): w399–415.
- Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.
- Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–500.

- Heim, Bradley T., Gillian Hunter, Adam Isen, Ithai Z. Lurie, and Shanthi P. Ramnath. 2016. "Income Responses to the Affordable Care Act: Evidence from the Premium Tax Credit Notch." Unpublished.
- Hendren, Nathaniel. 2016. "The Policy Elasticity." In *Tax Policy and the Economy*, Vol. 30, edited by Jeffrey R. Brown, 51–89. Chicago: University of Chicago Press.
- Hendren, Nathaniel. 2017. "Measuring Ex-Ante Welfare in Insurance Markets." Unpublished.
- Hu, Luoia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong. 2016. "The Effect of the Patient Protection and Affordable Care Act Medicaid Expansions on Financial Wellbeing." NBER Working Paper 22170.
- Jaffe, Sonia, and Mark Shepard. 2017. "Price-Linked Subsidies and Health Insurance Markups." Unpublished.
- Janicki, Hubert. 2013. *Employment-Based Health Insurance: 2010*. Washington, DC: US Census Bureau.
- Kaiser Family Foundation. 2016. "Marketplace Enrollment as a Share of the Potential Marketplace Population." State Health Facts. <https://www.kff.org/health-reform/state-indicator/marketplace-enrollment-as-a-share-of-the-potential-marketplace-population-2015>.
- Kleven, Henrik J. 2016. "Bunching." *Annual Review of Economics* 8: 435–64.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. 2012. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics* 127 (1): 199–235.
- Krueger, Alan B., and Ilyana Kuziemko. 2013. "The Demand for Health Insurance among Uninsured Americans: Results of a Survey Experiment and Implications for Policy." *Journal of Health Economics* 32 (5): 780–93.
- Kucko, Kavan, Kevin Rinz, and Benjamin Solow. 2018. "Labor Market Effects of the Affordable Care Act: Evidence from a Tax Notch." Unpublished.
- Long, Sharon K., Karen Stockley, and Heather Dahlen. 2012. *Health Reform in Massachusetts as of Fall 2010: Getting Ready for the Affordable Care Act and Addressing Affordability*. Boston: Blue Cross Blue Shield Foundation of Massachusetts.
- Mahoney, Neale. 2015. "Bankruptcy as Implicit Health Insurance." *American Economic Review* 105 (2): 710–46.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science* 341 (6149): 976–80.
- Massachusetts Health Connector and Department of Revenue. 2011. Data on the Individual Mandate: Tax Year 2011. Technical Report.
- Moore, Don A., and Paul J. Healy. 2008. "The Trouble with Overconfidence." *Psychological Review* 115 (2): 502–17.
- Mullainathan, Sendhil, and Eldar Shafir. 2014. *Scarcity: The New Science of Having Less and How It Defines Our Lives*. New York: Picador.
- Polyakova, Maria. 2016. "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics* 8 (3): 165–95.
- Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." *American Economic Review* 106 (1): 24–45.
- Smith, Vernon K., Kathleen Gifford, Eileen Ellis, Robin Rudowitz, Laura Snyder, and Elizabeth Hinton. 2015. *Medicaid Reforms to Expand Coverage, Control Costs and Improve Care: Results from a 50-State Medicaid Budget Survey for State Fiscal Years 2015 and 2016*. San Francisco: Kaiser Family Foundation.
- Spinnewijn, Johannes. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of the European Economic Association* 13 (1): 130–67.
- Tebaldi, Pietro. 2017. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." Becker Friedman Institute for Research in Economics Health Economics Series Working Paper 2017–05.
- Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.
- US Department of Agriculture. 2016. *The Food Assistance Landscape: FY 2015 Annual Report*. Washington, DC: US Department of Agriculture.
- US Department of Health and Human Services. 2015. *2015 Actuarial Report on the Financial Outlook for Medicaid*. Washington, DC: US Department of Health and Human Services.
- US Internal Revenue Service. 2015. *EITC Calendar Year Report*. Washington, DC: US Internal Revenue Service.