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

Amy Finkelstein, Nathaniel Hendren, Mark Shepard*

July 2017

Abstract

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% for each $40 increase in monthly enrollee premiums. Marginal enrollees tend to be lower-cost, consistent with adverse selection into insurance. But across the entire distribution we can observe – approximately the bottom 70% of the willingness to pay distribution – enrollee willingness to pay is always less than half of *own* expected costs. As a result, we estimate that take-up will be highly incomplete even with generous subsidies: if enrollee premiums were 25% of insurers' average costs, at most half of potential enrollees would buy insurance; even premiums subsidized to 10% of average costs would still leave at least 20% uninsured. We suggest an important role for uncompensated care for the uninsured in explaining these findings and explore normative implications.

JEL codes: H51; I13.

Keywords: Health insurance; subsidies; low-income.

*Finkelstein: Department of Economics, MIT and NBER; afink@mit.edu. Hendren: Department of Economics, Harvard and NBER; nhendren@fas.harvard.edu. Shepard: Harvard Kennedy School and NBER; Mark_Shepard@hks.harvard.edu. 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 discussant Jeff Clemens for thoughtful and constructive comments. We also thank Abhijit Banerjee, Amitabh Chandra, Keith Ericson, Josh Goodman, Jon Gruber, Jeff Liebman, Brigitte Madrian, Neale Mahoney, and seminar participants at Harvard Kennedy School, the AEA Annual Meetings, University of Texas, MIT, and NBER Summer Institute 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 No. T32-AG000186 (via the National Bureau of Economic Research).

1 Introduction

Governments spend an enormous amount of money on health insurance for low-income individuals. For instance, the U.S. Medicaid program (at $550 billion in 2015) dwarfs the size of the next largest means-tested programs – food stamps and the EITC ($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 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% 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%, 200% and 250% 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%, 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%) with each premium increase.

[1] See U.S. Department of Health and Human Services (2015), U.S. Department of Agriculture (2016), and U.S. 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 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 about willingness to pay and costs for these plans. 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% of insurers' average costs, at most half the population would purchase insurance if offered H. Subsidies making enrollee premiums 10% of insurers' average costs still leave at least 20% 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% of income for the benchmark plan in the ACA versus 0-5% 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 (2016); Avalere Health (2016)). They can also generate predictions of coverage under alternate reform proposals.

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, implying that plans are adversely selected (Einav, Finkelstein, and Cullen, 2010). But across the in-sample distribution – which spans the 6th to the 70th 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. This suggests individuals would not enroll even if prices were subsidized to reflect the expected cost of marginal enrollees. Enrollment is low not simply because of adverse selection, but because people are not willing to pay their own cost imposed on the insurer.

In the final section of the paper, we briefly explore potential explanations for our findings and analyze their normative implications. Back-of-the-envelope calculations using estimates from the prior literature suggest that moral hazard effects of health insurance cannot explain much of the gap between willingness to pay and costs. However, an estimate of the uncompensated care available to the uninsured suggests that this can account for nearly all of our estimated gap between willingness to pay and costs of insurance.

[2]These 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.

If we interpret our willingness to pay estimates normatively as individuals' value of insurance, they suggest that most low-income enrollees would prefer being uninsured to having to pay the cost they impose on the insurer. This suggests that subsidies in this market cannot be justified simply as a response to adverse selection. Our results suggest two potential such justifications for subsidies: 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 (see e.g., Einav, Finkelstein, and Levin (2010) for an overview), there is relatively little work providing such estimates for low-income adults on whom much public policy attention is focused. This is partly because, until recently, most of the low-income uninsured either were not offered health insurance or faced prices that were difficult to measure.[3] The introduction of the ACA in 2014 has given researchers an opportunity to study how low-income insurance take-up responds to subsidies (Frean, Gruber, and Sommers, 2017; Sacks, 2017). However, the ACA's subsidy schedule does not feature the sharp discontinuities present in Massachusetts, which we exploit for our research design.

Our estimates of a low-income population's willingness to pay for formal insurance compared to insurance costs are consistent with Finkelstein, Hendren, and Luttmer (2015), who use *ex-post* outcome distributions to infer low-income adults' WTP for Medicaid in the context of the Oregon Health Insurance Experiment. Both sets of results suggest that low-income enrollees value formal health insurance products at substantially below their average cost. Such findings contrast with those for higher-income populations. In particular, Hackmann, Kolstad, and Kowalski (2015) study the *unsubsidized* Massachusetts health insurance exchange for individuals *above* 300% of poverty. They also find evidence of adverse selection but (unlike our estimates for a low-income population) estimate that willingness to pay *exceeds* own costs over the entire population of potential consumers. One natural explanation for these ostensibly 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)).[4]

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 mis-pricing of contracts is relatively small. This literature however has "looked under the lamppost" – primarily focusing on

[3] In a creative effort to surmount this substantial obstacle, Krueger and Kuziemko (2013) conducted a survey experiment designed to elicit willingness to pay for hypothetical plan offerings among a large sample of of the uninsured. In a recent exception-that-proves-the-rule, Dague (2014) examines how duration enrolled in Medicaid responds to increases in monthly premiums. In addition, another body of work has studied low-income individuals' price sensitivity in their choice of health plans, both in Massachusetts (Chan and Gruber, 2010) and in the ACA exchanges (Tebaldi, 2016), though without formally estimating willingness to pay.

[4] Differences 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 et al., 2016); unlike low-income adults, low-income elderly do not have access to substantial uncompensated nursing home care (the primary healthcare covered by Medicaid), either in the De Nardi et al. (2016) model or in practice.

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 2 presents the setting and data. Section 3 presents the basic descriptive empirical evidence, documenting the level and responsiveness to price of both insurance demand and average insurer costs. Section 4 uses a simple model of insurance demand to translate the empirical results from Section 3 into estimates of willingness to pay and costs for insurance. Section 5 briefly considers potential explanations for low willingness to pay and our findings' normative implications for insurance subsidies for low-income adults. The final section concludes.

2 Setting and Data

2.1 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% 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,[5] the CommCare-eligible population consisted of adults aged 19-64 without access to employer coverage and who were either (1) childless and below 300% of FPL, (2) non-pregnant parents between 133% and 300% of FPL, or (3) pregnant women between 200% and 300% of FPL.

CommCare specified a required set of benefits (i.e., covered services and cost sharing rules) and then solicited competing private insurers to provide these benefits. Each insurer offered a single plan that had the required set of benefits but could differ in its network of hospitals and doctors. Between

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

4 and 5 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 co-pays.[6]

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 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.

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) 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.

Subsidy Structure

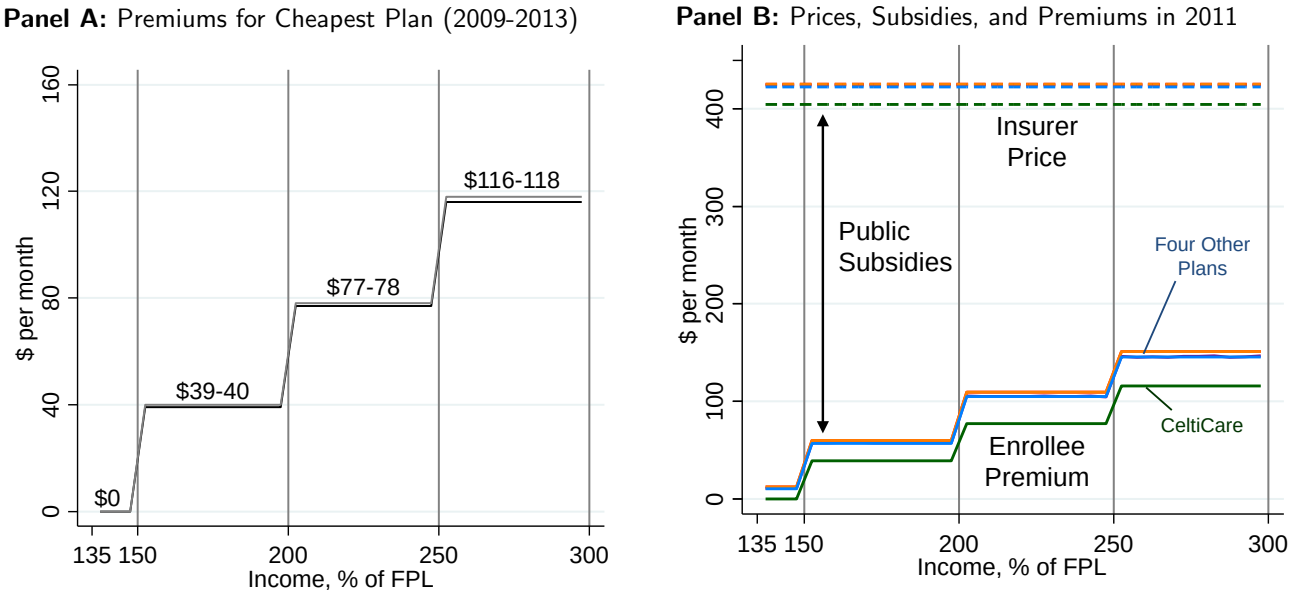
Insurers in CommCare set a base plan price that applied to all individuals, regardless of income (or age, region, or other factors). The actual payment the insurer received from CommCare equaled their base price times a risk score intended to capture predictable differences in health status.

Enrollees paid premiums equal to their insurer's base price minus an income-varying subsidy paid by the state.[7] 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%, 200%, and 250% of FPL. Figure 1, Panel A 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% of FPL and increases to $39 per month above 150% FPL, $77 per month above 200% FPL, and $116 per month above 250% 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% of income (specifically, 2.9% of income just above 150% FPL, 4.2% just above 200% FPL, and 5.1% just above 250% FPL).

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

[7]We 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.

Figure 1: Insurer Prices and Enrollee Premiums in CommCare Market



NOTE: Panel A plots enrollee premiums for the cheapest plan by income as a percent of FPL, noting the thresholds (150%, 200%, and 250% 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.

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. Figure 1, Panel B shows these 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 CeltiCare's premium is noticeably lower.

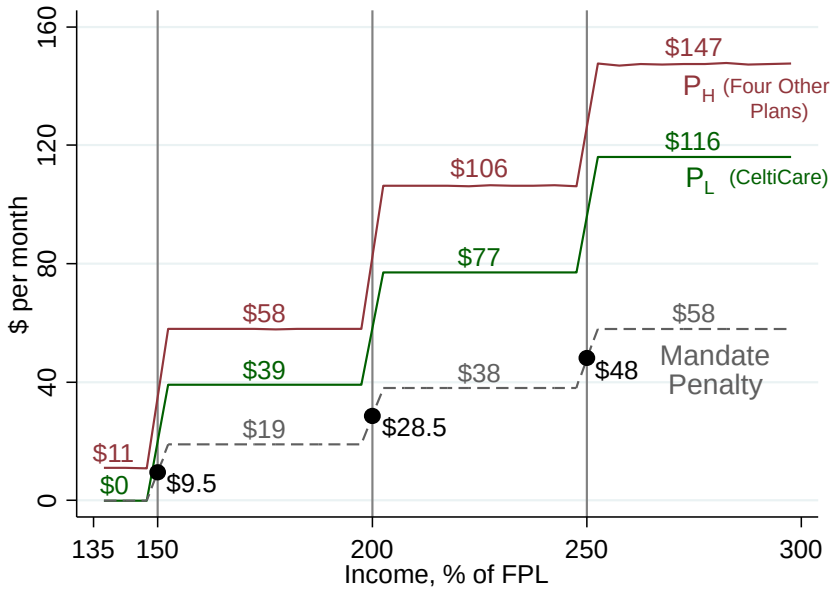
Along with its lower price, CeltiCare also had had a more limited network than other plans. We estimate that CeltiCare covered 42% of Massachusetts hospitals (weighted by bed size), compared to 79% or higher for the other three plans offered statewide.[8] Both because of this limited network and because of its lack of long-term reputation with consumers (it had entered the state insurance market only in 2010), CeltiCare was perceived by enrollees as less desirable, aside from its lower price.[9]

As a result, in much of our analysis that follows we pool the 2011 plans into two groups: CeltiCare

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

[9] Consistent with this perception, when all plans were available for free – which was the case for enrollees below 100% of FPL – 96% of enrollees chose a plan other than CeltiCare.

Figure 2: Premium and Mandate Penalty Variation, 2011



NOTE: 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.

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 4, we will further assume that H is perceived as higher quality than L. We also show in an extension in Section 4.4 below that we can generate fairly tight bounds on willingness to pay in a more general model that does not assume this vertical structure.

Figure 2 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; 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%, 200%, and 250% 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% FPL, to \$19 from 150-200% FPL, to \$29 from 200-250% FPL, and to \$31 above 250% FPL.

The final relevant option for people eligible for CommCare was to remain uninsured and pay the mandate penalty. The dotted gray line in Figure 2 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* 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 2 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.[10]

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 4.4 below, we also report the lower willingness to pay estimates that we find when we normalize premiums by the expected mandate penalty values (shown in black dots).

2.2 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 is at the individual level because CommCare only offers individual (not family) coverage.[11]

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 1 and 2). 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% level), which is the running variable for our RD analysis. This variable is

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

[11]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: CommCare Summary Statistics in 2011: Premiums, Enrollment and Costs

	Any Plan (1)	**H plan** (2)	**L Plan** (3)
Enrollment			
N (# 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%
Medical Claims (Monthly)			
Insurer-Paid	$358.5	$377.3	$197.9
Total	$377.3	$396.4	$213.4
Prices, Subsidies and Premiums (Monthly)			
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 show summary statistics from CommCare administrative data for fiscal year 2011 for enrollees between 135 and 300 percent of FPL.

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% of FPL because of the significant eligibility change at 133% 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 2). 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% 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% random sample of U.S. households. We use this data to estimate the number of people eligible for

CommCare in each 1% FPL bin between 135% of FPL and 300% FPL. To be coded as eligible, an individual must live in Massachusetts and be: a U.S. 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).[12] Because the ACS is a 1% sample (and because of clustering in reported income at round numbers), our raw estimates of the size of the eligible population by 1% FPL bin are relatively noisy. We therefore smooth the estimates using a regression of raw counts by 1% FPL bin on a polynomial in income. Appendix A 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% of eligible individuals enroll in CommCare. This number seems low compared to the take-up estimate in the ACS data itself, where we find that 63% of eligible individuals report having insurance. We suspect the ACS take-up estimate is closer to the truth since it closely matches estimates from a MA health insurance survey in the fall of 2010 (Long, Stockley, and Dahlen, 2012) and estimates from tax filing data.[13] We conservatively 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.

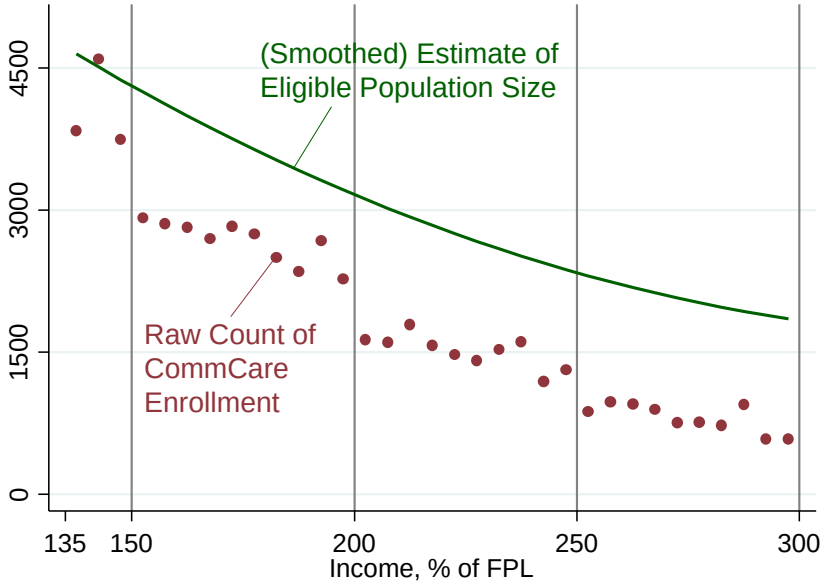
Specifically, we take our estimates of the number of eligible individuals in the ACS by FPL bin (see Panel A of Appendix Figure 15) and scale the whole distribution down by a constant multiple (of 0.59) so that dividing the administrative count of enrollees by our adjusted eligible population size yields an average take-up rate of 63% (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.

[12] The ACS does not distinguish Medicaid from CommCare coverage (both are coded as "Medicaid/other public insurance"), so we cannot directly exclude Medicaid enrollees.

[13] Specifically, Long et al. (2012) find a 90% insurance coverage rate for non-elderly adults below 300% of FPL, nearly identical to the 89% 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% of the FPL were uninsured at a typical point in time during 2011 (Connector and 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 midpoint 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% 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 3: Eligible and Enrolled Population, 2011



NOTE: 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% of the FPL.

Figure 3 shows our (smoothed) estimate of the size of the eligible population by FPL bin. For comparison we also show 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. 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).

3 Descriptive Analysis

3.1 Regression Discontinuity Design

We use the discrete change in enrollee premiums at 150%, 200% and 250% 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% of FPL level:

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

where Y_b is an outcome measure in that income bin b, Inc_b is income (as a % of FPL) at the midpoint of the bin, and $s(b)$ is the income segment on which bin b lies (either 135-150%, 150-200%, 200-250%, or 250-300% 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%, 200%, and 250% FPL). This would be violated if people strategically adjust (or misreport)[14] their income to get just below the thresholds and qualify for a larger subsidy.[15] While in principle such manipulation would be possible, in our setting individuals were largely shrouded during the application process from a clear illustration of how their reported incomes translated into the %FPL formula for determining subsidies. As we illustrate below in Section 4.4, we find minimal evidence of such manipulation. 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.

3.2 Evidence from Pooled Years (2009-2013)

Insurance Demand

Figure 1 showed that premiums increase discontinuously at 150, 200, and 250% of FPL. Figure 4, Panel A 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% to 40%. All three changes are statistically distinguishable from zero ($p < 0.01$).

Cost of Insurance and Adverse Selection

Figure 4, 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.

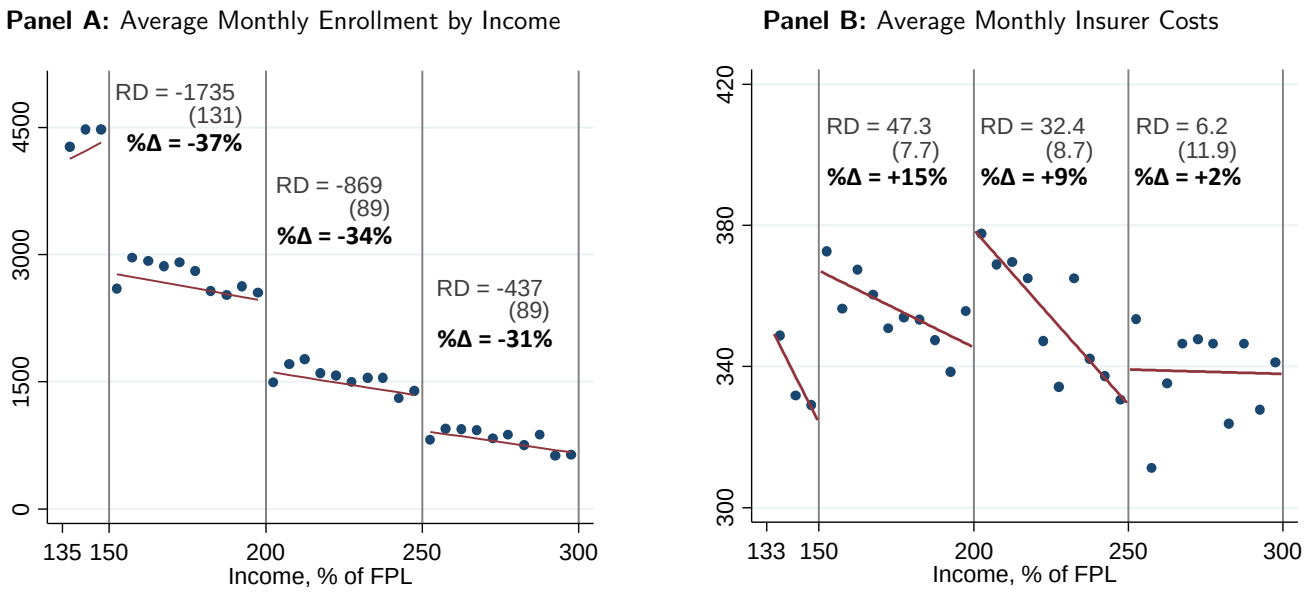
The figure shows that average costs of the insured rise as the enrollee premium increases. For example, we estimate a discontinuous increase in costs of $47.3 (s.e. $7.7) per enrollee per month at 150% FPL and of $32.4 (s.e. $8.7) at 200% of FPL. We find a smaller but noisier increase of $6.2 (s.e. $11.9) at 250% FPL; this imprecision may reflect the smaller size of the eligible and enrolled populations at 250% of FPL (see Figure 3).

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. In Appendix Figure 19 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

[14] Enrollees 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.

[15] In addition, there are minor changes in eligibility just above 200% 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 towards zero, since the eligible population grows just above 200% FPL. In sensitivity analysis, we show that our main results are robust to excluding this discontinuity.

Figure 4: CommCare Enrollment and Average Insurer Costs, 2009-2013



NOTE: The figure shows discontinuities in enrollment and average insurer costs at the income thresholds (150%, 200%, and 250% of FPL) at which enrollee premiums increase (see Figure 1). Panel A shows average enrollment in CommCare (total member-months, divided by number of months) by income over the 2009-2013 period our data spans. 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% 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 "%Δ =".

thresholds. In other words, in response to higher premiums, younger and lower-risk enrollees are more likely to leave the market.

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 pre-existing 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.

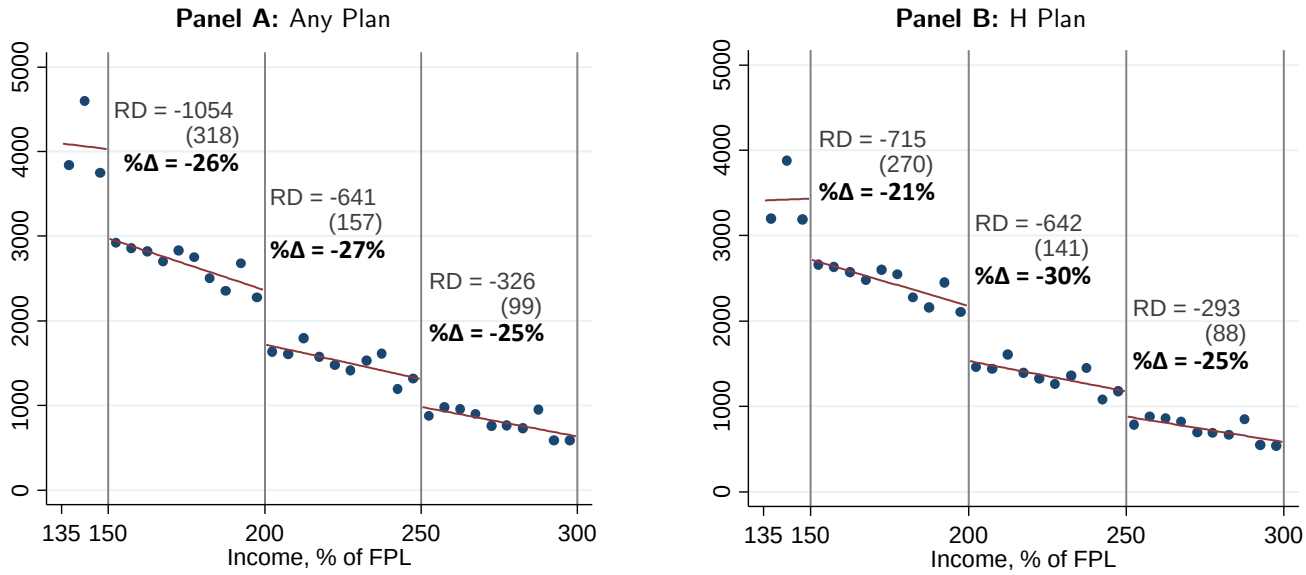
3.3 Evidence from 2011

In most of the rest of the paper, we study data from fiscal year 2011, which has 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 5 shows statistically significant (at the 1% level) declines in overall CommCare and H plan enrollment at each enrollee premium threshold (see Figure 2). The drops in enrollment do not occur only when premiums rise from zero to a positive amount (150% FPL threshold): enrollment falls by

Figure 5: CommCare Enrollment, 2011



NOTE: 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% 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 =".

20-30% at all three thresholds.

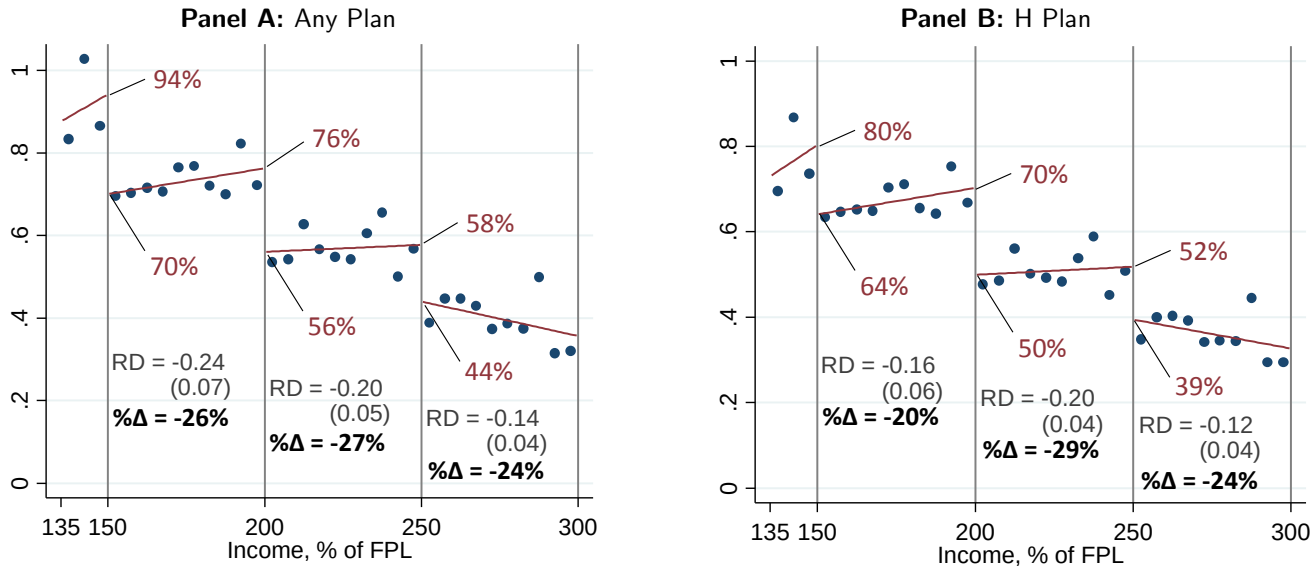
Figure 6 transforms these raw enrollment counts into market shares, using our estimate of the eligible population (see Figure 3) as the denominator. Panel A shows that the share enrolled in any CommCare plan falls by a statistically significant 24-27% at each discontinuity at which enrollee premiums 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 make inferences about the levels of take-up rates as a function of the enrollee premium. Take-up rates fall from 94% when insurance is free (below 150% FPL) to 70%where the cheapest premium increases to $39 per month (above 150% FPL). Take-up rates continue to fall with premiums, declining to below 50% as the cheapest premium rises to $116 per month (above 250% FPL).

Cost of Insurance

Figure 7 shows that average monthly insurer costs rise as enrollee premiums increase at each income threshold. Panel A indicates that at 150% of FPL, average costs for CommCare enrollees increase by $47 (about 14%); this is statistically distinguishable from zero at the 1% level. We also see increases in average costs at the 200% and 250% thresholds, but the increases are somewhat smaller ($31 and $15, or 9% and 4%) 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

Figure 6: Share of Eligible Population Enrolled, 2011



NOTE: NOTE: Figure shows share of eligible population enrolled (in bins of 5% 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% 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 =".

average costs at all three discontinuities; however, these are less precisely estimated.

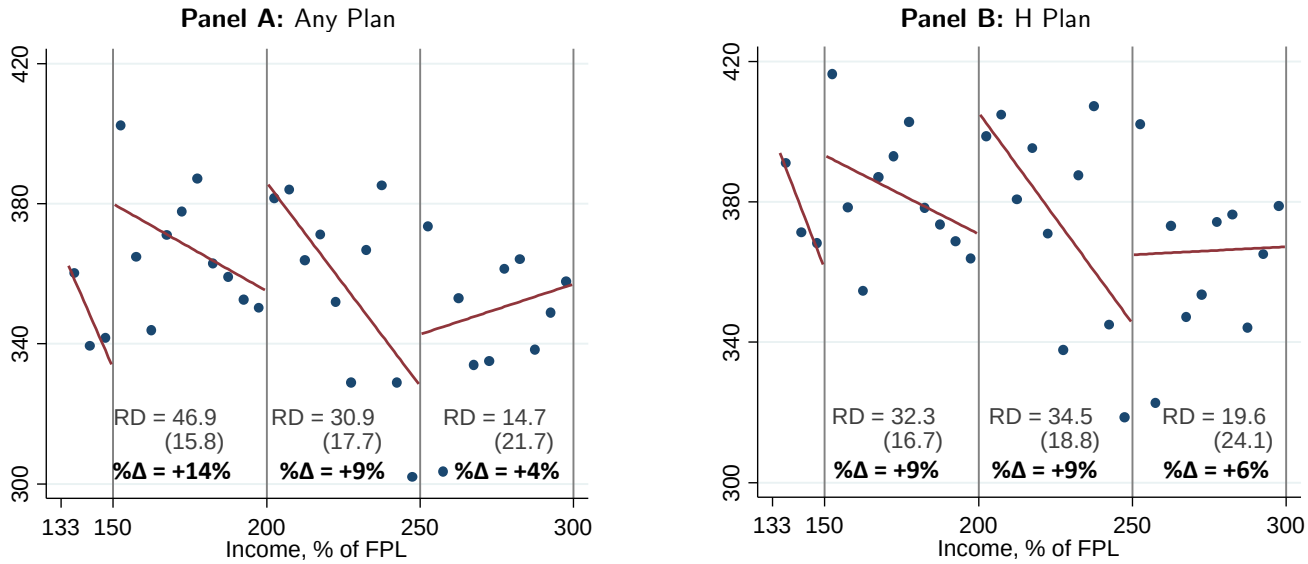
4 Willingness to Pay and Cost Curves

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 et al. (2010a). 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 4.4 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.

4.1 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

Figure 7: Discontinuities in Average Insurer Costs, 2011



NOTE: Figure shows average monthly insurer medical costs for enrollees (in bins of 5% 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% 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 "%Δ =".

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.

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

$$\begin{aligned} u_{ij} &= w(\alpha_j; i) - h - p_{ij} \qquad \text{for } j \in \{L, H\} \\ u_{iU} &= w(\alpha_U; i). \end{aligned}$$

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

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

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:

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

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* $d^2w(\alpha; s)/d\alpha d(1 - s) > 0$.

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 2 the data are consistent with this vertical assumption: when the price of H and L are the same – specifically CommCare enrollees below 100% of FPL for which all plans are free – 96% 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 in a single index of preferences, $1 - s$.[16] 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[17] and D_L depends on both $p_H - p_L$ and p_L.

Figure 8 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 the gap $W_H - W_L$ widens as W_L increases (as one moves left).

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

[16] Note 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.

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

those just to the left enroll in the L plan. Mathematically, these points s^*_{HL} and s^*_{LU} are defined by the equations:

$$\begin{aligned} \Delta W_{HL}\left(s^*_{HL}\right) \equiv W_H\left(s^*_{HL}\right) - W_L\left(s^*_{HL}\right) &= p_H - p_L \\ W_L\left(s^*_{LU}\right) &= p_L \end{aligned} \tag{3}$$

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

$$\text{Positive Demand Condition:} \qquad 0 < s^*_{HL} < s^*_{LU} < 1.$$

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,[18] these are given by:

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

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 above equations 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.[19] 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. We define average costs $AC_j(s)$ as the average costs of all individuals with type $\tilde{s} \leq s$:

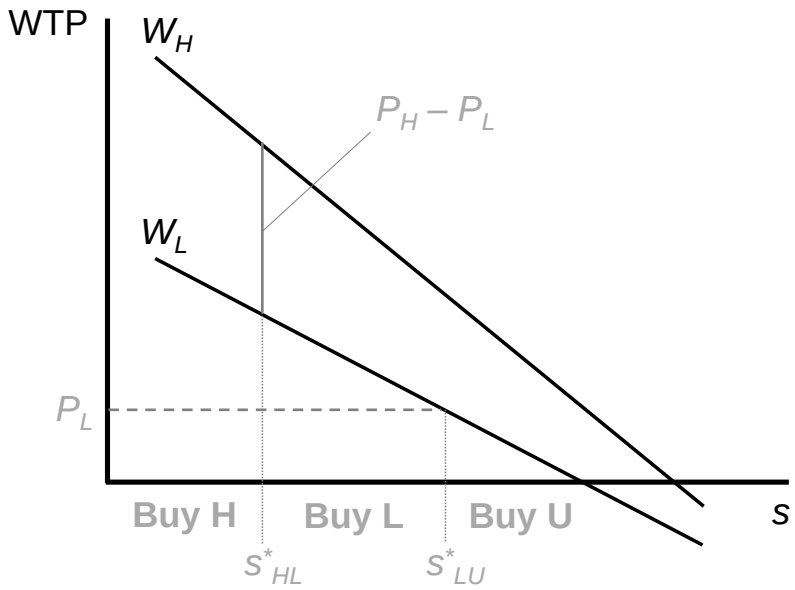
$$AC_j(s) = \frac{1}{s}\int_0^s C_j(\tilde{s})\, d\tilde{s}. \tag{5}$$

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 $AC_j(s)$.

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

[19] 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.

Figure 8: Willingness to Pay Curves: Model



NOTE: 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.

4.2 Constructing Willingness to Pay and Cost Curves

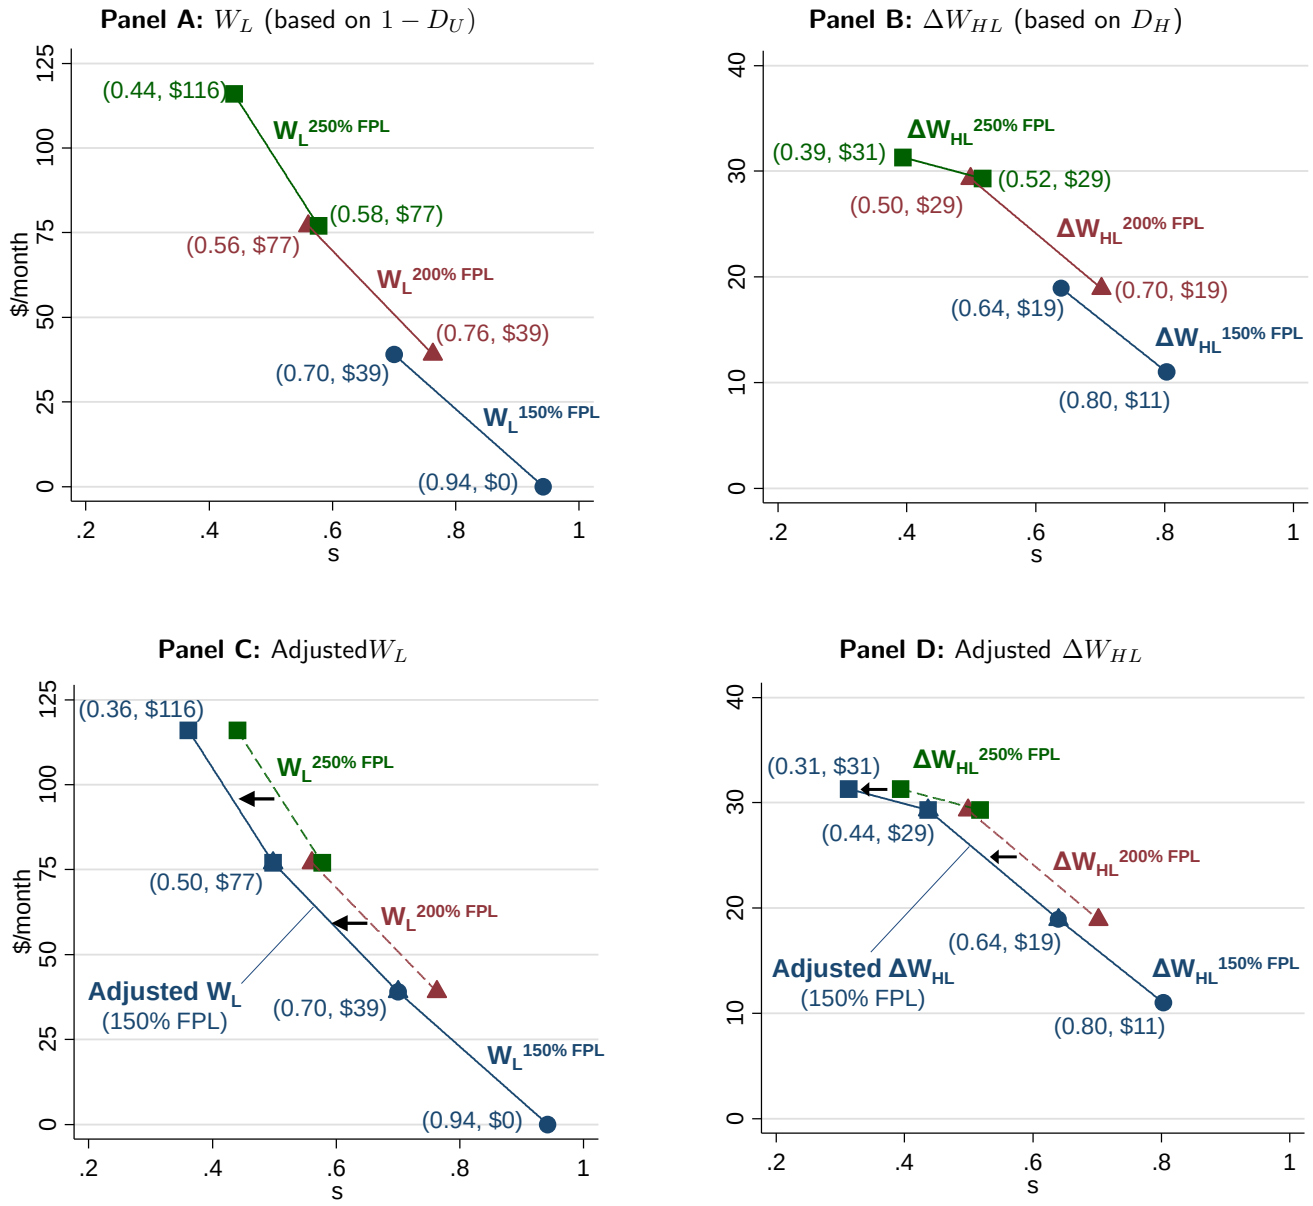
4.2.1 Willingness to Pay (W_j)

We combine the modeling assumptions above with the empirical patterns documented in Figure 6 to construct the empirical analogues of the $W_H(.)$ and $W_L(.)$ curves in Figure 8. Figures 9-10 walk through this exercise.

Panels A and B of Figure 9 plot the willingness to pay curves for the L contract ($W_L(.)$) and for the H contract relative to L ($\Delta W_{HL}(.)$), respectively. Each line segment represents points derived from our three income RDs at 150% FPL (in blue), 200% FPL (red), and 250% 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 2 and 6). 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 it 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

Figure 9: Willingness to Pay Curves: Empirical



NOTE: 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 6 and the premium variation at each discontinuity from Figure 2. 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%, 200%, and 250% 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% FPL line segment.

premiums that apply between 150-200% FPL identify one point on the demand curve for 150% FPL ("from the right") and one point on the demand curve for 200% 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 6.[20] As a result, the demand line segments for the three income groups (shown in different colors in Figure 9) 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. To allow us to combine demand information across income groups, we assume that income functions as a horizontal shifter of the value 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 9 illustrate the implications of this assumption graphically. Specifically, we horizontally shift the Panel A and B value curves estimated at the discontinuities at 200% FPL and 250% FPL so that willingness to pay (i.e. demand) lines up with the curve estimated at 150% FPL. We chose to line up the curves at 150% of FPL since that is the threshold with the greatest share of the eligible population (see Figure 3); results would be qualitatively similar if we instead created a demand curve at the 200% or 250% FPL threshold. In practice, since demand is relatively flat with respect to income, the shift is not very large: the 200% FPL curve is shifted leftward by 6% points in s space, and the 250% FPL curve is shifted leftward by an additional 2% points. The resultant willingness to pay curves consist of three piece-wise linear segments.

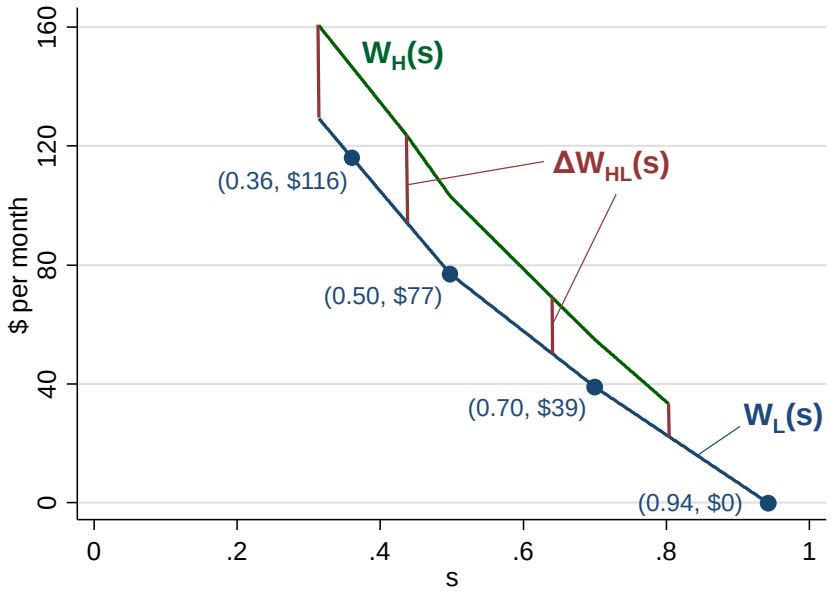
Finally, in Figure 10 we use our estimates of $W_L(.)$and $\Delta W_{HL}(.)$ from Panels C and D of Figure 9 to construct $W_H(s)$ as $W_L(s) + \Delta W_{HL}(s)$ using equation (2). The resulting W_L and W_H curves allow us to infer willingness to pay for L and H for people earning 150% of FPL. Willingness to pay for H is non-trivially higher than for L. The additional value (ΔW_{HL}) ranges from \$11 to \$31/month, or 11-30% of the median type's WTP for L. The median type has total WTP 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% and lowest 6% 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% and lowest 20% of the distribution.

4.2.2 Cost Curves for H: AC_H and C_H

Constructing AC_H In Figure 11 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 7) against the shares in the H plan at each discontinuity (from Panel B of Figure 6). For instance, just below the 150% FPL discontinuity, 80% are in the H plan, and

[20] This 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.

Figure 10: Final Willingness to Pay Curves



NOTE: Figure shows our final estimated WTP curves. The blue curve is the adjusted W_L curve shown in Panel C of Figure 9, 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 9) 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.

the average cost is \$361. Just above the discontinuity, 64% of people are in the H plan and average cost is \$393. Therefore, the average cost curve for 150% FPL flows through the points (64%, \$393) and (80%, \$361).

In Panel B, we once again adjust the average cost curves to obtain a single curve applicable to individuals at 150% 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% FPL and 250% FPL thresholds to align with the 150% FPL average cost curve.[21] 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% FPL curve.

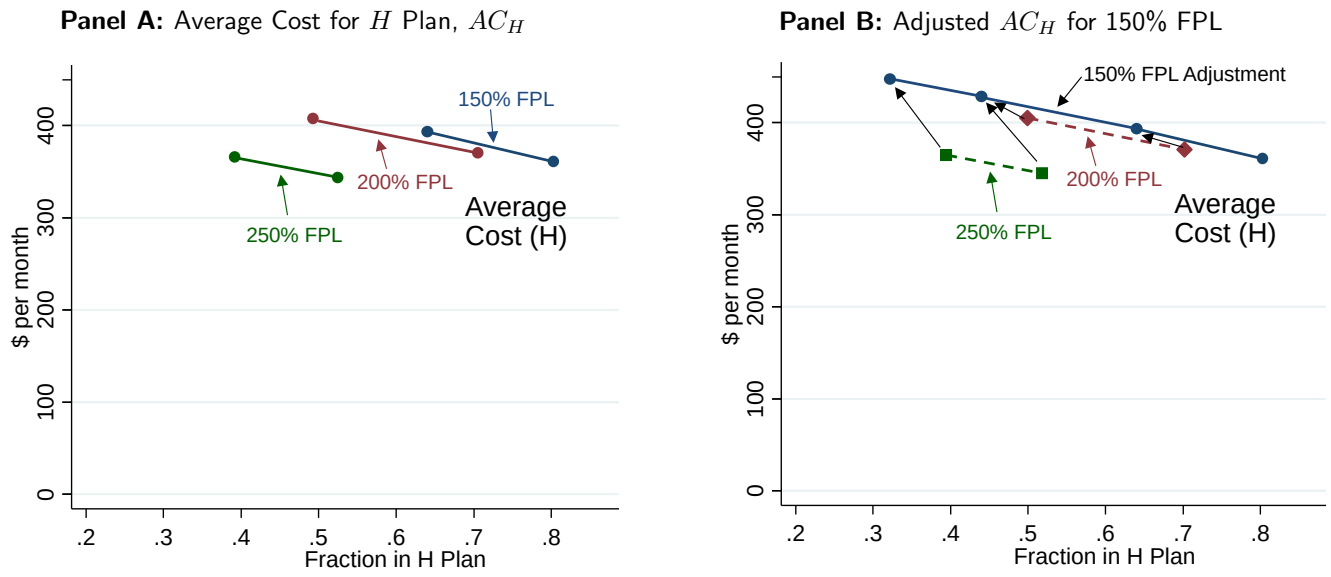
Constructing C_H In addition to the average cost curves, we construct the cost to the insurer of of marginal enrollees, C_H. To do so, note that the *total costs* to the insurer under the H contract at prices $p=\{p_L,p_H\}$ equals:

$$TC_H(p) \equiv \int_0^{s^*_{HL}(p)} C_H(s)\,ds = s^*_{HL}(p)\cdot AC(s^*_{HL})$$

where $AC(s)$ was defined in equation (5). This formula integrates over all the individuals who choose the H contract at these prices. Under the vertical model structure, this corresponds to types for which $s \leq s^*_{HL}(p_H - p_L)$.

[21] Specifically, 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 Figure 7 Panel B).

Figure 11: Construction of Average Cost Curves (H)



NOTE: 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 Figure 7 Panel B; for convenience, Appendix Table 6 summarizes those estimates). Panel B shows how we generate our adjusted AC_H for 150% FPL by translating the 200% and 250% FPL line segments to line up with the 150% FPL segment.

Now, consider the variation induced by the discontinuities, where both p_L and p_H may vary. To capture this, we introduce some additional notation. Let θ parameterize the price changes at the discontinuity so that p_L changes by $\frac{dp_L}{d\theta}$ and p_H changes by $\frac{dp_H}{d\theta}$. The policy induces a change in $p_H - p_L$ of $\frac{d(p_H - p_L)}{d\theta} = \frac{dp_H}{d\theta} - \frac{dp_L}{d\theta}$. Despite the fact that both p_H and p_L vary at the discontinuities, one can still use variation induced by the policy to estimate $C_H(s)$. To see this note that:

$$\underbrace{\frac{dTC_H}{d\theta}}_{\text{Change in total costs in } H} = \underbrace{\frac{ds^*_{HL}}{d\theta}}_{\text{Change in } D_H} \times \underbrace{C_H(s^*_{HL})}_{\text{Cost of marginal consumers}}$$

where $\frac{dTC_H}{d\theta} = \frac{dTC_H}{d(p_H - p_L)} \frac{d(p_H - p_L)}{d\theta}$ is the impact of the policy change (i.e. the discontinuity) on total costs of the H insurers and $\frac{ds_{HL}}{d\theta} = \frac{ds_{HL}}{d(p_H - p_L)} \frac{d(p_H - p_L)}{d\theta}$ is the net impact of the policy on demand for H (since $D_H = s^*_{HL}$). Given estimates of the policy change on total costs of H, $\frac{dTC_H}{d\theta}$ and on demand for H, $\frac{dD_H}{d\theta} = \frac{ds_{HL}}{d\theta}$, we can solve for the cost of the marginal type in the H contract, $s^*_{HL}(p)$,

$$C_H(s^*_{HL}) = \frac{\frac{dTC_H}{d\theta}}{\frac{dD_H}{d\theta}}. \tag{6}$$

Because the pricing change does not affect the costs of infra-marginal types, we can infer the costs of the marginal group by measuring the change in total costs and demand for H. This logic is identical to the two-plan case considered in past work (Einav, Finkelstein, and Cullen, 2010). The key requirement for equation (6) to be valid is that p_H and p_L do not change by the same amount at the discontinuities

(so that $d(p_H - p_L) \neq 0$.)[22]

Panel B of Figures 6 and 7 showed how shares (D_H) and average costs (or AC_H) changed at the pricing discontinuities. We map these – using the AC_H curve adjusted to 150% of FPL from Panel B of Figure 11 – into C_H in (6) using the identity that total costs equal average costs times demand: $TC_H = AC_H \cdot D_H$. The resulting $C_H(s)$ curve is shown in Figure 12, along with the previously estimated curves $W_{H,}$, AC_H and W_L. We place the C_H values along the horizontal axis points that correspond to the midpoint of the relevant average cost segment. The downward slope of each average cost curve in turn implies that the cost curve $C_H(s)$ lies below the average cost curve, $AC_H(s)$.

4.2.3 Cost Curve for L: C_L

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. Absent such independent variation in prices, it is difficult to separate the costs of those who enter/exit the L plan into the H plan, and those who enter/exit the L plan into uninsurance, U. Appendix Figure 20 shows the regression discontinuity results for enrollment in the L plan, the L plan's market share, and average monthly insurer costs among L enrollees. We see statistically significant decreases in enrollment and increases in costs at the 150% threshold. There is little evidence of changes at other thresholds.

However, we can draw some inferences about C_L by exploiting the fact that the market share of L is relatively small.This implies that the average L enrollee is similar to the marginal enrollee. For instance, just above 150% FPL, Appendix Figure 20 shows that just 6% of the population enrolls in the L plan, and we estimate (see Figure 9, Panel A) that the marginal individual who enrollees at that premium is at $s = 0.70$ in the WTP distribution. Thus, the 6% who buy at that premium span $s \in [0.64, 0.70]$, and the average cost of the L plan just above 150% FPL provides an approximation to $C_L(s)$ for individuals in this narrow range of WTP.

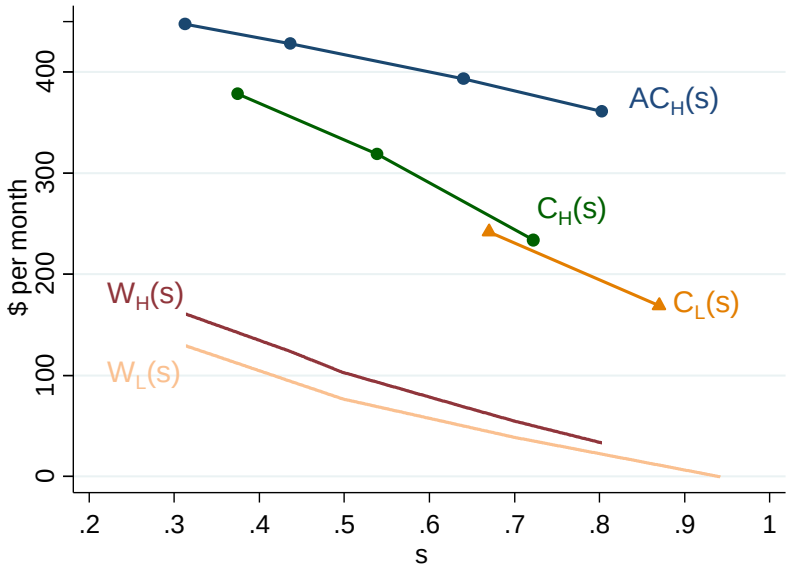
We use this strategy to estimate the $C_L(s)$ for individuals at 150% of FPL. A similar strategy could be used at other income thresholds but we focus on 150% for simplicity. In practice, this means that we use our estimates (see Appendix Table 6) of the average C_L of \$169 per month for those enrolled just below 150% FPL (where the relevant s range is $s \in [0.80, 0.94]$) and \$242 per month just above 150% FPL (where the relevant s range is $s \in [0.64, 0.70]$).[23]

We include these two $C_L(s)$ points (locating them at the midpoint of each s range) in Figure 12. 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).

[22]Note that although the impact on demand is driven both by changes in p_H and p_L, we only need to observe the *net* impact on demand and costs. Under the vertical model, there is only one type of marginal consumer for the H plan – i.e., those with $s = s^*_{HL}$.

[23]Note that the empirical distribution of average claims for those in L shown in Appendix Figure 20, Panel C is not the analog of the "average cost" concept defined in equation (5) since the figure shows average claims for all those enrolled in L; some individuals with higher WTP will in fact enroll in H.

Figure 12: Willingness to Pay and Cost



NOTE: Figure reproducesAC_H, W_H, and W_L from Figures 10 and 11b. The C_H curve is constructed using the formula in equation (6) 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 constructive using the approach described in the text.

4.3 Results and Implications for Take-up

Figure 12 displays our key findings for individuals at 150% of FPL; 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,[24] 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% 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*.[25] 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 69th 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 69th 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

[24] Our price variation spans roughly the 70th to the 6th percentile of the WTP distribution (essentially all but the top 30 percent of the WTP distribution).

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

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 of C_L we can observe. 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% of the population for which we can measure demand for H or L.

Take-up under Counterfactual Subsidies These results imply low take-up of even heavily subsidized insurance for low-income adults. For example, at 150% of poverty, individuals in MA face a \$39 enrollee premium for the L plan, which is a 90% subsidy relative to the insurer price (see Figure 1, Panel B). Our estimates of W_L and W_H indicate that, with a 90% price subsidy, only 69% of the market would enroll if offered L, and only 76% 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% FPL in 2011, the ACA would make the second-cheapest silver plan cost 5.2% of income (or \$83 per month),[26] even though the ACA silver plan has an actuarial value below the plans we study in Massachusetts.[27] Our estimates suggest that at an enrollee premium of \$83, only about half of the market would buy L and only about 60% 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% below average costs (i.e., a subsidy in excess of \$300 per month) only 50% of the population would enroll. Premiums would need to be 90% below average costs in order to induce 80% enrollment. Interestingly, the per-enrollee subsidy cost increases by only \$2 as subsidies move from 90% to 100% of average cost, reflecting the fact that average enrollee costs are declining steeply as healthier individuals are brought into the market.

4.4 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 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

[26] This calculation applies the ACA's subsidy rules (see https://www.irs.gov/irb/2014-50_IRB/ar11.html) – which specify the premium of the second-cheapest silver plan as a percent of income – to the FPL for a single individual in 2011.

[27] We estimate that CommCare plans have an actuarial value of about 97% for enrollees between 100-200% of poverty. In the ACA, the baseline silver plan has an actuarial value of 87% for enrollees between 150-200% FPL.

Table 2: Implications of Alternative Subsidies for H Plan

Enrollee Premium	Subsidy		Share Insured	Marginal Enrollee		Average Cost
	% of AC	$/month		WTP	Cost	
$246	50%	$246	3%	$246	$488	$492
$105	75%	$314	49%	$105	$335	$418
$36	90%	$328	79%	$37	$208	$364
$0	100%	$329	96%	$0	$139	$329

NOTE: Table summarizes the implications of our estimates for enrollment and costs for H under alternative subsidies (shown in the rows). We consider subsidies that lead to enrollee premiums for H of $246, $105, $36, and $0 per month (column 1); these correspond to subsidies of 50%, 75%, 90%, and 100% of equilibrium average costs (column 2). The third 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. Note that while our analysis of 75% and 90% subsidies uses "in-sample" demand and cost curves, analysis of 100% and 50% subsidies requires extrapolating demand and costs outside of our in-sample range, which spans all but the top 31% and bottom 6% of the willingness to pay distribution. Appendix B shows the simple linear extrapolation we use to approximate willingness to pay and costs out of sample for these two estimates.

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%, 200%, and 250% thresholds (see equation 1). In practice, this meant a bandwidth of 50% of FPL everywhere but to the left of the 150% FPL threshold. The first two rows of Table 3 show results using a narrower (25% 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% of FPL threshold.[28] As we discussed in Section 2, this threshold is potentially problematic because of two other small changes at 200% FPL that could affect enrollment or costs independently of the change in enrollee premium: eligibility expands slightly at 200% FPL (to cover pregnant women and HIV-positive individuals for whom Medicaid eligibility ceases) and copays increase slightly at 200% FPL, resulting in a decline in plan actuarial value from 97% to 95%.

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 higher subsidies. A standard way to look for such manipulation is to examine how the density of the population varies around the RD thresholds. Appendix Figure 15 shows the density of the eligible population in the ACS data for 2011, as well as for a similar ACS sample

[28] We fill in the (now missing) space between the 150% and 250% FPL line segments by extrapolating the 150% FPL segment linearly.

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 % of AC_H)	
		WTP_H	C_H	AC_H	75% Subsidy	90% Subsidy
Baseline Estimates	[0.31, 0.94]	$103	$333	$417	49%	79%
(1) Alternate RD Specifications						
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) Alternate Take-up Estimates						
Unscaled ACS Eligible Pop.	[0.19, 0.56]	$24	$186	$354	29%	46%
(3) Accounting for Mandate Penalty						
Use Normalized Premiums	[0.31, 0.94]	$93	$333	$417	32%	68%

NOTE: Top line ("Baseline Estimates") reproduces results from Figure 12 and Table 2. The remainder of the table shows analogous estimates with (1) alternate RD specifications, (2) alternate take-up estimates, and (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.

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. 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.

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 4 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. Appendix Figure 21 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. 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. We expect any manipulation that does occur may reflect learning how to "game" the system by existing enrollees; consistent with this, we find no evidence of such manipulation among new enrollees even in 2012 and 2013 (not shown).

Alternative Estimates of Take-up As we discussed in Section 2, 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% 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 by the raw ACS estimates of the size of the eligible population, we estimate only a 37% 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% subsidy of average costs, only 29% of eligible individuals would enroll in H, compared to 49% in our baseline analysis. Right below 150% FPL, using the the raw ACS estimates for the denominator suggests 56% take-up, compared to our baseline estimate of 94%. 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 or not to buy insurance. While we argue in Section 2 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.[29] For example, our estimates now imply that with a 75% subsidy of average costs, only about 32% 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.[30] 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 joined the exchange prior to 2010 did not have it as an option. However, our main findings are driven by a shift in demand from any formal

[29] Specifically, we normalize premiums by subtracting the expected mandate penalty values (shown by the black dots in Figure 2) 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. Appendix Table 6 reports the normalized premiums by FPL.

[30] Institutionally, 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.

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.[31] This motivated our focus 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 sub-sample 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 Appendix Figures 22 and 23 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, compare 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 Appendix D, 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.

5 Discussion and Normative Implications

Our results suggest that willingness to pay for insurance among a low-income population is extremely low relative to the costs they impose on the insurer. Why might this be the case? And what are the normative implications for health insurance subsidies? We briefly discuss both issues here.

[31] For instance, in 2010 for individuals in the 150-200% FPL group, enrollee premiums for the four H plans varied from \$39 to \$64 per month.

5.1 Why Is WTP so Low?

We estimate that willingness to pay is less than one-third of average costs. This cannot be fully explained by adverse selection driving a wedge between average costs and costs for marginal enrollees. We estimate that willingness to pay is also always less than half of marginal enrollees' *own* expected costs. This stylized fact runs counter to a standard assumption in textbook models of insurance demand – that willingness to pay for insurance equals expected costs *plus* a value of risk protection.

In principle, moral hazard offers a potential explanation. If the provision of health insurance leads individuals to consume care they wouldn't have consumed when uninsured ("moral hazard"), their willingness to pay for this additional care will be less than the cost they would have had to pay for it when uninsured. In practice, however, moral hazard is unlikely to be large enough to explain the results. Insurance would have to increase costs by a factor of at least 200% to explain the estimated gap between willingness to pay (W_j) and own costs (C_j). This seems well outside the plausible range of estimates. For example, the Oregon Health Insurance Experiment finds that Medicaid for low-income adults – the closest analog to CommCare – increases health care spending by about 25% relative to being uninsured (Finkelstein et al., 2012).

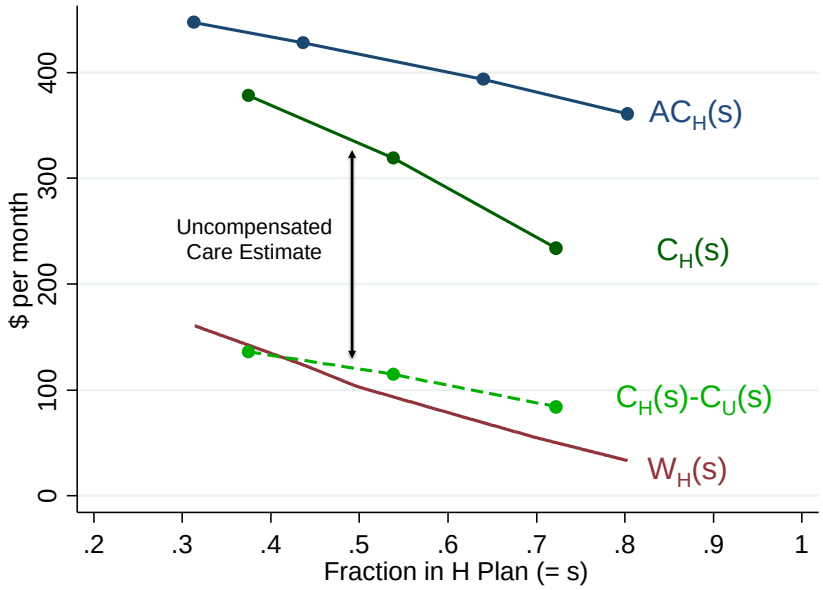
By contrast, uncompensated care for the low-income uninsured is large enough to rationalize their low willingness to pay for CommCare. Uncompensated care for the uninsured reduces willingness to pay relative to the gross cost of this insurance, since some share of those costs pays for care that, if uninsured, would have been paid for by third parties. Estimates suggest that the uninsured pay only about 20% to 35% of their cost of care (Coughlin et al. 2014; Hadley et al. 2008; Finkelstein, Hendren, and Luttmer, 2015), which is remarkably similar to our estimated ratio of WTP to own costs for the H plan. Not surprisingly therefore, a simple back-of-the-envelope calculation suggests that the magnitude of uncompensated care for the low-income uninsured could close most of the gap between willingness to pay (W_H) and own expected cost (C_H).

To illustrate the potential role of uncompensated care, we let $C_U(s)$ denote expected uncompensated care costs that type-s individuals receive if uninsured. Uncompensated care is notoriously difficult to measure. We form a rough approximation of it as:

$$C_U(s) = (1-x)\left(\frac{C_H(s)}{1+\phi}\right). \tag{7}$$

where x denotes the share of the uninsured's total health care costs that they pay out of pocket (thus $1-x$ reflects the fraction provided as uncompensated care) and ϕ denotes the percentage increase in costs that result from insurance coverage (i.e. moral hazard). Hence, total costs of the uninsured are $\frac{C_H(s)}{1+\phi}$. We assume ϕ is 0.25 and x is 0.2 based on results from the Oregon Heath Insurance Experiment (Finkelstein et al. 2012; Finkelstein, Hendren, and Luttmer, 2015). These values imply that $C_U(s) = 0.64 C_H(s)$. Thus the "net" cost of insurance, $C_H(s) - C_U(s)$, is only about one-third of gross costs $C_H(s)$. Figure 13 suggests that this magnitude of uncompensated care can account for nearly all of the gap between willingness to pay and own costs. For example, at the median of the willingness to pay distribution, we estimate net cost of \$120, which is quite close to the willingness to pay estimate of about \$100. At higher points in the willingness to pay distribution (e.g. 30th

Figure 13: Cost Curve Adjusted for Uncompensated Care



NOTE: Figure reproduces $W_H(s)$, $AC_H(s)$ and $C_H(s)$ curves from figure 12. In addition, it shows (in the green dashed line) our approximation of the net cost curve, which is constructed as costs minus estimated uncompensated care (or $C_H(s) - C_U(s)$).

percentile), willingness to pay is about \$10 higher than our approximation of net costs.[32] Hence, a plausible parameterization of uncompensated care shows its sizable role in the economics of health insurance demand: individuals may be willing to pay their net cost from additional coverage but not the gross cost they impose on the insurer that ignores savings for providers of uncompensated care.

A large role for uncompensated care also can potentially explain differential take-up findings for low- vs. high-income populations. The low willingness to pay we find for the low-income population in Massachusetts contrasts with the findings for higher-income individuals in Massachusetts: Hackmann, Kolstad, and Kowalski (2015) estimate that individuals above 300% of the federal poverty line in Massachusetts are willing to pay the (gross) cost they impose on the insurer. One parsimonious way to rationalize these findings is that only low-income individuals are able to obtain substantial uncompensated care, and this uncompensated care when uninsured reduces the willingness to pay for formal insurance. Consistent with this interpretation, our results are similar to Finkelstein, Hendren, and Luttmer (2015) who, based on *ex-post* valuation of Medicaid, find its value to enrollees to be substantially below (gross) insurer costs. Both of these estimates apply to a low-income adult population; the estimates here are for adults at 150% of FPL while Finkelstein, Hendren, and Luttmer (2015) study adults below 100% of FPL.

Figure 13 suggests that subsidies that offset the costs of uncompensated care will be necessary to generate substantial insurance coverage. They may not, however, be sufficient. For example, Figure

[32] For our illustrative purposes, we have focused on the H contract. As shown in Figure 12, the $C_L(s)$ appears similar to $C_H(s)$ over the range we can observe both, suggesting that the net cost estimate in equation (7) would be similar if we calculated it based on C_L(s) instead of $C_H(s)$. Likewise, if we repeated Figure 13 for the L policy, it would similarly suggest willingness to pay for L is similar to or above own expected net cost of L.

13 suggests that willingness to pay is above own net costs for about 40% of the market. However, if subsidies simply covered C_U, prices could still be driven higher than WTP because of adverse selection. The adverse selection we observe is not sufficient to explain low take-up given that WTP is less than own costs, but if WTP were closer to C_H, adverse selection would be more consequential.

More broadly, the analysis in Figure 13 does not preclude other factors that may reduce estimated willingness to pay below the individual's true value of insurance. For example, Hendren (2017) has recently emphasized that observed demand may understate the ex-ante value of insurance measured before the individual has learned something about her health type; in practice, however, as we show in Appendix C, our estimates suggest that even though an *ex-ante* willingness to pay measure may be significantly higher than our baseline willingness to pay estimates, the *ex-ante* willingness to pay measure would still be below own (gross) cost. Liquidity constraints are another potential factor creating a wedge between willingness to pay and value that does not appear quantitatively important. The majority of the 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. Finally, behavioral biases – such as lack of information about eligibility, inattention, inertia, or misperceptions (Spinnewijn (2017); Baicker et al. (2012)) – could also reduce willingness to pay below value. 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. What we can conclude is that the size of uncompensated care for low-income populations can rationalize their low willingness to pay even in the absence of behavioral or other frictions.

5.2 Normative Implications

To consider the normative implications of our findings, we need to make the additional assumption that willingness to pay is in fact the welfare-relevant metric for judging the welfare of health insurance subsidy recipients. This is a standard assumption in economics and has been widely used to estimate the surplus from private (often employer-provided) insurance Einav, Finkelstein, and Levin (2010). However, it does require assuming away behavioral or other frictions that may drive a wedge between estimated demand and the value individuals truly derive from the insurance.

For our normative analysis we follow Hendren (2016) and compute the marginal value of public funds ($MVPF$) for an incremental government subsidy for health insurance. This is given by the ratio of marginal benefit to marginal cost. The marginal benefit of an additional subsidy dollar is defined by the marginal willingness to pay for it; the marginal cost to the government of the additional subsidy incorporates any fiscal externalities resulting from behavioral responses to the subsidy. Thus when fraction s^* of the market is enrolled, we can derive:

$$MVPF\left(s^*\right) = \frac{s^*}{s^* + \frac{C_H(s^*) - W_H(s^*)}{-W_H'(s^*)}} = \frac{1}{1 + \frac{C_H(s^*) - W_H(s^*)}{-s^* W_H'(s^*)}}. \tag{8}$$

Appendix E provides a formal derivation. Intuitively, the marginal benefit of the subsidy is s^*: lowering the premium by \$1 is valued at \$1 for all inframarginal enrollees, and the marginal enrollees who choose

to purchase insurance because of the additional subsidy are indifferent between being insured and uninsured. The marginal cost of the subsidy is $s^* + \frac{C_H(s^*) - W_H(s^*)}{-W'_H(s^*)}$, where the first term is the increased cost for the s^* inframarginal enrollees and the second term follows from the fact that the subsidy expands the market by $\frac{ds^*}{-dp_H} = \frac{1}{-W'_H(s^*)}$ and the marginal enrollee imposes a cost of $C_H(s^*) - W_H(s^*)$. Hence, the cost of additional insurance subsidies is less (greater) than the mechanical cost, s^*, to the extent to which the marginal enrollee's premium $p_H = W_H(s^*)$ is higher (lower) than her cost $C_H(s^*)$.[33] The MVPF will be above 1 if individuals are willing to pay the cost to the government of additional subsidies. This could result from market failures (e.g. adverse selection) that leaves people uninsured despite having a willingness to pay above own expected costs.

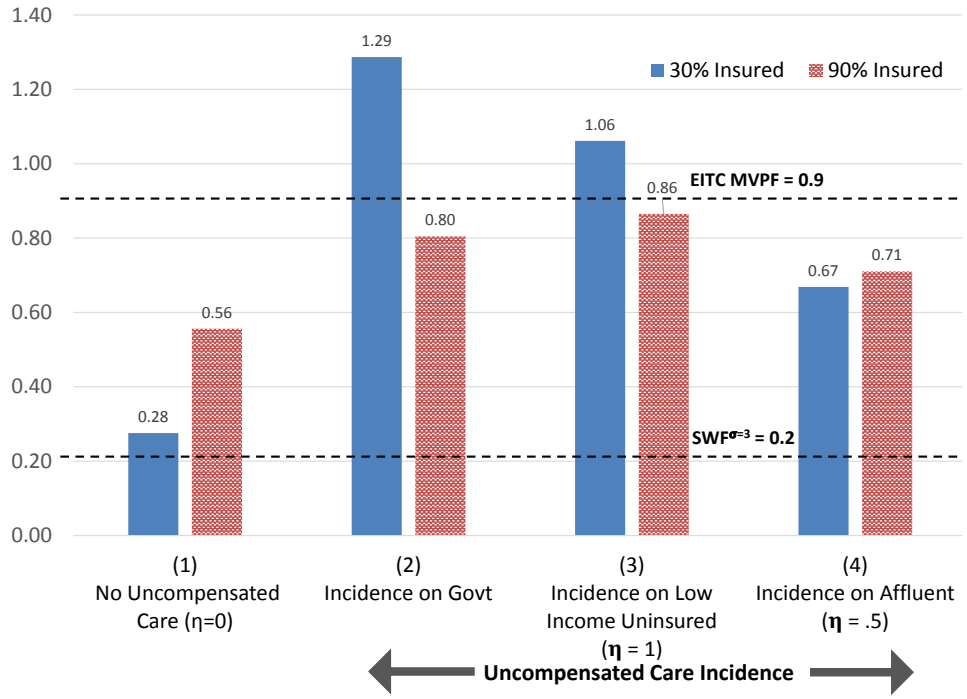
For illustrative purposes, we use the above formula and our estimates in Figure 12 to study the normative implications of subsidizing the H plan in a simplified world without the presence of an L plan. We compute the MVPF at two coverage levels: when 30% and 90% of the market is insured. This corresponds roughly to the in-sample range we observe. We estimate $MVPF(0.30) = 0.28$ and $MVPF(0.90) = 0.56$. An MVPF below 1 indicates that individuals would prefer a \$1 lump-sum transfer to \$1 of insurance subsidies. For example, $MVPF = 0.28$ indicates that recipients are willing to pay about 30 cents for every \$1 increase in government spending on subsidies.

This calculation, however, implicitly assumes that uncompensated care is either non-existent ($C_U(s^*) = 0$) or that it has no social value. The impact of accounting for uncompensated care on the MVPF depends on the incidence of uncompensated care. For example, suppose the government bears the incidence of uncompensated care payments – as would be the case if public funds were used to reimburse hospitals for uncompensated care. Then the cost to the government of providing insurance subsidies – which was $C_H(s^*) - W_H(s^*)$ in the denominator of equation (8) – would be replaced by the net cost, $C_H(s^*) - C_U(s^*) - W_H(s^*)$. If however, the incidence of uncompensated care were on the low-income beneficiaries themselves – as would be the case if reductions in uncompensated care costs on the newly insured allow providers to provide better care to those who are not enrolled – this would affect the numerator rather than the denominator of equation (8). Specifically, the marginal benefit of the insurance subsidies increases from the mechanical s^* impact on the inframarginal enrollees in (8) to $s^* + \frac{C_U(s^*)}{-W'_H(s^*)}$; the marginal benefit of the subsidy increases by $C_U(s^*)$ for each marginal enrollee, with a \$1 subsidy inducing $\frac{1}{-W'(s^*)}$ to enroll. Finally, if the incidence of uncompensated care is borne by third parties, such as hospital owners or those enrolled in employer-provided insurance, one might wish to account for the different social marginal utility of income of these third parties compared to the low-income beneficiaries whose insurance is subsidized. Let η denote the social marginal utility of income to those bearing the incidence of uncompensated care so that society is indifferent between \$1 to the low-income individuals receiving insurance subsidies and $\$\eta$ in payments to those paying the uncompensated care costs. In this case the marginal benefit of the subsidy would be $s^* + \eta \frac{C_U(s^*)}{-W'_H(s^*)}$.

Figure 14 illustrates the impact of these different assumptions about uncompensated care on our estimates of the MVPF of additional subsidies, starting from having either 30% or 90% of the market insured. We parameterize $C_U(s)$ using the methods described in Section 5.1. The first panel shows the

[33] This cost definition abstracts from any potential labor market activity response to the means-tested health insurance subsidies. If individuals reduce earnings in response to higher subsidies, this would introduce an added cost that would need to be incorporated into the denominator.

Figure 14: Marginal Value of Public Funds (MVPF) for insurance subsidies



NOTE: Figure shows estimates of the marginal value of public funds (MVPF) for insurance subsides, based on equation (8) and different assumptions about the initial level of insurance coverage and the presence and incidence of uncompensated care for the uninsured. See text for details.

MVPF ignoring uncompensated care (i.e $\eta = 0$). Panels (2)-(4) show the MVPF when, respectively, the incidence of uncompensated care payments is on the government, the low-income uninsured ($\eta = 1$), and more affluent third parties for whom $\eta = 0.5$.[34] Figure 14 shows that whether the MVPF from additional subsidies is above 1 depends on the incidence of uncompensated care, as well as the baseline share of the market insured.

Of course, the appropriate benchmark for the MVPF may not be 1. Insurance subsidies for low-income adults may serve redistributive aims, and redistribution inevitably involves some resource cost (Okun, 1975). In Figure 14 we therefore consider two other benchmarks. In the first we ask: given a specific social welfare function, individual utility function, and the consumption level of the eligible population relative to the general population, how high would the MVPF have to be to make additional subsidies desirable (Saez and Stantcheva, 2016)? We consider a utilitarian social welfare function over individual CRRA utility functions, and assume median consumption of our study population is 60% that of median consumption in the general population in Massachusetts.[35] With a coefficient of relative risk aversion of 3 (i.e. marginal utility of $u'(c) = 1/c^3$), this would suggest society should

[34] An $\eta = 0.5$ can be motivated by assuming the incidence of uncompensated care falls on more affluent individuals whose social marginal utility is 50% that of the direct subsidy recipients. Alternatively, it can be motivated by Hendren (2014)'s estimate that, because of the distortionary cost of taxation, $1 at the top of the income distribution can be turned into only $0.50 for those at the bottom.

[35] This is derived from authors' calculations from the Consumer Expenditure Survey. We compare median consumption in the recipient population to the general population under the assumption that the subsidy is financed through a uniform lump sum tax on the population, and we ignore fiscal externalities from income effects that result from such a tax.

prefer policies with an MVPF of at least 0.20, a threshold crossed by all of our MVPF estimates. Naturally, more concave social welfare functions would further increase social value, while more (less) concave individual utility functions would increase (decrease) social value.

A second benchmark, which avoids having to take a stand on the curvature of the social welfare function or the individual utility function, asks whether these subsides are more or less costly than other redistributive programs (Hendren, 2014). A natural comparison is to the Earned Income Tax Credit (EITC) – the primary form of tax subsidies to low-income individuals in the U.S. Hendren (2016) draws on existing empirical estimates of the impact of the EITC on labor supply to estimate a MVPF of 0.9 for the EITC. Figure 14 indicates that whether additional health insurance subsidies for low-income individuals deliver positive welfare benefits relative to increasing the EITC (i.e., whether the MVPF is above 0.9) once again depends on the share of the market already insured and the incidence of uncompensated care.

6 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% of the low-income eligible population, we find that willingness to pay for insurance is far below insurers' average costs. Adverse selection exists, despite the presence of the coverage mandate, but is not the driving force behind low take up. We estimate that willingness to pay is less than half of enrollees' own expected costs; thus, even if insurers could set prices conditional on an individual's willingness to pay, at least 70% of the market would be uninsured.

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% of insurer costs would generate only 80% coverage. This reality may underlie the incomplete take-up of insurance under the ACA, despite a coverage mandate and generous subsidies. We provide suggestive evidence of an important role for uncompensated care in explaining why willingness to pay is substantially below (gross) insurance costs. This in turn suggests a potential economic rationale for subsidies as a response to the "Samaritan's dilemma" to offset the implicit tax that uncompensated care imposes on formal insurance (Coate, 1995).

If we are willing to interpret our estimates of willingness to pay as revealing individuals' value for insurance, our results also have normative implications for such subsidies. In particular, they suggest that if individuals had to pay the full cost imposed on the insurer, mandating these low-income individuals to buy health insurance would not make them better off. We find robust evidence that low-income individuals value health insurance below the cost they impose on the insurance company. *A priori*, this need not have been the case if adverse selection were the primary source of lack of insurance. Our results therefore suggest that adverse selection, while it exists in this market, is not by itself provide a sufficient rationale for health insurance subsidies. Further, we show that normative analysis of the insurance subsidies depends critically on the economic incidence of the provision of

uncompensated care. Better understanding this incidence is an important, albeit challenging, area for future work.

References

Baicker, K., S. Mullainathan, and J. Schwartzstein (2012). Behavioral hazard in health insurance. *NBER Working Paper Series* (18468).

Bhargava, S. and D. Manoli (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an irs field experiment. *The American Economic Review 105*(11), 3489–3529.

Brooks, T., K. Wagnerman, S. Artiga, E. Cornachione, and P. Ubri (2017). Medicaid and chip eligibility, enrollment, renewal, and cost sharing policies as of january 2017: Findings from a 50-state survey. *Kaiser Family Foundation Report*.

Buchanan, J. (1975). The samaritan's dilemma. *In Altruism, Morality, and Economic Theory, edited by Edmund S. Phelps.*.

Chan, D. and J. Gruber (2010). How sensitive are low income families to health plan prices? *The American Economic Review 100*(2), 292–296.

Chandra, A., J. Gruber, and R. McKnight (2011). The importance of the individual mandate – evidence from massachusetts. *New England Journal of Medicine 364*(4), 293–295.

Coate, S. (1995). Altruism, the samaritan's dilemma, and government transfer policy. *The American Economic Review*, 46–57.

Connector, M. H. and D. of Revenue (2011). Data on the individual mandate: Tax year 2011. Technical report.

Coughlin, T. A., J. Holahan, K. Caswell, and M. McGrath (2014). Uncompensated care for the uninsured in 2013: A detailed examination. *Kaiser Family Foundation*.

Dague, L. (2014). The effect of Medicaid premiums on enrollment: A regression discontinuity approach. *Journal of Health Economics 37*, 1–12.

De Nardi, M., E. French, and J. B. Jones (2016). Medicaid insurance in old age. *The American Economic Review 106*(11), 3480–3520.

Diamond, R. and P. Persson (2016). The long-term consequences of teacher discretion in grading of high-stakes tests. *NBER Working Paper Series*.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2016). The economic consequences of hospital admissions. *NBER Working Paper Series*.

Einav, L. and A. Finkelstein (2011). Selection in insurance markets: Theory and empirics in pictures. *The Journal of Economic Perspectives 25*(1), 115–138.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics 125*(3), 877–921.

Einav, L., A. Finkelstein, and J. Levin (2010). Beyond testing: Empirical models of insurance markets. *Annual Review of Economics 2(1)*, 311–336.

Ericson, K. M. M. (2014). Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *American Economic Journal: Economic Policy 6*(1), 38–64.

Finkelstein, A., N. Hendren, and E. F. Luttmer (2015). The value of medicaid: Interpreting results from the oregon health insurance experiment. *NBER Working Paper Series*.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, and K. Baicker (2012). The Oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics 127(3)*, 1057–1106.

Frean, M., J. Gruber, and B. 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, C., T. Gross, and M. J. Notowidigdo (2015). Hospitals as insurers of last resort. *NBER Working Paper Series*.

Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2015). Adverse selection and an individual mandate: When theory meets practice. *The American economic review 105*(3), 1030–1066.

Hadley, J., J. Holahan, T. Coughlin, and D. Miller (2008). Covering the uninsured in 2008: current costs, sources of payment, and incremental costs. *Health Affairs 27*(5), w399–w415.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *The American Economic Review 103*(7), 2643–2682.

Health, A. (2016). The state of exchanges: A review of trends and opportunities to grow and stabilize the market. Technical report.

Hendren, N. (2014). The inequality deflator: Interpersonal comparisons without a social welfare function. *Working Paper*.

Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy 30*(1), 51–89.

Hendren, N. (2017). Measuring ex-ante welfare in insurance markets. *Working Paper*.

Hu, L., R. Kaestner, B. Mazumder, S. Miller, and A. Wong (2016). The effect of the patient protection and affordable care act medicaid expansions on financial well-being. *NBER Working Paper Series*.

Institute, M. M. P. (2012). Masshealth: The basics – facts, trends, and national context. Technical report.

Jaffe, S. and M. Shepard (2017). Price-linked subsidies and health insurance markups.

Janicki, H. (2013). Employment-based health insurance: 2010. Technical report, U.S. Census Bureau.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics 8*, 435–464.

Krueger, A. and I. 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)*, 780–793.

Long, S. K., K. Stockley, and H. Dahlen (2012). Health reform in massachusetts as of fall 2010: Getting ready for the affordable care act and addressing affordability. Technical report, Blue Cross Blue Shield Foundation of Massachusetts.

Mahoney, N. (2015). Bankruptcy as implicit health insurance. *The American Economic Review 105*(2), 710–746.

Okun, A. (1975). *Equality and Efficiency.* Brookings Institution Press.

Polyakova, M. (2016). Regulation of insurance with adverse selection and switching costs: Evidence from medicare part d. *American Economic Journal: Applied Economics 8*(3), 165–195.

Ruggles, S., K. Genadek, R. Goeken, J. Grover, and M. Sobek (2015). Integrated public use microdata series: Version 6.0 machine-readable database. Minneapolis: University of Minnesota.

Sacks, D. W. (2017). How do the mandate and the premium tax credit affect the individual insurance market?

Saez, E. and S. Stantcheva (2016). Generalized social marginal welfare weights for optimal tax theory. *The American Economic Review 106*(1), 24–45.

Smith, V. K., K. Gifford, E. Ellis, R. Rudowitz, L. Snyder, and E. 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. Technical report.

Spinnewijn, J. (2017). Heterogeneity, demand for insurance and adverse selection. *American Economic Journal: Economic Policy 9*(1), 308–343.

Tebaldi, P. (2016). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca.

Tirole, J. (1988). *The Theory of Industrial Organization.* MIT press.

U.S. Department of Agriculture (2016). The food assistance landscape: FY 2015 annual report.

U.S. Department of Health and Human Services (2015). 2015 actuarial report on the financial outlook for Medicaid.

U.S. Internal Revenue Service (2015). EITC calendar year report.

Online Appendix

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

Amy Finkelstein, Nathaniel Hendren, Mark Shepard

Appendix A. Estimating the CommCare Eligible Population in the ACS

We estimate the size of the eligible CommCare population in Fiscal Year 2011 using data from the 2010-2011 American Community Survey (ACS). This appendix describes how we estimate the size of the population eligible for CommCare using the American Community Survey (ACS). Our estimation has two main steps: first, we apply CommCare's eligibility criteria to the ACS data to limit to the sample of individuals likely eligible for CommCare. Second, we estimate the eligible population by income bin, using a regression to smooth the raw counts in each bin. Below, we describe each step in detail.

Applying CommCare Eligibility to ACS Data

We begin with ACS data from 2010-2011, using both years because our CommCare year of interest, fiscal year 2011, spans July 2010-June 2011.[36] To take an average of the population across years, we multiply sample weights by 1/2.

We begin by defining family income as a share of the poverty line, analogous to the measure used by CommCare. Specifically, we sum total personal income across all members of an individual's "health insurance unit" (HIU), a variable defined by the University of Minnesota's SHADAC to approximate family unit definitions used by public insurance programs. We divide this total income by the FPL defined by the year and the HIU size.

We then define people as CommCare eligible if they are U.S. Citizens in the relevant age range (19-64) and income range (less than 300% FPL) who are not enrolled in another form of health insurance (specifically, employer insurance, Medicare, or Tricare) and are not eligible for Medicaid (based on income and demographics). We discuss each of these restrictions in turn. The top panel of table 4 shows the ACS sample size, the (unweighted) number of individuals dropped at each of stages, and the estimated population size (scaled up using 1/2 of the ACS person weights).

Restricting to the relevant age range (row 1) and income range (row 2) is straightforward. In row 3 we restrict to U.S. citizens. Nearly all non-citizens are ineligible for CommCare, with the exception of long-term green card holders (longer than 5 years). Because we cannot separately measure this latter group in the ACS, we exclude all non-citizens. In row 4, we exclude any individual who reports

[36] We obtained ACS data from the IPUMS-USA website (Ruggles et al., 2015).

having employer-sponsored insurance (ESI), Medicare, Tricare, privately purchased insurance, or VA insurance.

The last two rows of the table show how we exclude individuals eligible for Medicaid (MassHealth) instead of CommCare. We cannot directly measure Medicaid eligibility in the ACS.[37] Instead, we approximate it by excluding the two largest groups we know are eligible: parents with income below 133% of FPL, and the disabled (proxied by under 65 and SSI receipt). Parents with dependents under 18 are eligible for Medicaid below 133% FPL and eligible for CommCare above this cut-off. We focus on income groups above the cutoff so that results are not affected by this large compositional change in eligibility, and use the 135% FPL cutoff to avoid ambiguity right at the 133% FPL cutoff and to maintain equal-size 5% FPL bins for later analysis. This approach misses a few groups who are Medicaid eligible but whom we cannot easily measure in the ACS - specifically, pregnant women below 200% FPL and HIV-positive people below 200% FPL. Based on the number of women below 200% FPL with a child under one year old, we estimate that pregnant women may constitute 0.4% of our eligible sample. The HIV-positive group is likely to be even smaller.[38]

The final sample in the ACS includes 2,856 observations. Scaling this up to a population size using the ACS's person weights, we estimate a CommCare-eligible population size of 168,041 for FY 2011.

Estimating Smoothed Eligible Population

We next use this restricted sample to estimate the CommCare-eligible population by income bin.

Figure 15 shows the raw estimates, with each point representing the estimated eligible population size for a 5% of FPL bin. These estimates are quite noisy, both because the ACS is a 1% sample and because of clustering in reported income at round numbers (which is emphasized in the very high outlier points). To prevent this noise from introducing error into our estimates of market shares from the administrative data, we construct a smoothed estimate of the eligible population by income bin. Specifically, we regress the raw population counts by 1% of FPL bin on a quadratic polynomial in income as a percentage of FPL. The predicted values from this regression (multiplied by 5 to match the 5% FPL bins we use in our analysis) are shown in the red curve in Figure 15.

We use the value of this curve at the midpoint of each 5% FPL income bin for our smoothed estimate of the CommCare-eligible population size.

Appendix B: Extrapolating out of sample

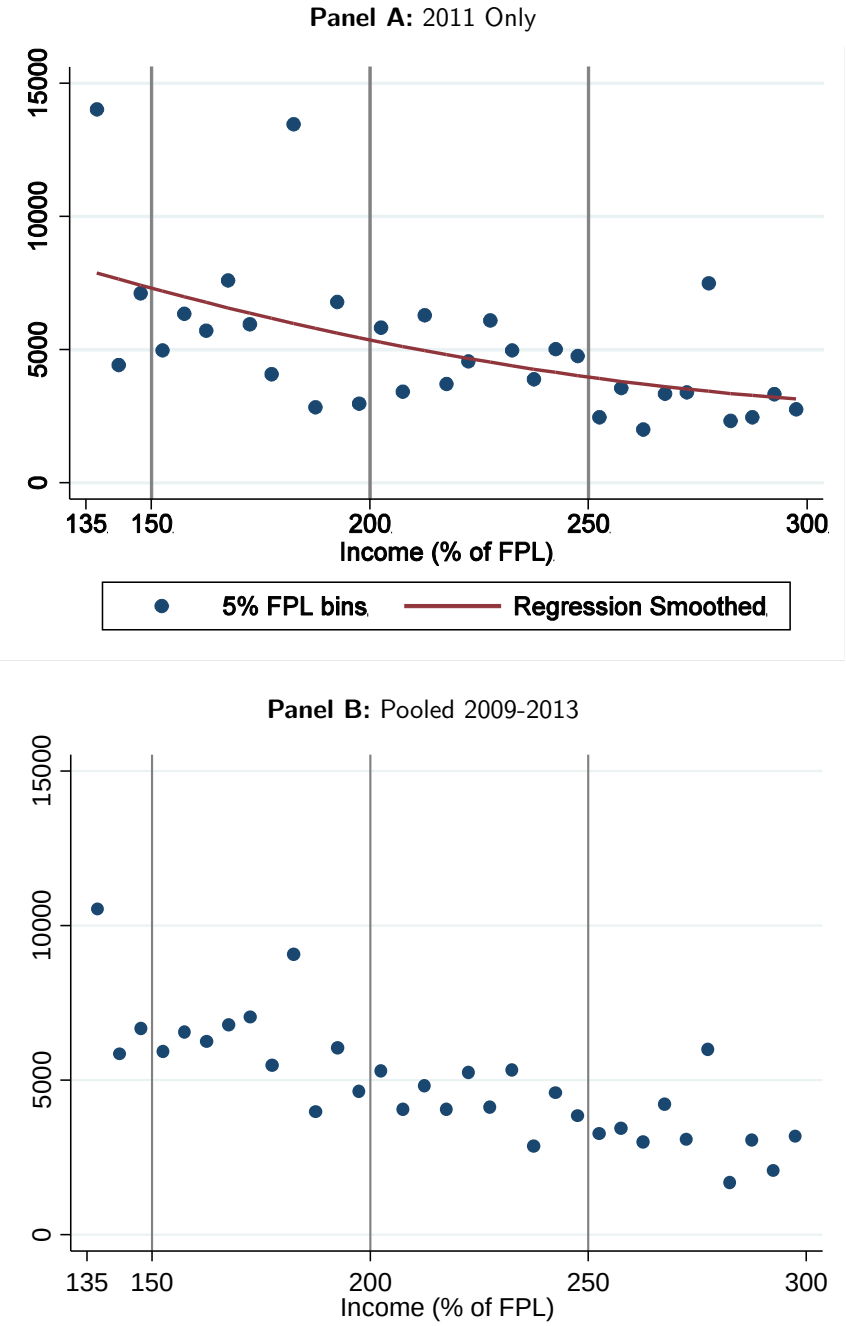
Appendix Figure 16 shows the simple extrapolation we used to approximate willingness to pay and cost out of sample; we use this out-of-sample extrapolation only for the 100% and 50% subsidy counterfactuals in Table 2.

We extrapolate by extending the left-most segments of W_H and AC_H linearly to the left (and

[37] We cannot even directly measure Medicaid enrollment; the ACS does not distinguish between Medicaid and CommCare (both are coded as "Medicaid/other public insurance").

[38] Details of Medicaid eligibility rules are based on MMPI (2012).

Figure 15: Estimate of CommCare Eligible Population from ACS Data



NOTE: Panel A shows our smoothed estimates of the CommCare-eligible population from the 2010-2011 ACS data. The dots are raw estimates of the annual eligible population size (weighting by the ACS "person weight" to generate a population estimate) by 5% of FPL bin. Because these data are relatively noisy – especially at high outlier points that reflect round income numbers like $20,000 – we use a quadratic regression to generate a smoothed estimate of the eligible population size. The resulting estimates are shown in the curve. Panel B shows the raw estimates of the annual eligible population size from the pooled 2009-2013 ACS data.

Table 4: ACS Sample Construction

	ACS Sample Size			Est. Population
Sample / Exclusion	# Dropped	% Dropped	# Remaining	# Remaining
Full ACS Mass. Sample (2010-11)			**135,009**	**6,572,395**
Drop Age <19 or ≥ 65	52,362	39%	82,647	4,147,512
Drop Income > 300% FPL	47,790	35%	34,857	1,863,450
Drop Non-Citizens	3,994	3%	30,863	1,602,621
Drop People with Medicare, ESI, Tricare	19,961	15%	10,902	601,145
Sample Eligible for CommCare or Medicaid	--	--	**10,902**	**601,145**
Limit to 135-300% FPL	7,862	72%	3,040	178,772
Drop Disabled (under 65 and receiving SSI)	184	2%	2,856	168,041
Final Sample	--	--	**2,856**	**168,041**

NOTE: The table shows how we construct our ACS sample of the population eligible for CommCare, as described in the text of Appendix A. It starts from the full ACS 2010 and 2011 (pooled) samples of Massachusetts residents and shows the number of observations dropped and remaining at each step. The final column refers to the estimated population size, applying the appropriate ACS "person weights" (and dividing in half to compute an annual estimate from the two years of pooled data).

recalculating c_H using equation (6)).[39] Likewise, we extend the right-most segments of W_H and AC_H out to its value above the right-most value of W_L that we observe ($s = 94\%$) using a linear extrapolation of the right-most segment of the ΔW_{HL} and AC_H curves.[40]

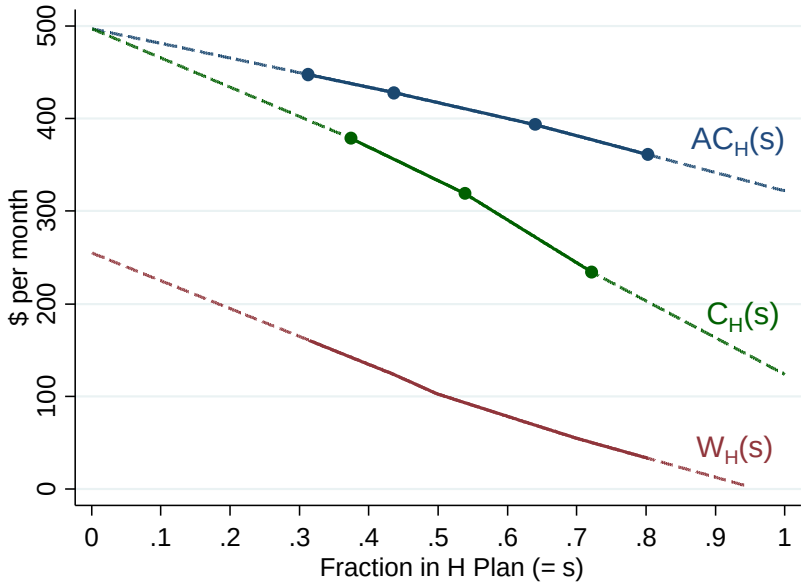
This extrapolation leads to estimates of willingness to pay that are still everywhere far below average and own costs. Because demand lies everywhere below average costs (falling short by more than \$225; W_H never exceeds 52% of average costs), our extrapolation suggests the market would fully unravel in the absence of large government subsidies. Because individuals' WTP lies everywhere below their own expected cost they impose on the insurer (by at least \$140; W_H never exceeds 53% of C_H), this suggests that even if insurers were able to price discriminate on WTP type (i.e., based on s), the market would unravel. In this sense, adverse selection cannot explain the low take up of H by low-income individuals. Even if the quantile of WTP, s, were known to insurers and they were allowed to price on this information, they would still not be able to profitably sell insurance.

Naturally, a concern with this linear extrapolation is that it assumes away the possibility that there is a subset of the population with much higher demand than other types, so that demand increases non-linearly. While we cannot of course rule this out, the most natural source for willingness to pay increasing non-linearly would be if the variance of costs (i.e. risk) were higher for higher willingness to pay individuals, and this does not appear to be the case. To test for this, Appendix Figure 17 explores how the standard deviation of costs changes around our pricing discontinuities. While the

[39] It would also be possible to extrapolate c_H linearly and recalculate AC_H accordingly. In practice, this produces even higher estimates of both cost curves, meaning that our conclusion that WTP is entirely below costs is unchanged.

[40] Interestingly, W_H is quite close to zero at this point (about \$4). This does not occur by construction but is consistent with the idea that these people have virtually no willingness to pay for health insurance, whether H or L. The model implies that W_L is zero or negative for the rest of the s distribution, and based on the estimates, the same is also true of W_H; in the context of the model, this can be explained by transaction costs in enrolling (even at a zero "sticker" price), which reduces the net willingness to pay for a formal contract below the uninsured option. Of course, it is also possible that these 6% who do not enroll at zero price are uninformed about their eligibility (e.g. Bhargava and Manoli (2015)).

Figure 16: Value and Cost of H – Extrapolation



NOTE: The figure shows out-of-sample extrapolations for W_H, AC_H and C_H. The solid lines are our in-sample estimates, identical to those shown in Figure 12. The dashed lines are the extrapolations. Both W_H and AC_H are extrapolated linearly using the slope of the left-most and right-most line segment. C_H is extrapolated by calculating C_H based on the implied values from AC_H, applying the formula for C_H in equation (6).

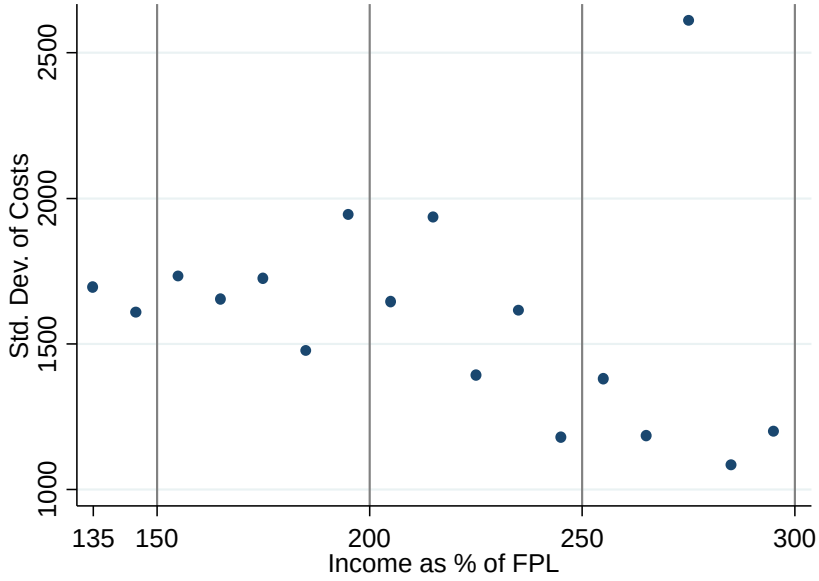
results are fairly noisy, there is no evidence that the standard deviation of costs is increasing at the price discontinuities where willingness to pay of those enrolled is also increasing. While this does not guarantee that the linear extrapolation is appropriate, it does suggest that it is not entirely inconsistent with the underlying cost variation in the data.

Appendix C. Willingness to Pay Behind the Veil of Ignorance

One potential concern with comparing willingness to pay to costs is that demand is measured after some information about health risk may potentially have been revealed to the individual. For example, suppose demand is measured after one learns whether or not she has a chronic condition. In this case, observed demand will understate the *ex-ante* value of insurance that would be measured before the individual has learned their risk. Hendren (2017) provides a method for calculating willingness to pay for insurance from behind the veil of ignorance. Instead of using the observed market demand curve, $W(s)$, one uses an "*ex-ante*" demand curve, $W(s) + EA(s)$, where $EA(s)$ captures the value of expanding the size of the insurance market from the perspective of behind the veil of ignorance.

For a linear demand curve, the formula in Hendren (2017) for the *ex-ante* component of willingness to pay is given by $EA(s) = (1-s)(C(s) - W(s) - sW'(s))\gamma\frac{1-s}{2}(12 * W'(s))$ where γ is the coefficient of absolute risk aversion and the factor of 12 translates our monthly demand estimates into yearly units. We apply this formula in our context using a conservatively high coefficient of absolute risk aversion of $5x10^{-4}$ (which corresponds to a coefficient of relative risk aversion of 5 if individuals have

Figure 17: Standard Deviation of Insurer Costs across Enrollees, by %FPL (H Plan)



NOTE: The graph shows the standard deviation of insurer costs, by 10% of FPL bin. The standard deviation is calculated across individuals in the data for 2011, using each individuals average insurer-paid cost per month enrolled. As we discuss in the text, there is little evidence that the standard deviation jumps discretely at the income thresholds where subsidies and take-up changes.

$10,000 of annual consumption). Our estimates suggest that even using an *ex-ante* demand measure and a high value of γ, willingness to pay would still be below own cost. For example, at $s = 0.50$ we estimate that the marginal welfare impact from behind the veil of ignorance of expanding the size of the insurance market is roughly $EA\left(0.5\right) = 0.5\left(333 - 103 - 0.5 * (-239)\right) 0.0005 * \frac{0.5}{2} * 12 * (-239) = \$63$ higher (i.e., $163 instead of $100) than the marginal welfare impact implied by observed demand. Although non-trivial, this is small relative to the approximately $300 difference between marginal cost and observed demand at $s = 0.5$. The intuition for this is that the "risk" of learning that one is a high risk type and must purchase insurance is not exceedingly large when premiums are already heavily subsidized.

Appendix D. WTP Estimates without a Vertical Model

Our vertical model involves non-trivial assumptions about the nature of the market and preferences of consumers. These assumptions – while reasonable for CommCare in 2011 – will not be reasonable in all settings, including for CommCare in other years. In this section, we develop a model with fewer assumptions that gives us bounds on the WTP for access to a given set of insurance plans.

Setup and WTP for Insurance The setup is very general. Consider an insurance market with plan options $j = 1, ..., J$ and an outside option of uninsurance, $j = U$. Let W_{ij} be the willingness to pay of consumer i for plan j, where we normalize WTP relative to $W_{iU} = 0$. Let p_{ij} be the premium of each plan (which can vary across consumers), where we also normalize $p_{iU} = 0$. Consumers choose

among available options to maximize their utility, which equals:

$$u_{ij} = W_{ij} - p_{ij}.$$

Note that by our normalizations, $u_{iU} = 0$.

We would like to estimate the willingness to pay for (any) CommCare insurance (W_i^{Ins}), defined as the willingness to pay for each consumer's most preferred plan:

$$W_i^{Ins} = \max_{j \neq U} \{W_{ij}\}$$

A challenge in measuring the WTP for insurance is that while we want the maximum value of W_{ij} across j options, in CommCare consumers choose plans to maximize $W_{ij} - p_{ij}$. However, we can use choices in CommCare to get bounds on W_i^{Ins}. To do so, note that if someone chooses U in the CommCare setting, it implies that $W_{ij} \leq p_{ij}$ for all $j \neq U$. This in turn implies that

$$\text{Choose } U: \qquad W_i^{Ins} \leq \max_{j \neq U} \{p_{ij}\} \equiv p_i^{max} \tag{9}$$

Thus, p_i^{max} is an upper bound on the value of access to CommCare for people who choose not to buy into it. Similarly, if an individual chooses to take up CommCare, we know that $W_{ij} \geq p_{ij}$ for at least one plan. Therefore, we can bound their W_i^{Ins} from below by the cheapest plan's price:

$$\text{Choose } j \neq U: \qquad W_i^{Ins} \geq \min_{j \neq U} \{p_{ij}\} \equiv p_i^{min} \tag{10}$$

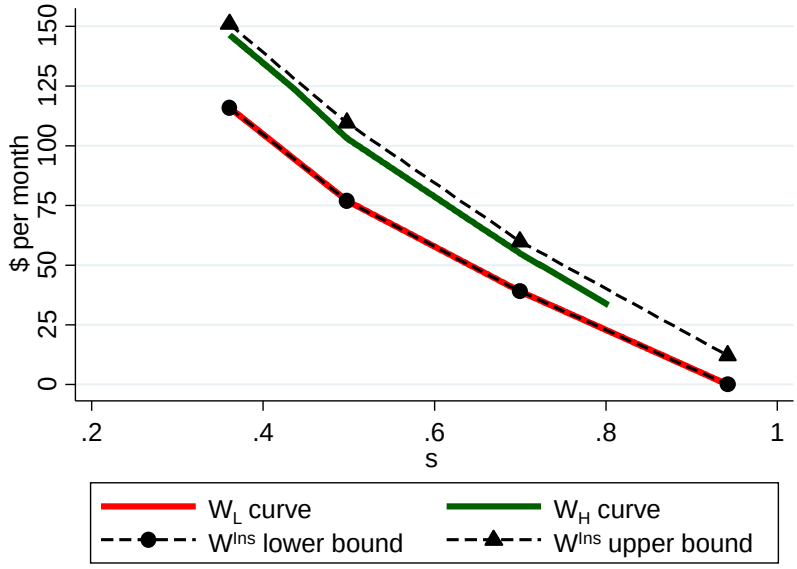
We can now map these bounds into bounds on a W^{Ins} curve. Let s be an index that orders people according to decreasing W_i^{Ins}. Without loss of generality, let the distribution of s be uniform on $[0, 1]$. Let $W^{Ins}(s)$ denote the WTP for insurance for someone with index s (i.e., the $(1-s)$th quantile of WTP). Suppose that at a given vector of premiums we observe that $1 - s^*$ share of people choose U, while s^* choose formal insurance. For the marginal type s^*, both conditions (10) and (9) hold, so we can say that:

$$p^{min} \leq W^{Ins}(s^*) \leq p^{max}$$

We use this result, along with our variation in premiums, to estimate bounds on the $W^{Ins}(s)$ curve. Specifically, we use the same income discontinuities in premiums discussed above. At each side of the discontinuity, we measure p^{min} and p^{max} and estimate $1 - D_U$. We then plot $1 - D_U$ (as x-values) against the bounds $\left\{p^{min}, p^{max}\right\}$ (as y-values) using the points on either side of the discontinuity. As with our estimates of W_L and W_H, we implement this exercise at each income level separately and then shift the curves horizontally to line up with the curve for 150% of FPL.

Appendix Figure 18 shows the resulting bounds for W^{Ins}, with our baseline estimates of W_L and W_H from the vertical model shown for comparison. The estimated bounds are relatively tight and quite close to the W_L and W_H curves. Indeed, the lower bound on W^{Ins} is identical to W_L by construction; both are generated by plotting the share purchasing formal insurance versus the premium of the cheapest plan (L). The upper bound on W^{Ins} is above W_H, but only slightly higher – a result

Figure 18: Bounds on Willingness to Pay for CommCare Insurance, 2011



NOTE: Figure shows our estimated bounds for WTP for CommCare (W^{Ins}), whose construction is described in the text of Appendix D. The lower and upper bounds for $W^{Ins}$are shown in black dashed lines (with the point estimates shown in circles and triangles, respectively). For comparison, the graph shows our baseline estimates of W_L (solid red line) and W_H (solid green line) that were derived using the vertical model.

that does not occur by construction but reflects the fact that the premiums of the H plans are quite similar and that few people choose L. From this exercise, we conclude that our basic estimates of (low) willingness to pay for CommCare insurance are robust to relaxing the vertical model assumptions.

Relaxing the vertical model assumptions for estimating costs is more challenging. Intuitively, our vertical model assumes we can pool the four non-CeltiCare plans into a single composite H option for which a type-s individual has a single expected cost $C_H(s)$. As premiums for H increase slightly, individuals of a single s type (s^*_{HL}) leave the plan, and we can estimate $c_H(s^*_{HL})$ using average costs before and after the change. However, it is also possible that some individuals may switch *among* the plans within H as premiums change. If we weaken the composite plan assumption, this switching could have an independent effect on AC_H and TC_H. In practice, however, we expect any bias to our cost estimates from any switching among H plans to be small. First, there is little reason to expect significant switching, since the premiums of the H plans are nearly identical to each other on both sides of the income thresholds (see Appendix Table 5). Second, given the similarity of $C_H(s)$ and $C_L(s)$ (see Figure 12), it seems unlikely that cost differences among the (much more similar) H plans for a given s type would be large.

Appendix E: Deriving the MVPF for additional subsidies

We derive the expression for the marginal value of public funds (MVPF) for additional subsidies for health insurance shown in equation (8). The MVPF is defined as the ratio of the marginal benefit of

an additional subsidy to the marginal cost. The marginal benefit of the additional subsidy is defined by the marginal willingness to pay for the additional subsidy. The marginal cost to the government of the additional subsidy incorporates any fiscal externalities resulting from a behavioral response to the subsidy; in this case, this comes from an enrollment response to the subsidy.

To derive the MVPF, we define aggregate welfare of beneficiaries when premiums p_H are set such that a fraction s^* of the market is enrolled as:

$$V(s^*) = \int_0^{s^*} (W_H(s) - p_H) ds$$

where $p_H = W_H(s^*)$ is the price required for a fraction s^* to choose to enroll. The term $W_H(s) - p_H$ captures the consumer surplus of an individual s enrolled in the H plan; this surplus equals her willingness to pay $W_H(s)$ minus the premium she pays, p_H. We define the cost to the government of enrolling a fraction s^* as:

$$\chi(s^*) = \int_0^{s^*} (C_H(s) - p_H) ds$$

where $C_H(s) - p_H$ denotes the net-of-premium cost of enrolling an s type.

We can then use these two expressions to define the marginal value of public funds when s^* of the market is enrolled:

$$MVPF(s^*) \equiv \frac{\frac{-d}{dp_H} V(s*)}{\frac{-d}{dp_H} \chi(s*)}$$

and to derive:

$$MVPF(s^*) = \frac{s^*}{s^* + \frac{C_H(s^*) - W_H(s^*)}{-W_H'(s^*)}} = \frac{1}{1 + \frac{C_H(s^*) - W_H(s^*)}{-s^* W_H'(s^*)}}.$$

Consider first the expression for the numerator: $\frac{-d}{dp_H} V(s^*) = s^*$. Lowering the premium by a dollar is valued by \$1 for the s^* fraction of the market already purchasing insurance. In addition to the transfer to the existing (inframarginal) enrollees, the market also expands in response to the subsidy by $\frac{ds^*}{-dp_H} = \frac{1}{-W_H'(s^*)}$, or in other words by the inverse of the slope of the demand curve. But the value of insurance to the marginal enrollee is by definition equal to the premium (i.e. $W_H(s^*) \equiv p_H$), so that the marginal benefit of insurance for the marginal enrollee is 0. Thus by the standard envelope theorem argument, the marginal benefit term consists only of the transfer to existing enrollees.

Consider next the expression for the denominator: $\frac{-d}{dp_H} \chi(s^*) = s^* + \frac{C_H(s^*) - W_H(s^*)}{-W_H'(s^*)}$. The first term, s^*, follows from the fact that $\frac{dW_H(s^*)}{dp_H} = -1$, so that the s^* inframarginal enrollees cost \$1. The second term follows from differentiating the extremum of the integral, which accounts for the cost of new enrollees: $\frac{ds^*}{dp_H} [C_H(s^*) - W_H(s^*)]$, where we use the fact that $p_H \equiv W_H(s^*)$ and $\frac{ds^*}{-dp_H} = \frac{1}{-W_H'(s^*)}$. The second term arises because the subsidy expands the market by $\frac{ds^*}{-dp_H} = \frac{1}{-W_H'(s^*)}$ and the marginal enrollees impose a cost of $C_H(s^*) - p_H$.

Thus the key term for determining whether the MVPF is above or below one is whether the cost of the marginal enrollee $C_H(s^*)$ is lower or higher than their willingness to pay ($W_H(s^*)$. If the cost

Table 5: Premiums by CommCare Plan

CommCare Plan	H / L Plan	Insurer Price (pre-subsidy)	Enrollee Premium (post-subsidy) by Income Group (% of FPL)			
			100-150%	150-200%	200-250%	250-300%
BMC HealthNet	H	$425	$11	$57	$105	$146
Fallon	H	$426	$12	$60	$110	$151
Neighborhood Health Plan (NHP)	H	$426	$12	$60	$110	$151
Network Health	H	$423	$10	$57	$105	$146
CeltiCare	L	$405	$0	$39	$77	$116

NOTE: The table shows enrollee premiums (by income group range) for each CommCare insurer in the market in fiscal year 2011. The top four plans – which we pool into an "*H*" plan in our analysis – all have very similar premiums because their (pre-subsidy) price bids were nearly identical, having been constrained by a binding price ceiling.

of the marginal enrollee is above (below) the willingness to pay of the marginal enrollee, the marginal subsidy increase will be welfare decreasing (increasing). It is important to note that, *a priori*, the sign is unclear. In particular, if adverse selection were the sole or primary source of low take-up – so that willingness to pay exceeded own costs for the marginal enrollee – subsidies for insurance financed through lump sum taxes on enrollees could increase aggregate welfare (Einav and Finkelstein, 2011).

Appendix F. Additional Results

Table 6: Summary of Estimates, 2011

Variable	150% FPL			200% FPL			250% FPL		
	Below	Above	Δ	Below	Above	Δ	Below	Above	Δ
Sticker Premium (Monthly)									
P_U (Expected)	$9.5		--	$28.5		--	$48.0		--
P_L	$0.0	$39.0	+$39	$39.0	$77.0	+$38	$77.0	$116.0	+$39
P_H	$11.0	$57.9	+$47	$57.9	$106.3	+$48	$106.3	$147.3	+$41
P_H - P_L	$11.0	$18.9	+$8	$18.9	$29.3	+$10	$29.3	$31.3	+$2
Normalized Premium (Monthly)									
P_L	-$9.5	$29.5	+$39	$10.5	$48.5	+$38	$29.0	$68.0	+$39
P_H	$1.5	$48.4	+$47	$29.4	$77.8	+$48	$58.3	$99.3	+$41
Market Shares									
Any Insurance	0.94	0.70	-0.24	0.76	0.56	-0.20	0.58	0.44	-0.14
H Plan	0.80	0.64	-0.16	0.70	0.50	-0.20	0.52	0.39	-0.12
L Plan	0.14	0.06	-0.08	0.06	0.06	0.00	0.06	0.05	-0.01
Average Cost (Monthly)									
Any Insurance	$333	$380	+$47	$355	$386	+$31	$328	$343	+$15
H Plan	$361	$393	+$32	$371	$405	+$35	$345	$365	+$20
L Plan	$169	$242	+$73	$170	$225	+$55	$175	$153	-$22

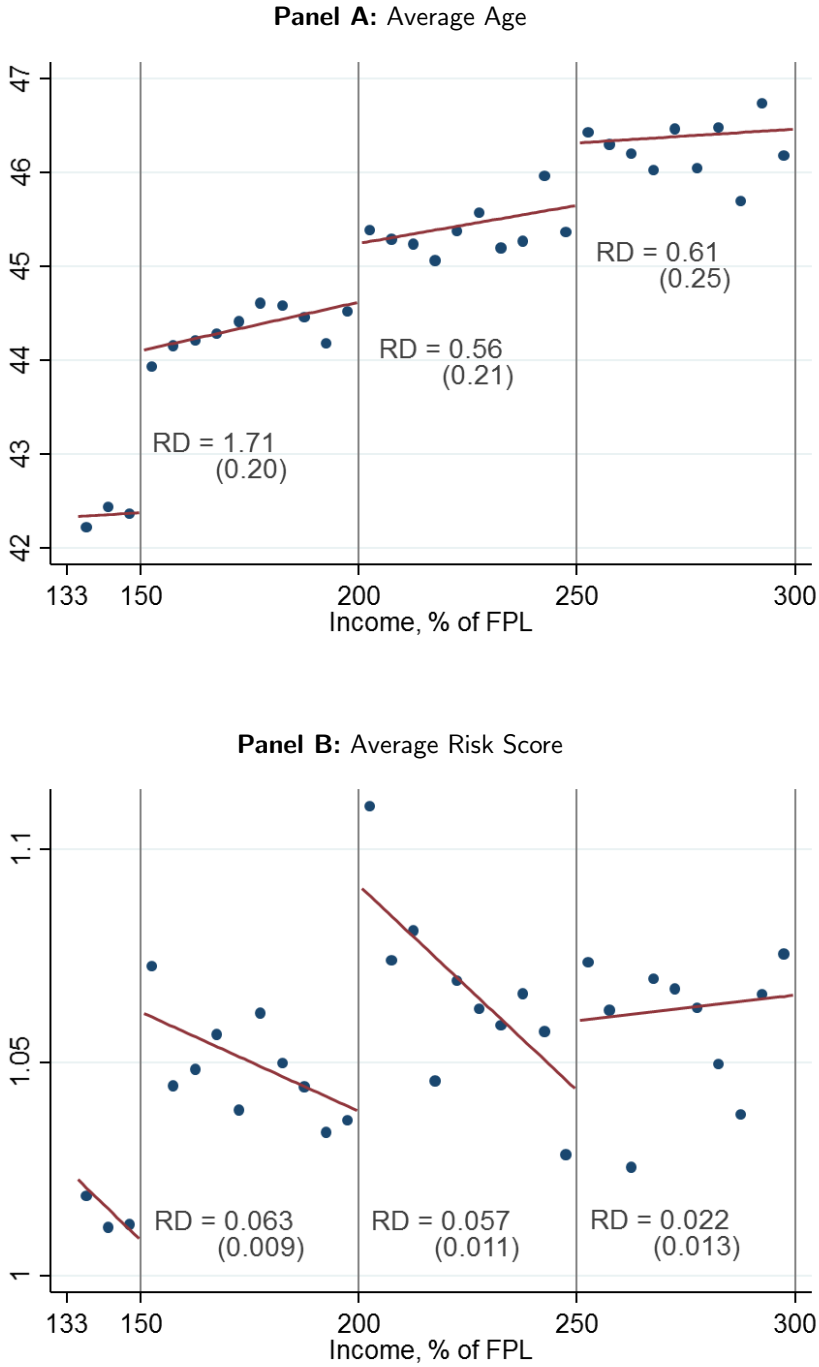
NOTE: This table summarizes the inputs into our estimates of willingness to pay and cost curves. For each income threshold at which premiums change, table shows the monthly enrollee premium, estimated market share, and average monthly insurer costs just below and above the threshold, as well as the change across the threshold. The premiums reported in the first two panels were previously presented in Figure 2. The first panel shows sticker premiums for U (the expected mandate penalty), the L plan, H plan, and the difference $p_H - p_L$; we use these for our main demand estimates. The second panel shows "normalized" premiums, where p_U has been normalized to $0, which we use for robustness analysis. The third and fourth panels report changes in market shares and average insurer costs based on RD estimates of equation (1); these results for any insurance and H plan were previously shown in the main text in Figures (6) and (7), respectively. The results for the L plan are in Appendix Figure 20.

Table 7: Willingness to Pay and Costs ($ per month)

Point in WTP Distribution		WTP		Costs	
		$W_L(s)$	$W_H(s)$	$C_H(s)$	$AC_H(s)$
Min In-Sample	(s = 0.94)	$0	$5	$148	$334
20th Percentile	(s = 0.80)	$23	$34	$203	$362
40th Percentile	(s = 0.60)	$58	$79	$291	$400
Median	(s = 0.50)	$77	$103	$333	$417
60th Percentile	(s = 0.40)	$105	$135	$369	$434
Max In-Sample	(s = 0.31)	$131	$162	$399	$448

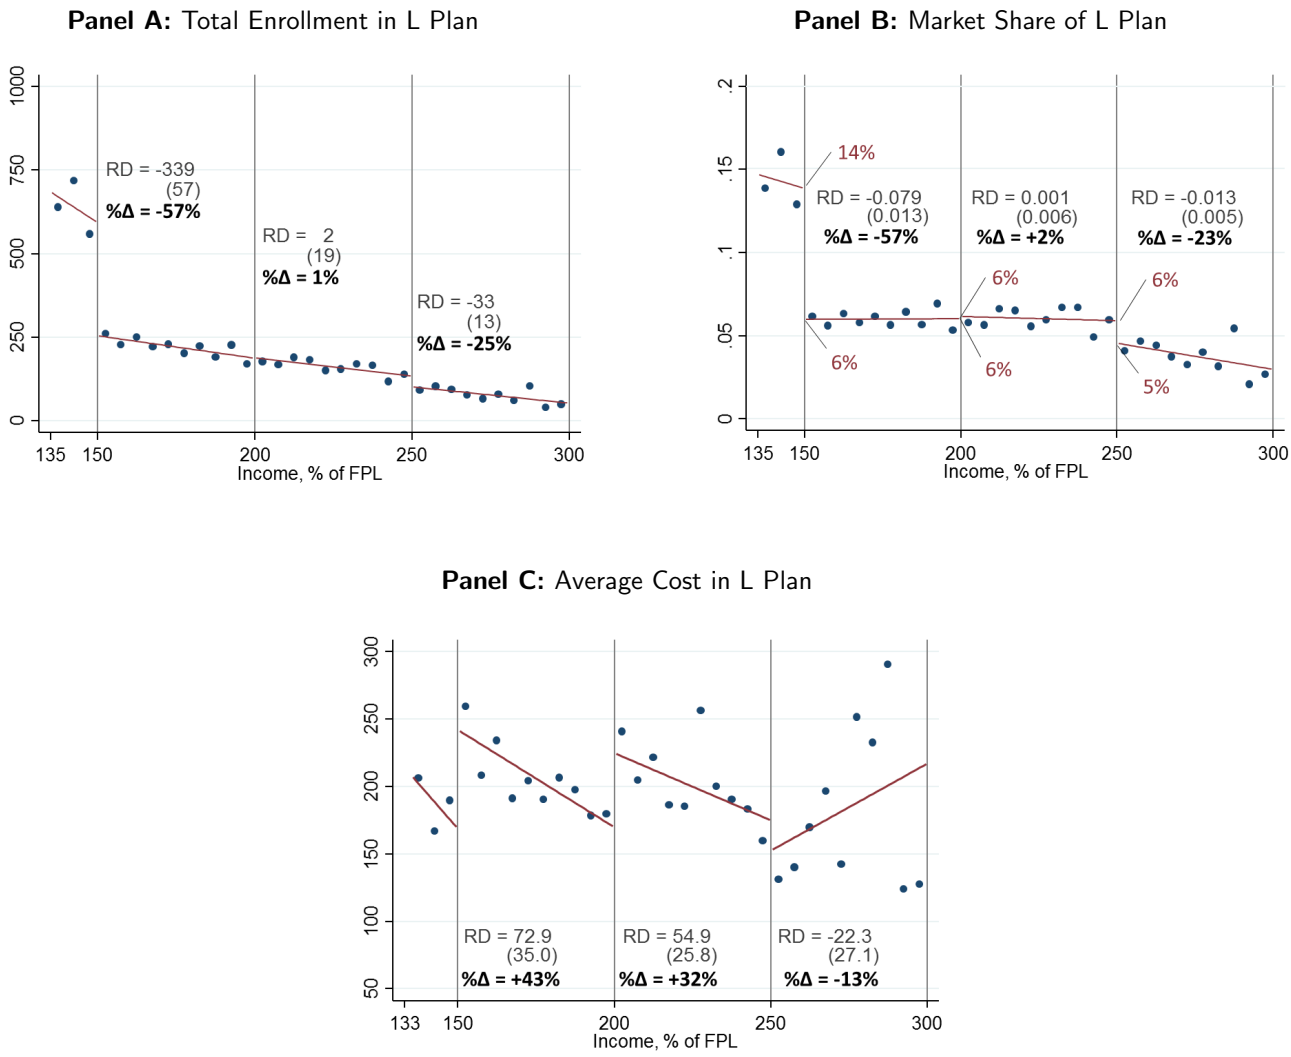
NOTE: Table summarizes our estimates of willingness to pay and costs for individuals at 150 percent of the FPL shown in Figure 12.

Figure 19: RD for Age and Risk Scores for Enrollees in All Plans, 2009-2013



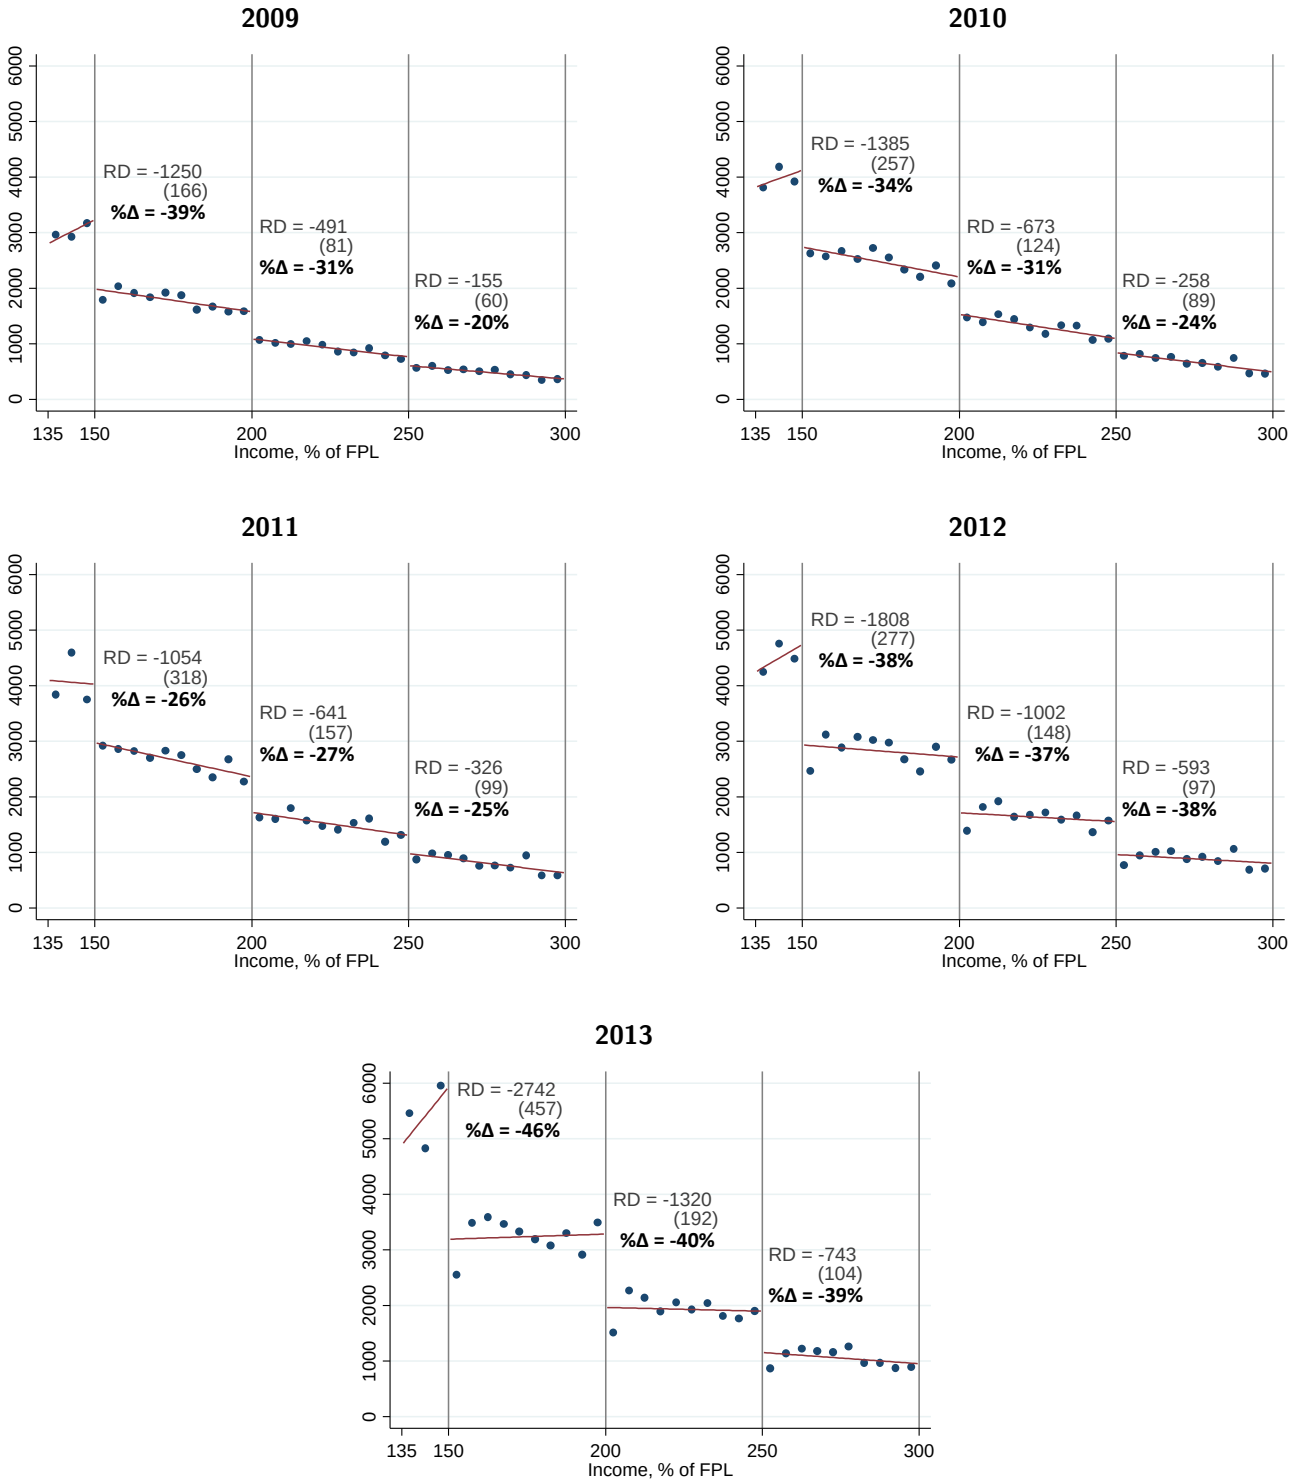
NOTE: These graphs show RD estimates for the average age (panel A) and risk score (panel B) of CommCare enrollees in all plans, pooled over the 2009-2013 period of our data. Risk scores are calculated by CommCare to reflect a person's expected medical spending based on their age, sex, and medical diagnoses. They are used by CommCare to adjust payments to insurers for their enrollees. Risk score values are relative to an average enrollee (whose risk score is 1.0) – e.g., a risk score of 1.05 indicates expected costs 5% above average.

Figure 20: RD Estimates for L Plan, 2011



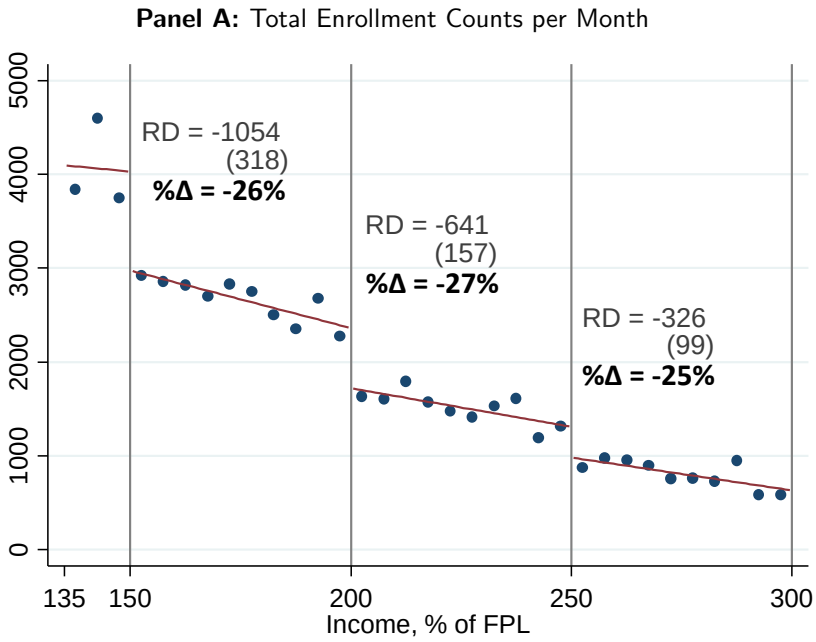
NOTE: Figures show our RD estimates for total enrollment, market shares, and average costs for the L plan in 2011, analogous to the estimates for all plans and the H plan in Figures 5-7 of the main text.

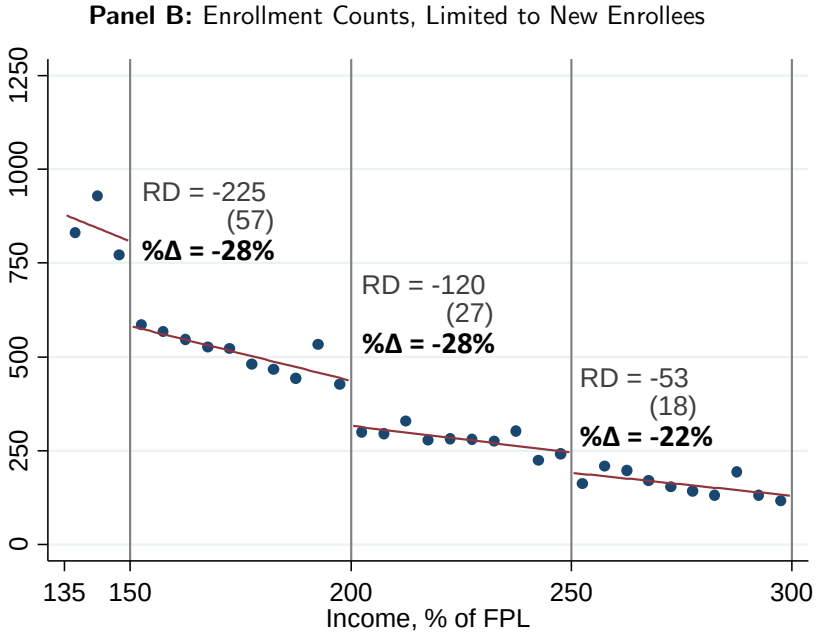
Figure 21: Total Enrollment Counts in CommCare, by Year



NOTE: The graph shows our baseline regression discontinuity analysis for total enrollment counts per month from Figure (4), Panel A (which showed results pooled for 2009-2013). Here, we show results separately by year.

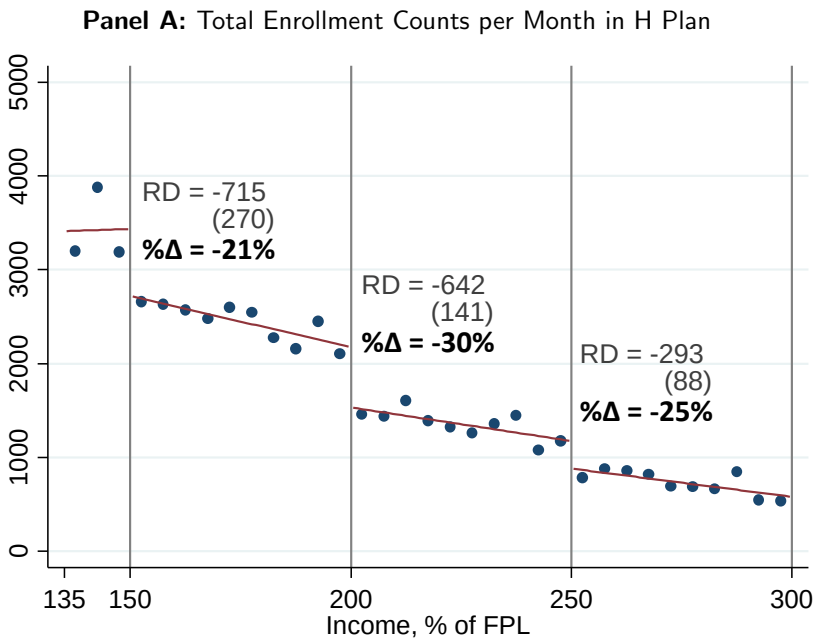
Figure 22: Enrollment Counts in 2011, by Income

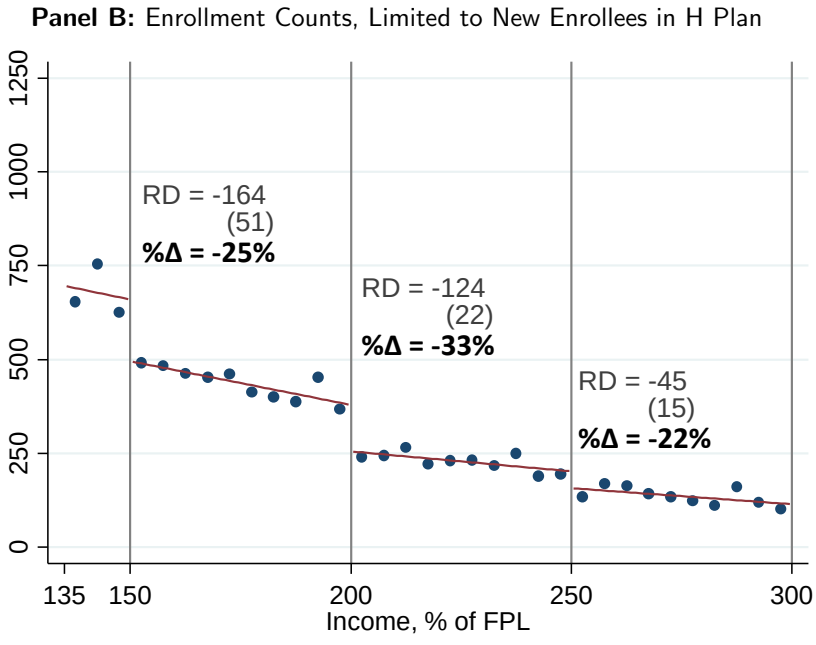




NOTE: These graphs are identical to Figure 21 but applied to counts of enrollees, but Panel B is limited to just new enrollees in CommCare during 2011.

Figure 23: Enrollment Counts in 2011 for H Plan, by Income





NOTE: These graphs are identical to Figure 22 but applied to counts of enrollees in the H plan only.