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

Faculty Research Working Paper Series

Amy Finkelstein
Massachusetts Institute of Technology

Nathaniel Hendren
Harvard University

Mark Shepard
Harvard Kennedy School

April 2017

RWP17-017

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

Amy Finkelstein, Nathaniel Hendren, Mark Shepard*

April 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 three to four times below *own* expected medical 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, and even premiums subsidized down to 10% of average costs would still leave at least 20% uninsured. We briefly consider explanations for this finding – which suggests an important role for uncompensated care for the uninsured – and explore normative implications for insurance subsidies for low-income individuals.

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 and Ray Kluender for helping with several calculations. We thank the Massachusetts Health Connector (and particularly Marissa Woltmann) 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, Josh Goodman, Jeff Liebman, Brigitte Madrian and seminar participants at Harvard Kennedy School 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 this 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.[2] Partial subsidies are also a textbook policy response to adverse selection or other market imperfections in settings where 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 offers heavily-subsidized private plans to non-elderly adults below 300% of poverty who do not have access to insurance through an employer or another public program. Public subsidies are 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 is also a mandate to have health insurance, backed up by financial penalties.

We use a regression discontinuity design, together with administrative data on enrollment and medical costs, to estimate demand and cost for CommCare plans. 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. Our main analysis focuses on fiscal year 2011, when the insurance options had a convenient vertical structure, but we present some complementary results for the 2009-2013 period as well.

Our analysis begins by documenting two main descriptive patterns. First, enrollee demand is highly sensitive to premiums. With each discrete increase in enrollee premium, enrollment in CommCare falls by 20 to 24 percentage points, or about 25% of total enrollment. For example, we estimate that take-up is nearly complete (94%) for individuals below 150% of poverty for whom the cheapest plan is free. But

[1] See Department of Health and Human Services (2015, 2016)U.S. Department of Health and Human Services (2015), Department of Agriculture (2016)U.S. Department of Agriculture (2016), and Internal Revenue Service (2015)U.S. Internal Revenue Service (2015).

[2] As of 2017, 30 states charge premiums or enrollment fees for some Medicaid/CHIP enrollees above 133% of poverty (Brooks et al., 2017). In addition, five states (Arkansas, Indiana, Iowa, Michigan, and Montana) have received federal waivers allowing them to charge premiums to enrollees newly eligible under the ACA (Smith et al., 2015).

take-up falls to just 70% when the enrollee premium for this cheapest plan rises to \$39 per month just above 150% of poverty – even though this still represents a subsidy of more than 90% off of the price charged by the insurer. 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.

We use a simple model to analyze the implications of these empirical 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").[3] Using this vertical model and the framework of Einav et al. (2010a), we show how to map our data on prices and market shares into WTP curves for H and for L, and average cost and marginal enrollee cost curves for H – which is chosen by the vast majority of enrollees.

The model allows us to translate the descriptive patterns into two main results. First, even large insurance subsidies are insufficient to generate near-complete take-up of insurance by low income adults. At the median of the WTP distribution, WTP for H is about \$100 per month – less than one quarter of average costs of \$420 per month if all those with above-median WTP 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 H. Subsidies making enrollee premiums 10% of insurers' average costs still leave at least 20% uninsured.

Second, although adverse selection exists, it is not the primary driver of low take-up. Because of adverse selection, the cost of marginal consumers who enter as premiums decline is less than the average costs of inframarginal consumers (Einav et al., 2010a). But for our entire in-sample distribution – which spans the 6th to the 70th percentile of the WTP distribution – the WTP of marginal enrollees still lies far below their own expected costs imposed on insurers for either the H or L plans. For example, a median WTP individual imposes a cost of \$340 on the insurer for the H plan, but is willing to pay only about \$100 for the H plan. This suggests that textbook subsidies offsetting adverse selection would be insufficient for generating take-up for at least 70% of this low-income population. Coverage is low not simply because of adverse selection but because people are not willing to pay their own cost they impose on the insurer.

We demonstrate the robustness of these two main results to a number of econometric and modeling choices. In the final section of the paper, we briefly explore potential explanations for our findings and analyze their normative implications. A back-of-the-envelope calculation suggests that the substantial uncompensated care available to the uninsured (as documented in previous literature) can account for nearly all of the gap between willingness-to-pay and costs of insurance. Low-income adults may be willing to pay their own "net costs" of formal insurance, which equals the cost to the insurer minus any crowd-out of uncompensated care.[4]

[3] In practice, the H and L options represent groups of CommCare insurers that cover the same (state-specified) benefit package but differ in their networks of medical providers. The L plan has a narrower network (and lower premium), and accordingly, we find that enrollees have lower willingness to pay for it. However, our main results focus on demand for formal insurance (in H or L) versus uninsurance (U), since this is the main pattern of substitution that we see in the data as subsidies decline.

[4] Estimates suggest that the uninsured pay about 20 to 35% of their cost of care (Coughlin et al., 2014; Hadley et al.,

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 a significant fraction of their expected cost of coverage. This suggests that subsidies in this market cannot be justified simply as a response to adverse selection and, more broadly, that the search for market surplus and welfare benefits of health insurance subsidies needs to look beyond the enrollees themselves. Our results suggest two potential such justifications for subsidies: as an offset to the "tax" that uncompensated care imposes on formal insurance (i.e. the Samaritan's dilemma (Buchanan, 1975; Coate, 1995)), or as a means of redistribution to low-income households.

Related Literature Our results relate to a growing literature estimating the demand and costs for health insurance. However, there is relatively little work on health insurance demand and costs among low-income adults, since until recently, most of the low-income uninsured were either not offered health insurance or faced prices that were difficult to measure.[5] The ACA's introduction in 2014 has given researchers an opportunity to study how low-income insurance take-up responds to subsidies (Frean et al., 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 WTP for formal insurance compared to insurance costs are consistent with Finkelstein et al. (2015), who use *ex-post* outcome distributions (instead of revealed preference) 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 et al. (2015) study the *unsubsidized* Massachusetts health insurance exchange for individuals *above* 300% of poverty. They also find evidence of adverse selection but estimate a WTP curve for insurance that *exceeds* own costs over the entire population of potential consumers. By contrast, we find that at least 70% of low-income individuals in the same state have WTP far below own costs. One simple way to reconcile the divergent findings is that low-income individuals likely have much greater access to uncompensated care, which is crowded out by formal insurance. Indeed, a growing empirical literature has documented the large role of uncompensated care for the (predominantly low-income) uninsured and the impact of insurance in decreasing unpaid bills (see e.g., Finkelstein et al., 2012; Mahoney, 2015; Garthwaite et al., 2015; Dobkin et al., 2016; Hu et al., 2016).[6]

Our results also have broader implications for health insurance design and recent health insurance

2008; Finkelstein et al., 2015), which is remarkably similar to our estimated ratio of WTP to own costs.

[5] 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.

[6] Differences in the availability of uncompensated care may also help reconcile our findings with results from a calibrated life cycle model suggesting that the low-income elderly's WTP for Medicaid is above their costs (De Nardi et al., 2016); unlike for 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.

reforms. The price sensitivity we observe suggests 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 where income-specific premiums are significantly higher than in CommCare – 2-10% of income in the ACA versus 0-5% of income in CommCare.[7] Thus, our results may help understand coverage outcomes in the ACA exchanges, where early evidence suggests highly incomplete take-up (Tebaldi (2016); Avalere Health (2016)). They also generate predictions of coverage under alternate reform proposals. More generally, our results suggest a fundamental challenge in enrolling low-income people into health insurance markets, even with an insurance mandate: take-up is low not simply because of adverse selection but because people are not willing to pay the (gross) cost of coverage they impose on the insurer.

Finally, our results have implications for understanding the impact of adverse selection in health insurance markets – which is often highlighted as a key economic rationale for subsidies and mandates. The empirical literature has extensively documented the presence of adverse selection in health insurance markets but concluded that the welfare cost of the resultant mispricing of contracts is relatively small. This literature however has "looked under the lamppost" – primarily focusing on selection across contracts that vary in limited ways, rather than selection that causes a market to unravel, leaving open the possibility of larger welfare costs on this margin (Einav et al., 2010b). 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 et al., 2011; Hackmann et al., 2015). But it also finds that adverse selection is not the full story and that other factors must be important in explaining limited demand for formal coverage 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 WTP and costs for insurance; we document implications for take-up of subsidized insurance and the role of adverse selection. Section 5 briefly considers potential explanations for low WTP and potential normative implications of our findings for the desirability of subsidies for health insurance 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. It laid the foundation for many of the sub-

[7]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 plans.

sidized 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 (which ran from July 2010 to June 2011) but also present some descriptive results from a larger sample comprising the entire set of fiscal years (2009-2013) for which the necessary data are available. Except where otherwise indicated, the market rules described below apply to 2011; the rules for other years in the 2009-2013 period are similar but differ 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). Eligibility for CommCare is thus similar to the criteria for receiving subsidies on the ACA exchanges. Given Medicaid eligibility rules in Massachusetts,[8] 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.[9] About 70% of CommCare enrollees have incomes below 150% FPL (about $16,000 per year for a single individual).

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 to cover these benefits (the specified services and consumer cost-sharing rules) but could differ in its network of hospitals and doctors. In each year from 2009-2013, between 4 and 5 insurers participated in the market. Benefit design and participating insurers were very similar to the Massachusetts Medicaid program. In particular, CommCare enrollees faced very modest co-pays, similar or slightly less generous than in Medicaid.[10]

Eligible individuals could enroll at any time if they experienced a qualifying event (e.g., job loss or income change) or for any reason during the annual open enrollment period at the start of the fiscal year. To enroll, individuals filled out an application with information on age, income, 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. If approved for CommCare, individuals were notified (by mail and/or email) and 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.

[8] 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.

[9] In addition, a special "CommCare Bridge" program covered certain legal immigrants ineligible for Medicaid with incomes up to 300% of FPL. Our data exclude this group.

[10] 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 that our main results are not sensitive to excluding this discontinuity.

Subsidy Structure

Insurers in CommCare set a base price for their plan 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 of their enrollee pool.

Enrollees pay premiums (to CommCare) equal to their insurer's base price minus an income-varying subsidy paid by the state.[11] 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 then increases to $39 per month above 150% FPL, to $77 per month above 200% FPL, and to $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 comprised a small share of income. For instance, for a single individual in 2011 (whose FPL equaled $908 per month), these premiums range 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). We will use the 2009-2013 data to analyze how the number of enrollees and the average costs of those enrolled change with these discrete increases in enrollee premiums.

2011 Plan Options

We will also undertake more detailed analyses of the market in fiscal year 2011, when a relatively simple market outcome facilitates identification of demand and cost curves. Specifically, the market in 2011 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 lone 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 somewhat lower. As discussed above, insurer prices are constant across incomes, while all enrollee premiums jump (as subsidies decrease) at 150%, 200%, and 250% of FPL.

Along with its lower price, CeltiCare also had had a much 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.[12] 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.[13]

[11] 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.

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

[13] A piece of evidence consistent with this perception comes from enrollees below 100% of FPL, who (distinct from the subsidy structure for above-poverty enrollees) were fully subsidized and could choose among all available plans for free.

Figure 1: Insurer Prices and Enrollee Premiums in CommCare Market

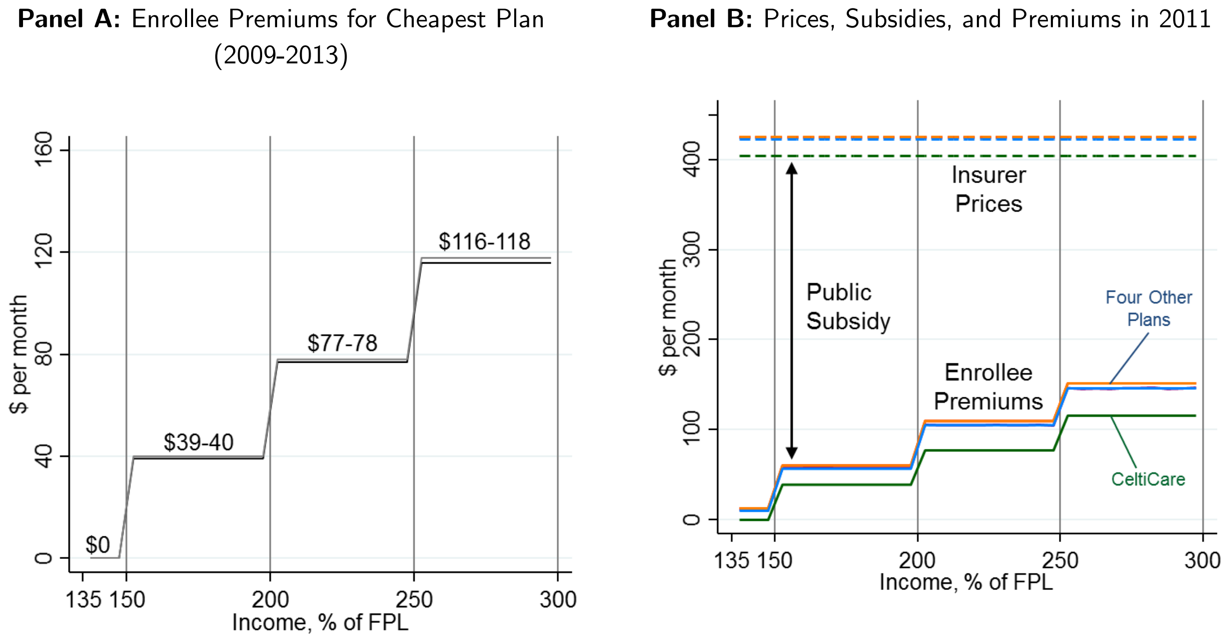

NOTE: The figure illustrates CommCare's insurance subsidy structure. As described in the text, CommCare sets subsidies so that the post-subsidy enrollee premium for the cheapest plan equals a target "affordable amount" for each income group. Panel A plots these affordable amounts by income as a percent 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 CommCare in the single year of 2011. As described in the text, each insurer sets a single price that applies across all income groups. In 2011, there was a price cap at $426/month, and four insurers set prices within $3 of this cap, while CeltiCare set a lower price ($405). Public subsidies are applied to reduce these prices to the premiums that enrollees pay. These subsidies vary by income within 50% of federal poverty level (FPL) bins, leading to discontinuous changes in enrollee premiums at income thresholds of 150%, 200%, and 250% of FPL, which are marked on the graph.

As a result, in much of our analysis that follows we pool the 2011 plans into two groups: CeltiCare as a low coverage "(L)" option and the other four plans (pooled together) as a high coverage "(H)" option. 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 feel comfortable pooling the four H plans offered in 2011 because their prices were so similar (a fact not true in other years), allowing us to interpret H as a composite contract that gives enrollees a choice among the four component insurers. The utility of the H option should be interpreted as the max of the utility over these four insurers. 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 6 reports enrollee premiums separately for each component plan. As previously discussed, enrollee premiums for the cheapest plan L ("p_L") – set to equal a target affordable amount – jump at

Among this group, 96% of enrollees chose a plan other than CeltiCare.

Figure 2: Premium and Mandate Penalty Variation, 2011

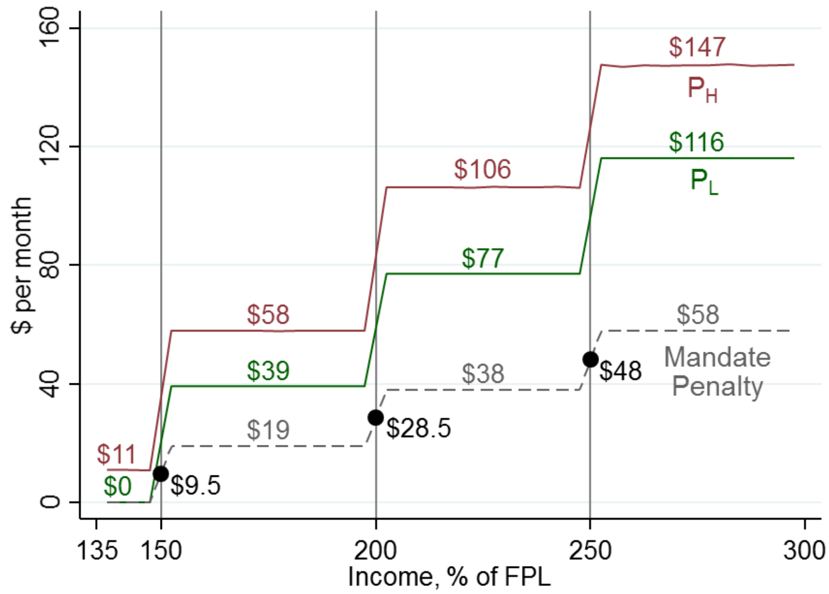

NOTE: The figure shows how enrollee premiums and the mandate penalty vary across incomes (as a percent of the federal poverty level, FPL). All numbers are for the CommCare market in fiscal year 2011. "P_L" denotes the enrollee premium for the L plan (CeltiCare), while "P_H" is the share-weighted average of the enrollee premiums in the four H (non-CeltiCare) plans. Appendix Table 6 shows how prices for each individual non-CeltiCare plan vary with income. "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. Because the mandate penalty is assessed based on total annual income, a person near the discontinuity is equally likely to end the year with income in the higher or lower range; therefore, the black dots are averages of the two mandate values.

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 for lower-income groups. Importantly for our demand estimation, this subsidy structure creates variation in both the level of premiums and premium 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). However, in practice, the actual penalty an individual in our data would owe likely diverges from the gray line for two reasons. First, the mandate is assessed based on total annual income (reported in tax filings), whereas the measure in our data that is used to determine enrollee premiums is from income reported on the program application (which is an estimate of current monthly income). 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.[14]

It is unclear how to use the mandate penalty when calculating revealed willingness to pay. For the reasons discussed above, the 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 prices of 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 the 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 the 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.[15]

We observe each enrollee's chosen plan at a monthly level. We define plan enrollment as the annualized number of enrollee months in the 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 (as summarized in Figures 1 and 2). Enrollee-linked insurance claim data allows us to directly 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 one decimal point), which is the running variable for our RD analysis. This variable is calculated by the regulator from information on family income and composition that enrollees report in their initial CommCare application. 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.[16]

[14]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 months or less (Connector and of Revenue, 2011).

[15]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).

[16]The risk score is a measure of predicted costs derived by the regulator based on enrollee's age, sex, and observed diagnoses. It is based on data from the previous 12 months and updated quarterly. For new enrollees, it is based only on age and sex. that is used for risk adjusting payments to insurers for enrollees. It is used by CommCare to adjust payments to insurers for their enrollees.

Survey Data: Eligible Population for CommCare

To translate our estimates of the number enrolled into the fraction of eligibles who take up insurance, we supplement the administrative data with the 2010 and 2011 American Community Survey (ACS), an annual 1% random sample of U.S. households.[17] This provides an estimate of the size of the population eligible for CommCare, as a function of family income as percent of FPL. Specifically, we use data from Massachusetts to estimate the number of people whose observables make them eligible for CommCare: i.e., they are a U.S. citizen, age 19-64, have family income below 300% of FPL, 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 further 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.

The ACS data give us information on both the number of eligible people and the shape of their distribution over the income range we examine. 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 income 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 smoothing regression fit.

Among those who we estimate to be eligible for CommCare in the ACS data, we find that 63% of these individuals report having insurance in the ACS data. This statistic also closely matches estimates from tax filing data.[18] However, 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. We believe this latter estimate is too low, arguably the result of measurement error in the classification variables used to generate the eligible population – for example it is plausible that some individuals mis-report the type of insurance they have (e.g. they don't know they actually have Tricare and therefore accidentally appear eligible in the ACS but are not).

Therefore, we proceed under the assumption that 63% of eligibles take up CommCare. We take the initial (smoothed) ACS estimates 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%. This method effectively uses the ACS to estimate two statistics: the shape of the eligible income distribution and the average take-up rate. But, it does not use the ACS for the total size of the eligible population.

Measuring the eligible population is difficult, and our approximation is, of course, imperfect. For-

[17] We obtained ACS data from the IPUMS-USA website (Ruggles et al., 2015). We average estimates over the 2010 and 2011 ACSs because CommCare's fiscal year 2011 runs from July 2010 to June 2011.

[18] Specifically, 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. We therefore use the within-ACS estimate of 63% take-up of CommCare among the eligible population.

tunately, 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 admin data alone), we can generate our key result: that WTP is far below costs for marginal enrollees who drop coverage at each discontinuity (and the average cost of remaining enrollees).[19] 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. In the sensitivity analysis in Section 4.4below, therefore, we show how our results change if we use the raw ACS estimates of the eligible population size (without re-scaling). This produces a substantially lower take up rate, which in turn yields even lower estimates of the share of the low-income population that would be insured under a given subsidy scheme.

Summary Statistics

Table 1 reports some summary statistics. The table, and our subsequent analyses focus on individuals enrolled in or eligible for CommCare above 135% of FPL.

The first two columns summarize the CommCare-eligible population in the ACS – both the full eligible population (column 1) and the subset who report being insured (column 2). The next three columns summarize CommCare enrollees (in any plan, in H, and in L) based on the administrative data. Panel A shows sample sizes and estimated population sizes, which we have just discussed. The remainder of the table shows average demographics and (for enrollees) premiums and costs. In theory, both column 2 (the eligible insured population in the ACS) and column 3 (enrollees in the admin data) represent the set of people enrolled in CommCare; in practice, they look roughly comparable but not identical.

Panel B allows us to compare demographics between all eligibles (column 1) and those who enroll (column 2 or column 3). We see that income as a share of FPL is lower among enrollees than all eligibles, consistent with greater take-up among lower-income groups who get larger subsidies. Average age is higher for enrollees than all eligibles, consistent with adverse selection into insurance. And males are much less likely to be enrolled in CommCare.

Panel C shows data on prices, premiums and costs. CommCare's subsidies are quite large. Average enrollee premiums ($70/month) covers less than 20% of insurer-paid medical costs ($359/month) or prices ($422).

The vast majority (about 90%) of CommCare enrollees are in one of the H plans, despite higher enrollee premiums (see Figure 2). Again consistent with adverse selection, these H enrollees are older and have much higher risk scores and medical costs than L enrollees.

[19]Intuitively, we can measure average cost and enrollment on either side of each discontinuity, which can be used to calculate marginal costs for the enrollees who drop coverage. These cost measures can be compared to WTP for the marginal enrollees, which we know is less than the higher premium (just above the discontinuity). We formalize this argument using our model in Section 4.

Table 1: Summary Statistics, 2011

	Eligible Poplation (ACS)		Enrolled in CommCare (Admin Data)		
	All Eligible (1)	**Insured** (2)	**Any Plan** (3)	**H plan** (4)	**L Plan** (5)
N (# of unique individuals)	2,856	1,849	107,158	96,391	14,828
Panel A: Populations					
Scaled Population estimate	99,151	62,096	---	---	---
Average monthly enrollment	---	---	62,096	55,599	6,497
Share of Eligible Population	100%	63%	63%	56%	7%
Panel B. Demographics					
Income					
Dollars (annual)	$29,629	$31,037	---	---	---
% of FPL	203.8%	200.1%	193.0%	193.5%	188.5%
Share below 150% FPL	15%	16%	20%	18%	29%
Household Size	1.98	2.23	---	---	---
Age	39.9	40.8	44.4	44.9	40.2
Male	50%	43%	41%	40%	47%
Risk Score	---	---	1.066	1.084	0.906
Panel C. Premiums and Costs					
Prices and Subsidies					
Insurer Price	---	---	$422.2	$424.2	$404.9
Enrollee Premium	---	---	$70.0	$72.8	$46.0
Public Subsidy	---	---	$352.2	$351.4	$358.9
Medical Costs					
Insurer-Paid	---	---	$358.5	$377.3	$197.9
Total	---	---	$377.3	$396.4	$213.4

NOTE: This table show summary statistics for CommCare enrollees and our estimates of the CommCare-eligible population in fiscal year 2011 for our analysis sample: eligible individuals between 135 and 300 percent of FPL. The first two columns show statistics for the eligible population, as estimated from the 2010-11 American Community Survey (ACS). The first column is for the full eligible population, while the second is for the subset of eligible individuals who have insurance (which in theory, should be CommCare). The next three columns show statistics from the CommCare administrative data for enrollees in any plan, an H plan, and the L plan. Entries are omitted where a variable is unavailable in one of the datasets.

3 Descriptive Analysis

3.1 Regression Discontinuity Design

We use the discrete change in enrollee premiums for H and L at 150%, 200% and 250% of FPL (recall Figure 2) 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 analyze the enrollment share data at the bin level (with robust standard errors), and the other outcomes at the enrollee-month level, and report robust standard errors, clustered at the bin level.

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 mis-report)[20] their income to get just below the thresholds and qualify for a larger subsidy.[21] In practice, two factors suggest that any such strategic income shifting, if it exists at all, is quantitatively unimportant.

The first is institutional: manipulation of income to end up just below the threshold would be challenging, as it would require that people know their income relative to the thresholds (which are not round numbers but percents of FPL). This involves a fairly sophisticated calculation for this population – they must understand the subsidy rules (including the definition of income used), know the FPL specific to their family size, and calculate the ratio of the two. None of this information is listed on the CommCare application form.

The second is empirical: we see little evidence of bunching either in the ACS income distribution or in the shape of CommCare's administrative enrollment data just below the threshold. The raw ACS income distribution are shown in Appendix Figure 14 for the pooled 2009-2013 data as well as the 2011 data separately. They show no evidence of such bunching, although are too noisy to be definitive. The administrative data on plan enrollment, which we turn to next, also shows no upward spike in the number of enrollees just below thresholds (or decline just above thresholds) as one would expect if there were strategic manipulation.

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

[21]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 – which 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. We explore this in the sensitivity analysis below.

3.2 Evidence from All Years (2009-2013)

3.2.1 Insurance Demand

As we showed in Figure 1, premiums increase discontinuously at 150, 200, and 250% of FPL in each of our years 2009-2013. To what extent does this reduce the fraction who buy insurance? Panel A of Figure 3 presents average enrollment in the CommCare market over the 2009-2013 period, by income bin. It also super-imposes the estimates from the linear RD model in equation 1, where the dependent variable is the number of CommCare enrollees in a given bin. The plot shows clear evidence that higher premiums lead to significant decreases in insurance take-up. At each of the three discontinuities (150, 200, and 250% FPL), enrollment falls by 31-36%, with each change statistically significant at the 1% level.

3.2.2 Cost of Insurance and Adverse Selection

Do the set of people who choose insurance after facing higher premiums impose a higher cost on the insurer? We use the same RD design to estimate the costs of insurance among those enrolled and how this varies with enrollee premium. We define the average cost of insurance as the average claims paid by the insurer for the set of people who enroll at a given price.

Panel B of Figure 3 plots the average cost for all CommCare enrollees by income bin. It also superimposes the estimates from the linear RD model in equation 1, where the dependent variable is average insurer costs among enrollees in a given bin. It shows that average costs of the insured rise as the enrollee premium increases – i.e. adverse selection. For example, we estimate a discontinuous increase in costs at 150% FPL of \$45.3 (s.e. \$7) per enrollee per month. We also see a robust increase in average costs at the 200% FPL threshold of \$33.6 (s.e. \$8.8). We find a smaller but noisier increase of \$5.9 (s.e. \$11.9) at 250% FPL, which is no longer statistically distinguishable from zero; this may reflect the smaller sample size of enrollees at this discontinuity.

Taken together, these patterns provide robust evidence of adverse selection: Intuitively, if average costs among the enrolled rise with premiums, this indicates that the marginal enrollees (who exit in response to the premium increase) are less costly than the average enrollee who remains. *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 the presence of a pre-existing condition); 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, was not sufficient to prevent adverse selection.

3.3 Evidence from 2011

The preceding results provide evidence that insurance demand declines as enrollee premiums rise and that there is adverse selection into CommCare despite the presence of a coverage mandate. To yield estimates of demand and costs for specific plans, in this section and most of the remaining paper, we study the data from fiscal year 2011 in which we have the vertical differentiation of the H and L

Figure 3: Discontinuities in CommCare Enrollment and Average Insurer Costs, 2009-2013

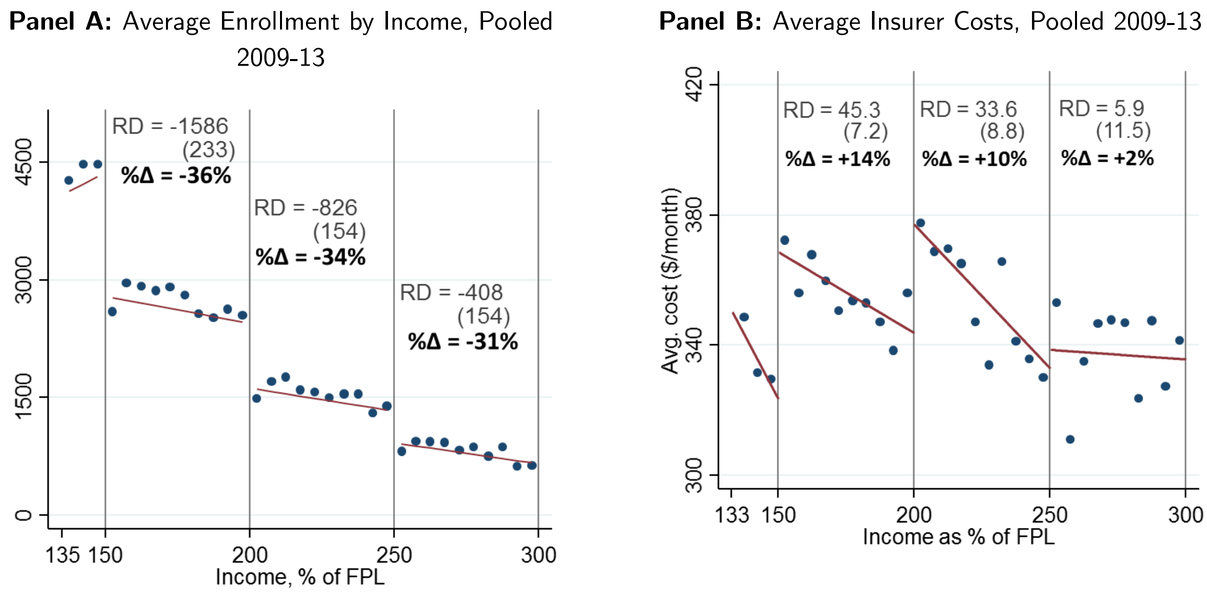

NOTE: The figure shows discontinuities in enrollment and average insurer costs at the income thresholds (150%, 200%, and 250% of FPL) at which enrollee premiums increase (see Figure 1). Panel A shows average enrollment in CommCare (total member-months, divided by number of months) by income over the 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, blue dots represent raw averages for a 5% of FPL bin, and the red lines are predicted lines from our linear RD specification. RD estimates and standard errors (in parentheses) are labeled just to the right of each discontinuity; percent changes relative to the value just below the discontinuity are shown in bold (labeled "%Δ =").

Figure 4: Raw Enrollment Counts, 2011

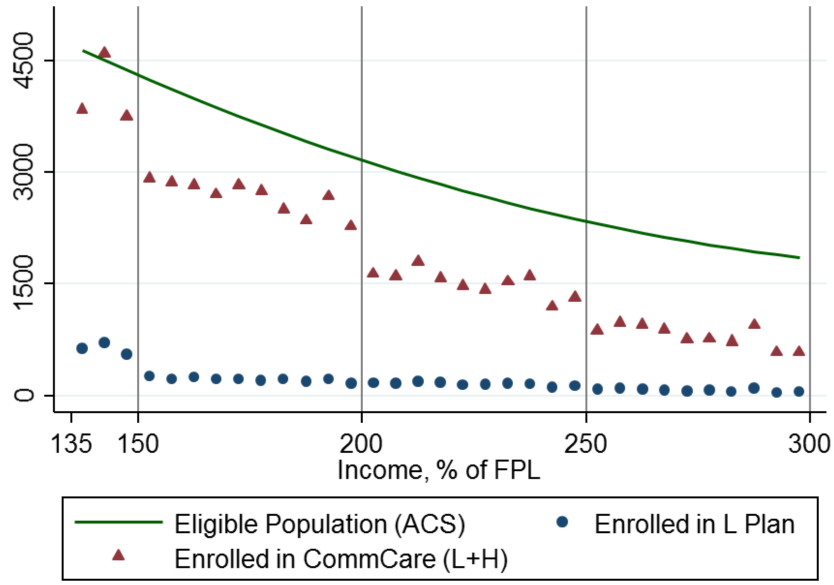

NOTE: Figure shows enrollment counts by income (in bins of 5% of the federal poverty level) in any CommCare insurance plan ($H + L$) and in in the L plan. It also shows our (smoothed) estimate of the CommCare-eligible population based on the American Community Survey (ACS). The construction of the eligible population size is described in Appendix A.

plans discussed above. Here, we present the reduced form evidence of demand and costs; in the next section, we illustrate how the vertical model structure allows us to translate these estimates into the willingness to pay and cost curves associated with the H and L plans.

3.3.1 Demand for Insurance

Figure 4 shows the raw counts of the number of people eligible for CommCare, enrolled in any CommCare plan, and enrolled in the L plan as a function of income (as a % of FPL). As described in Section 2, the estimate of the eligible population is based on a smoothed estimate from the American Community Survey (ACS), and the enrollment numbers are counts from the administrative data. The uninsured population U is the calculated as the difference between the eligible population and the number of CommCare enrollees.

There is evidence of sharp changes in enrollment in CommCare at the enrollee premium thresholds (150%, 200% and 250% of FPL, see Figure 2). There is also some evidence of a decline in the number in the L plan, particularly at 150% FPL. The eligible population is smooth (by construction) through the discontinuities; as we discussed above, there is no evidence of manipulation of the running variable in either the ACS or CommCare data.

Figure 5 transforms these raw counts into market shares, using the eligible population as the denominator. It also superimposes the estimates from the RD model in equation 1. In Panel A we report the share enrolled in any CommCare plan, and in panel B we report the share in the H plan (the share enrolled in plan L can be inferred from the difference between these two; we show the corresponding RD graph in Appendix Figure 16). The figures show clear drops in both enrollee shares at each pricing discontinuity, all of which are statistically significant at the 1% level. The percent changes in demand at each discontinuity are similar to (though somewhat smaller than) the estimates for the pooled 2009-2013 period.[22]

Several results from Figure 5 stand out. Panel A shows insurance coverage falling by about 25% when enrollee premiums rise by \$38-\$39 per month at each of the the three discontinuities; the drop in enrollment is not limited to when premiums go from zero to positive at the first (150% FPL) threshold. Note that the size of these *percent* drops are identified directly from the fall in enrollment in the administrative data. They do not depend on the eligible population size at the discontinuity that we approximate using the ACS data. For completeness, Panel A of Appendix Figures 18 and 19 show the RD in the enrollment counts directly.

If we also use the estimated eligible population size from the ACS, we have two additional findings. First, take-up appears to be nearly complete (94%) just below 150% FPL where insurance is free.[23] Second, take-up falls to 70% above 150% FPL where the cheapest premium increases to \$39 per month and continues to fall, ending up below 50% as the cheapest premium rises to \$116 per month for enrollees just above 250% of poverty. Since the take-up share is relatively flat in income within segments, a (not unreasonable) extrapolation suggests that requiring relatively modest enrollee premiums shifts the coverage rate from nearly universal to less than half. Indeed, this pattern of relatively flat coverage as income rises within 50% FPL segments is itself interesting and will be useful in our extrapolation exercises in Section 4.2.[24]

The shares of people taking up CommCare and choosing the H and L contracts as premiums vary will be the key inputs into our model to estimate the WTP for health insurance. The first three panels of Table 2 summarize these basic results from the demand analysis. For each discontinuity, we show the enrollee premiums and market shares just to the left and right of each discontinuity (along with the change at the discontinuity), with the shares derived from the RD estimates above.

[22]This difference does not arise because of the use of shares instead of enrollment counts, as we get nearly identical percent changes from running the 2011 regression on raw total enrollment counts (see Appendix Figure 4). Rather, the fall in demand in 2011 is slightly smaller than in other years. This suggests that our 2011 results are not an outlier year but if anything an under-estimate of how much low-income insurance demand falls as premiums rise, and thus result in higher willingness to pay estimates than we might find in other years.

[23]This may seem high relative to previous findings of lower take-up of (free) Medicaid by a similar low-income adult population. However a crucial distinction between Medicaid and CommCare is that, absent a "qualifying event," eligible individuals can enroll in CommCare only during the annual open enrollment period. In addition, CommCare enrollees can only receive coverage starting the month *after* they apply and qualify. By contrast, Medicaid enrollment can happen at any point and can be applied *retroactively* to recent health care use. This has led researchers to observe that many of the people "not enrolled" in Medicaid are in fact "conditionally covered" (Cutler and Gruber, 1996) – in the sense that they would enroll if a health event necessitated expensive care.

[24]The interpretation of this fact is not obvious because the types of people eligible for CommCare also changes with income – e.g., because higher income people are more likely to have employer coverage and therefore be ineligible. In particular, it does not immediately imply that income effects in the demand for insurance are small.

Figure 5: Market Shares, 2011

Panel A. Share Enrolled in Formal Insurance (L + H)

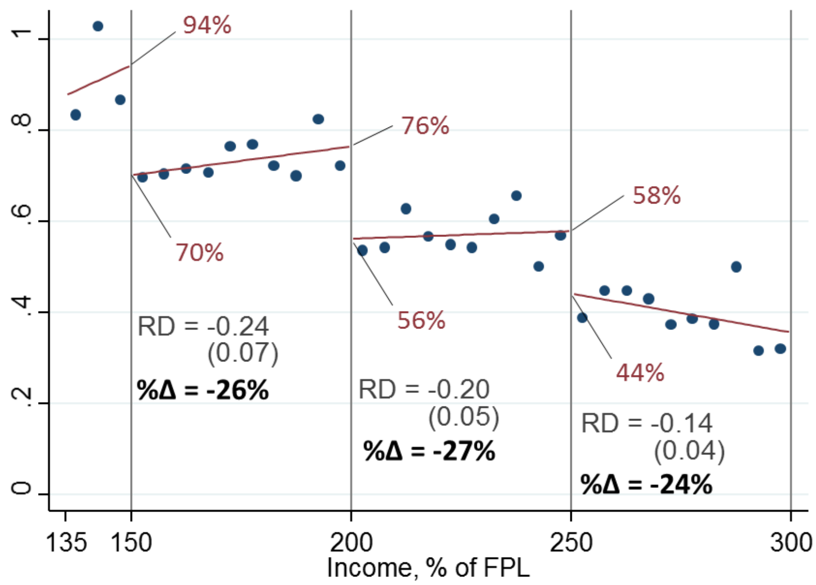

Panel B. Share with H Contract

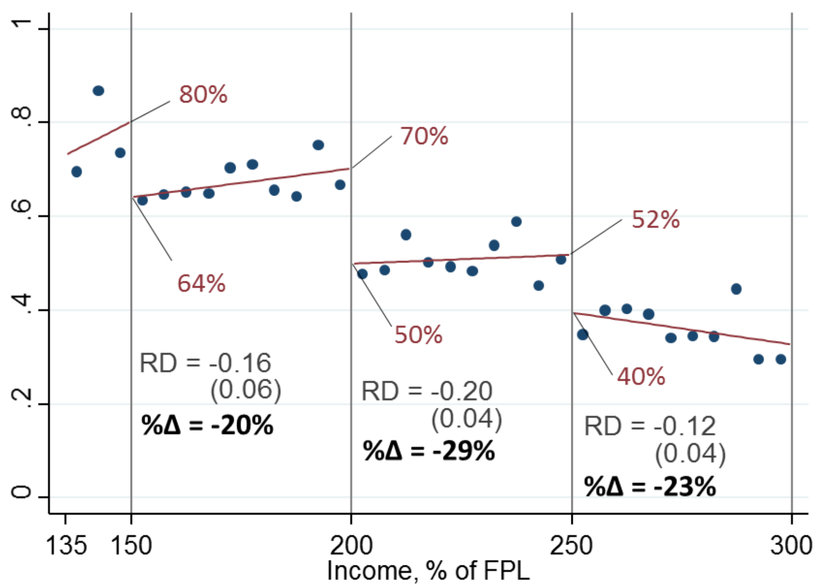

NOTE: These graphs show market shares by 5% of FPL income bin for any CommCare plan (Panel A) and for the *H* plan (Panel B). These market shares are constructed by dividing the raw enrollment counts by the estimate of the eligible population size, both shown in Figure 4. The red lines (and labels) show predicted value from the regression discontinuity specification discussed in the text. The RD estimates (for the change in share) and robust standard errors are shown in gray below each discontinuity; the %Δ values are the *percent* change in the share at the discontinuity.

Table 2: Summary of Demand Estimates, 2011

Variable	150% FPL Below	150% FPL Above	150% FPL Δ	200% FPL Below	200% FPL Above	200% FPL Δ	250% FPL Below	250% FPL Above	250% FPL 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	$335	$382	+$46	$353	$385	+$32	$330	$339	+$9
H Plan	$363	$395	+$32	$369	$404	+$35	$347	$362	+$15
L Plan	$177	$241	+$64	$170	$229	+$59	$177	$142	-$35

NOTE: This table summarizes the results of the demand estimates discussed in Section 3. For each income threshold at which premiums change (150%, 200%, and 250% FPL), it shows the monthly enrollee premium, estimated market shared, and average costs just below and above the threshold, as well as the change across the threshold. The first horizontal panel shows sticker premiums for U (the expected mandate penalty), the L plan, H plan, and the difference $p_H - p_L$, which we use for our main demand estimates. These are identical to the premiums shown in Figure 2. The second panel shows "normalized" premiums, where p_U has been normalized to $0, which we use for robustness analysis. The third panel shows market shares estimated via the RD regressions shown in Figure 5; the share in L is calculated as the difference between the share with any insurance and the share with H. The bottom panel shows average costs, estimated via the RD regressions shown in Figure 6.

Figure 6: Discontinuities in Average Insurer Costs, 2011

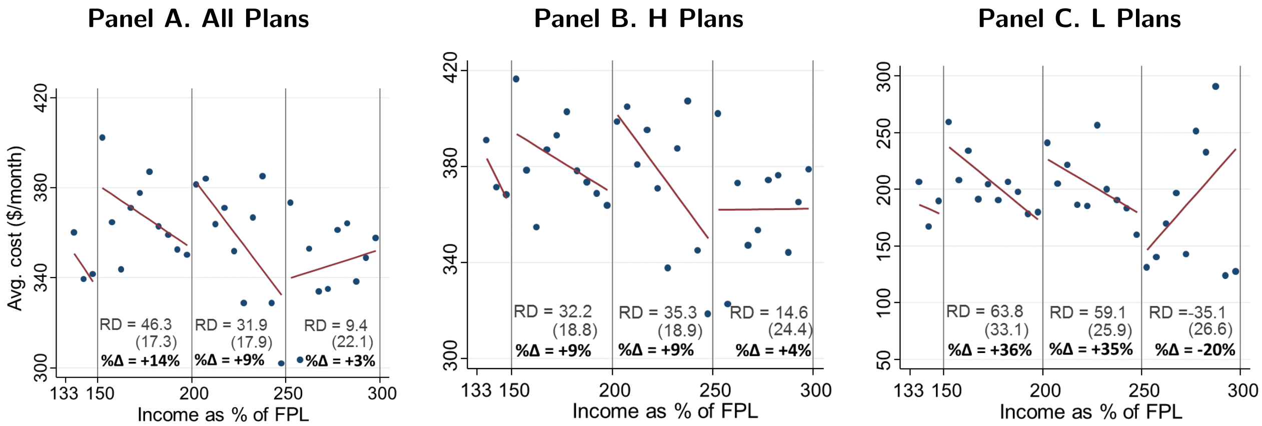

NOTE: These graphs show average insurer medical costs (from claims data) for enrollees by 5% of FPL income bin for all plans (Panel A), the H plans (Panel B), and the L plan (Panel C). The red lines show predicted value from the regression discontinuity specification discussed in the text. The RD estimates and standard errors (robust and clustered at the 1% bin level) are shown in gray below each discontinuity; percent changes relative to the cost just below the discontinuity are shown in bold (labeled "%Δ").

3.3.2 Cost of Insurance

Figure 6 shows the average cost estimates, which are summarized in the bottom panel of Table 2. Panel A plots the average cost for all CommCare enrollees by income bin. As in the pooled 2009-2013 data, it shows that average costs of the insured rise as the enrollee premium increases. For example, the RD estimates indicate that with the rise in enrollee premiums at 150% of FPL, average costs increase by $46 (about 14 percent); this is statistically significant at the 5% level. We also see increases in average costs at the 200% and 250% thresholds, but the increases are somewhat smaller ($10 to $30 or about 3 to 10 percent) and less precisely estimated. These magnitudes are similar to the more precise estimates for the pooled 2009-2013 years shown in Figure 3, Panel B above.

Panels B and C show analogous estimates for the 2011 enrollees in the H plan and the L plan, respectively. Average cost estimates are higher at all points for the H plan than the L plan – consistent with H plans being more generous and also attracting sicker enrollees. The figures show increases in average costs at most of the discontinuities; however only a few are statistically significant.

4 Willingness to Pay and Cost Curves

To interpret the 2011 descriptive results and use them for counterfactual analysis, we use a simple vertical model of insurance demand and cost to map the empirical objects from the previous section into estimates of willingness to pay and cost curves for CommCare. We follow the standard approach in the literature (Einav et al., 2010a), with the added complication that while prior work has focused on settings with two plan options, we have three options: the H plans, the L plan, and uninsured (U). Without additional structure, demand and costs are (potentially complicated) functions of both p_L and p_H, which would require flexible and independent variation in both prices to identify.

Motivated by our institutional setting and definition of the H and L contracts, we apply a vertical model of insurance demand that ranks H, L, and U in terms of quality. This model is helpful for tractability by reducing the required pricing variation and allowing us to define a single index of consumer heterogeneity over which willingness to pay for insurance varies. While we present our baseline results using this vertical model, in the robustness analysis in Section 4.4 we derive bounds (instead of point estimates) for WTP without the vertical model assumption; these bounds are fairly tight and our primary results continue to hold.

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 $w_i(\alpha)$ be the (dollar) value of consumer i for an α-generosity contract. Let p_{ij} be the enrollee premium of contract j, and normalize $p_{iU} = 0$ so that premiums are defined relative to U. Finally, there is an (additively separable) "hassle cost" of the enrollment process for contract j, h_j. We normalize $h_U = 0$ and assume the formal insurance contracts H and L involve the same hassle cost $h_H = h_L \equiv h$. This h will be positive if enrolling in formal insurance involves greater hassle cost (e.g., due to the hassle of applying for insurance and making an active plan choice) and negative if staying in U involves greater hassle (e.g., if outreach groups continually browbeat the uninsured about not having coverage).

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

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

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

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

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:

$$\Delta W_{HL}(i) \equiv W_H(i) - W_L(i) = w_i(\alpha_H) - w_i(\alpha_L).$$

To develop a tractable analytic approach, we adopt model similar to the standard vertical demand model (see Tirole, 1988) by imposing the following two assumptions about preferences:

Assumption 1. ***Vertical preferences for generosity:*** *Everyone prefers more generous contracts: $w_i(\alpha)$ is increasing in α.*

Assumption 2. ***Single dimension of heterogeneity in value for generosity (increasing differences):*** *$w_i(\alpha) = w(\alpha; s)$, where $s \in [0, 1]$ indexes increasing value for generosity, with $dw(\alpha; s)/ds > 0$ and $d^2w(\alpha; s)/d\alpha ds > 0$.*

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. Note also that because of the possibility that $h > 0$, the model rationalizes $W_L(i) < 0$; in other words, it allows for the phenomenon (in our context and others) of individuals not taking up public insurance at zero (or negative) monetary costs.

Assumption 2 imposes increasing differences – i.e. both the value for H relative to L and value for L relative to U are increasing in a single index of preferences, s, which ensures that a single crossing property holds.[25] Of course, it is always possible to define an index s_L that orders people according to increasing W_L and a separate index s_{HL} that orders people according to increasing ΔW_{HL}. Assumption 2 implies that these indices are the same: $s_{HL} = s_L = s$. This rules out cases in which people value L a lot relative to U but only care marginally about H relative to L, or cases in which people value H a lot but consider L to be only slightly better than U.

Demand Curves

Assuming that prices are such that there is positive demand for all contracts, these two assumptions imply that the highest type-s people buy H, middle s people buy L, and the lowest s people buy U. Moreover, with positive demand for all prices, the model has the appealingly 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[26] and D_L depends on both $p_H - p_L$ and p_L.

Figure 7 shows the relationship between willingness to pay and demand graphically. The graph plots W_L and W_H against an x-axis of $1 - s$, so that the highest-value types are on the left. The vertical preference assumption (#1) is captured by the fact that $W_H > W_L$ at all points. The increasing differences assumption (#2) is captured by the fact that the the gap $W_H - W_L$ widens as W_L increases (as one moves left).

Given this setup, the s at which the vertical distance between $W_H - W_L$ equals $p_H - p_L$ defines s^*_{HL}, the person indifferent between L and H. Everyone to the left of this buys H, and the demand for H equals $1 - s^*_{HL}$. Likewise, the s at which p_L intersects the $W_L(s)$ curve determines the person who is indifferent between L and U, or s^*_{LU}. Everyone to the right of this buys U, and people just to the left of this buy L. Mathematically, these points s^*_{HL} and s^*_{LU} are defined by the equations:

$$\begin{aligned} W_H(s^*_{HL}) - W_L(s^*_{HL}) \equiv \Delta W_{HL}(s^*_{HL}) &= p_H - p_L \\ W_L(s^*_{LU}) &= p_L \end{aligned} \tag{2}$$

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

[25] Note that Assumption 2 generalizes the standard assumption in a vertical model of demand. The standard vertical model assumes that $v(\alpha_j; s) = \beta(s)\alpha_j$ – so that choice-specific utility equals $\beta(s)\alpha_j - p_j$ – where $\beta(s)$ is the value of generosity for type s (with $\beta'(s) > 0$). This model satisfies our Assumption 2.

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

is for:

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

In the logic of Figure 7, the positive demand condition implies that prices and demands are such that some types prefer H ($W_H(s) > p_H$), other types prefer L ($W_L(s) > p_L$ but $W_H(s) < p_H$), and some types prefer to remain uninsured ($W_L(s) < p_L$ and $W_H(s) < p_H$). Without loss of generality (since it is an arbitrary index), assume a uniform distribution over s types. Because $\Delta W_{HL}(s)$ and $W_L(s)$ are monotonically increasing functions (by Assumption 2), the equations in (2) implicitly define $s^*_{LU}(p_L) = W_L^{-1}(p_L)$ and $s^*_{HL}(p_H - p_L) = \Delta W_{HL}^{-1}(p_H - p_L)$. 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, these are given by:

$$\begin{aligned} D_H(p_H - p_L) &= 1 - s^*_{HL}(p_H - p_L) \\ D_L(p_L, p_H - p_L) &= s^*_{HL}(p_H - p_L) - s^*_{LU}(p_L) \\ D_U(p_L) &= s^*_{LU}(p_L) \end{aligned} \tag{3}$$

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. demand for H or L) which is easily calculated from the above equations as $1 - D_U$, which depends only on p_L not p_H. 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).

Insurer Costs

We denote by $C_j(s)$ the expected costs to the insurer of enrolling type s in plan j.[27] 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, although we note that this rules out the possibility that the availability of insurance impacts *ex-ante* health behaviors (Ehrlich and Becker (1972)). We define average costs $AC_J(s)$ as the average costs of all individuals with type $\tilde{s} \geq s$:

$$AC_j(s) = \frac{1}{1-s} \int_s^1 C_j(\tilde{s})\, d\tilde{s}. \tag{4}$$

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

[27] In a setting with a binary contract choice (as in Einav et al. (2010a)), the variation in $C_j(s)$ with 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 7: Willingness to Pay Curves: Model

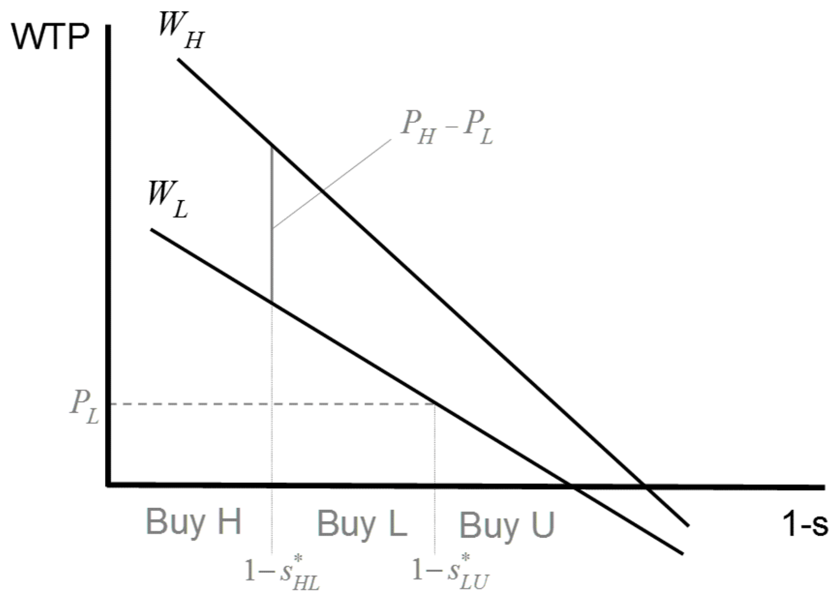

NOTE: This graph shows the theoretical implications of our vertical model for the willingness to pay (W) curves for the H and L plans. The model assumptions imply that both W_H and W_L are downward sloping (i.e., decreasing with $1 - s$) and that the gap between $W_H - W_L$ is also narrowing as $1 - s$ increases. Under the positive demand condition for prices (which this graph assumes), the highest-s types (furthest left on the x-axis) buy H, middle-s types (between s^*_{LU} and s^*_{HL}) buy L, and the lowest-s types choose U.

4.2 Results: Willingness to Pay and Average Cost Curves

4.2.1 Willingness to Pay (W_j)

To construct the willingness to pay curves for the H and L policies, we combine the modeling assumptions above with the empirical patterns documented in Figure 5 to construct the empirical analogues of the $W_H(.)$ and $W_L(.)$ curves in Figure 7. Figures 8-9 walk through this exercise.

Panels A and B of Figure 8 plot the willingness to pay curves for formal insurance and the H contract, respectively, using the enrollee premium and market share estimates from Table 2. These are simple translations of Figures 5 (which is in income-by-market shares space) and Figure 2 (which is in income-by-enrollee premiums space) into market shares-by-premium space. Equation (3) shows that the $W_L(.)$ curve is identified by observations of $(1 - D_U, p_L)$, and the $\Delta W_{HL}(.)$ curve (the gap between $W_H(.)$ and $W_L(.)$) is identified by observations of $(D_H, p_H - p_L)$.

Notice that 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. The demand curve for individuals at 150% of FPL need not, for example, be the same as that for individuals at 200% of FPL. The data allows us to see how different these demand curves are empirically. This is because the premiums that apply between 150-200% FPL are used to 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"). Likewise, the premiums

Figure 8: Willingness to Pay Curves: Empirical

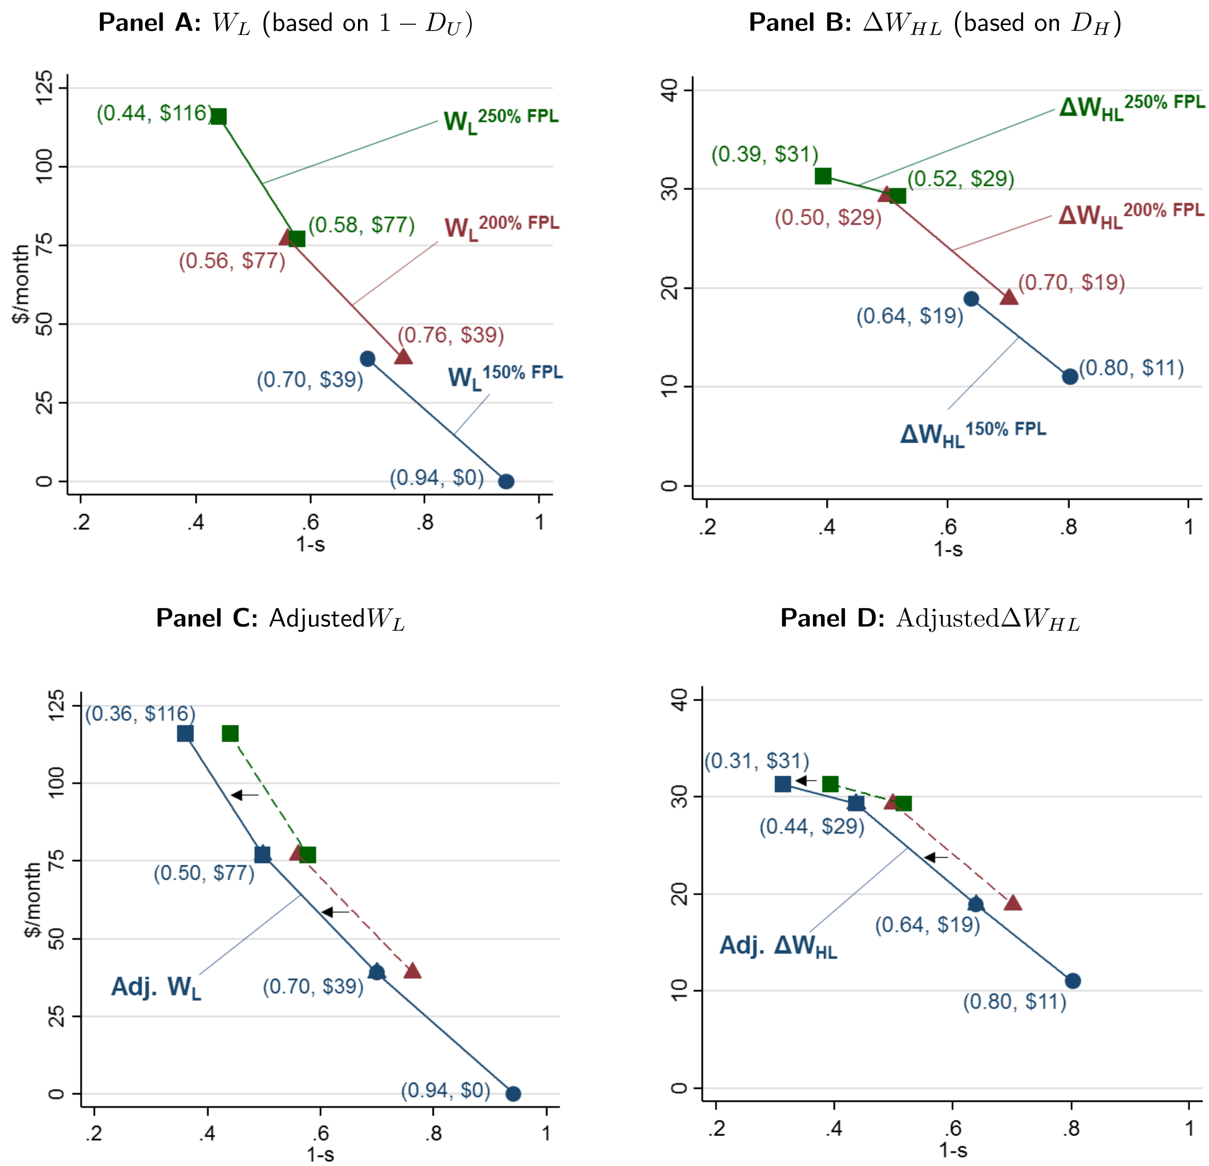

NOTE: These figures show our construction of the willingness to pay curves (W_L and ΔW_{HL}) from the demand points in our RD estimates (as summarized in Table 2). Panel A shows the W_L points, each of which represents an observation of $(1 - D_U,\, p_L)$ drawn from either side of our income discontinuities 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. These adjusted curves are what we use for our final WTP curves.

that apply between 200-250% FPL identify one point on the 200% FPL demand curve and one point on the 250% FPL curve.

In principle, we could identify part of a willingness to pay curve using only one discontinuity. In practice, we prefer to use the data from all three discontinuities because it lets us observe demand at a wider range of prices. Using the three pricing discontinuities, we observe demand for formal insurance spanning about 44 to 94 percent (panel A), and demand for the H contract spanning 39 to 80 percent (panel B). Moreover, the figure illustrates that demand in fact varies little with income over the range of 150% to 300% of FPL. In other words, the market shares are relatively flat within an income range that experiences a given set of premiums (as was evident in Figure 5). As a result, the demand line segments for the three income groups (shown in different colors in Figure 8) are close to intersecting.

To adjust for 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. Motivated by our finding of relatively flat demand shares as a function of income and to allow us to combine demand information across income groups, we assume that income functions as a horizontal shift 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 8 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. In practice, the shift is not very large: for both curves, 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. This reflects the fact that demand does not vary much with income over the range we study. The resultant willingness to pay curves consist of three piece-wise linear segments. The W_L curve shows some convexity, with valuation rising faster than linearly as $1 - s$ decreases; the ΔW_{HL} curve, by contrast, is somewhat concave.

Next, we construct the $W_H(.)$ curve. To do so, we use the fact that our uni-dimensional heterogeneity assumption lets us write the willingness to pay of an s type for H as the sum of the willingness to pay for L, $W_L(s)$ and the difference, $\Delta W_{HL}(s)$. Thus we construct $W_H(s)$ as $W_L(s) + \Delta W_{HL}(s)$. Figure 9 reproduces the adjusted W_L curve from Panel C of Figure 8, and shows the W_H curve constructed by adding the adjusted $\Delta W_{HL}(.)$ curve from Panel D to the adjusted W_L curve. The results are the empirical analogs of the theoretical curves from Figure 7.

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.[28] Using our in-sample variation, we can infer W_L from the 6th percentile of eligible individuals ranked by WTP for insurance (i.e., $1 - s = 0.94$) up to the 64th percentile (i.e., $1 - s = 0.36$). Similarly, our variation lets us infer W_H from the 20th percentile of WTP (i.e.,

[28] Alternatively, we could have shifted the curves to align with other %FPL income levels.

Figure 9: Final Willingness to Pay Curves

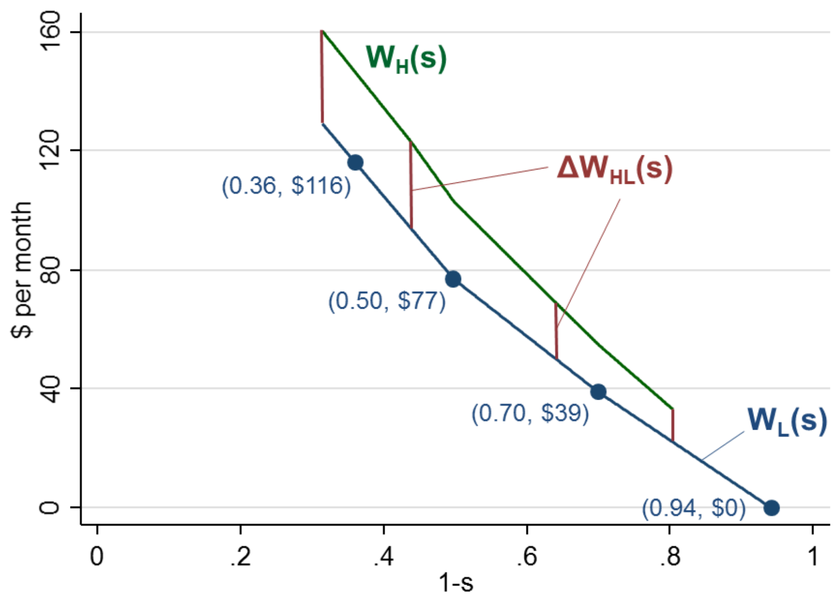

NOTE: This figure shows our final estimated willingness to pay curves (W_L and W_H). The blue curve is the adjusted W_L curve shown in Panel C of Figure 8, with the large dots representing the observed points and the lines interpolating linearly between these values. The W_H curve (in green) is constructed by vertically summing the W_L and ΔW_{HL} curves (as shown in Panel D of Figure 8) at each x-axis value (of $1-s$). The red vertical bars in the figure represent the observed points in the adjusted ΔW_{HL} curve. Note that we extrapolate the W_L curve slightly out of sample (down to $1-s=0.31$) to be able to add on the final point of the ΔW_{HL} curve.

$1 - s = 0.8$) up to the 69th percentile ($1 - s = 0.31$). Willingness to pay for H is non-trivially higher than for L – with the additional value (ΔW_{HL}) ranging from \$11 to \$31/month, or 11-30% of the median person's WTP for L (relative to uninsurance). In total, we infer a WTP for the H plan that ranges from \$33/month (at the 20th percentile) to \$161/month (at the 69th percentile). The median eligible individual has WTP for H of \$103/month.

4.2.2 Average Cost Curve (AC_j)

In Figure 10 we construct the average cost curve for the H policy, AC_H.[29] Panel A plots three line segments representing the average cost curves implied by our RD estimates at the 150%, 200%, and 250% FPL thresholds. The x-values of these points are the shares in the H policy around each discontinuity (from Panel B of Figure 5); the y-values are the estimated average costs (from Panel B of Figure 6). For instance, just below the 150% FPL discontinuity, 80% are in the H policy and the average cost is \$363. Just above the discontinuity, 64% of people are in the H policy and average cost is \$395. Therefore, the average cost curve for 150% FPL flows through the points (64%, \$395) and (80%, \$363).

Figure 10 also shows that the average cost curves are lower at higher-income discontinuities, consistent with health care costs falling with income. As a result, we need to adjust the relevant curves to get an estimate of the marginal cost curve for 150% FPL (to compare to our V_H estimated for the 150% FPL group). 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.[30] To be consistent with how we adjusted the value curves, we also shift the x values to line up with the 150% FPL curve. Panel B shows the resulting adjusted average cost curve.

4.2.3 Cost Curve for Marginal Enrollees (C_j)

Why is willingness to pay below the average cost of the plan? One potential reason is the adverse selection indicated by the downward sloping average cost curve AC_H (see Figure 10). This suggests comparing an enrollee's WTP to his own insurance costs, rather than average costs of all enrollees. To do this, we need to estimate type-specific costs $C_j(s)$.

Constructing C_H The vertical model lets us use the price discontinuities above to identify the $C_H(s)$ curve. To see this, note that the *total costs* to the insurer under the H contract at prices $p = \{p_L, p_H\}$ equals:

$$TC_H\,(p) \equiv \int_{s^*_{HL}(p)}^{1} C_H\,(s)\,ds = (1 - s^*_{HL}\,(p))AC(s^*_{HL})$$

[29] We focus on AC_H because AC_L policy depends on both p_H and p_L, and we do not have independent variation in these two enrollee premiums. We will not be able to directly estimate the average cost curve for L but, as we will discuss below, we can exploit the fact that its market shares for L are relatively small to draw some inferences about C_L.

[30] This is not an innocuous assumption, as there is clear evidence that cost is declining in income – our assumption is that the slopes of the average cost curves with respect to type, s, are stable with income even though the levels of costs are different.

Figure 10: Construction of Average Cost Curves (H)

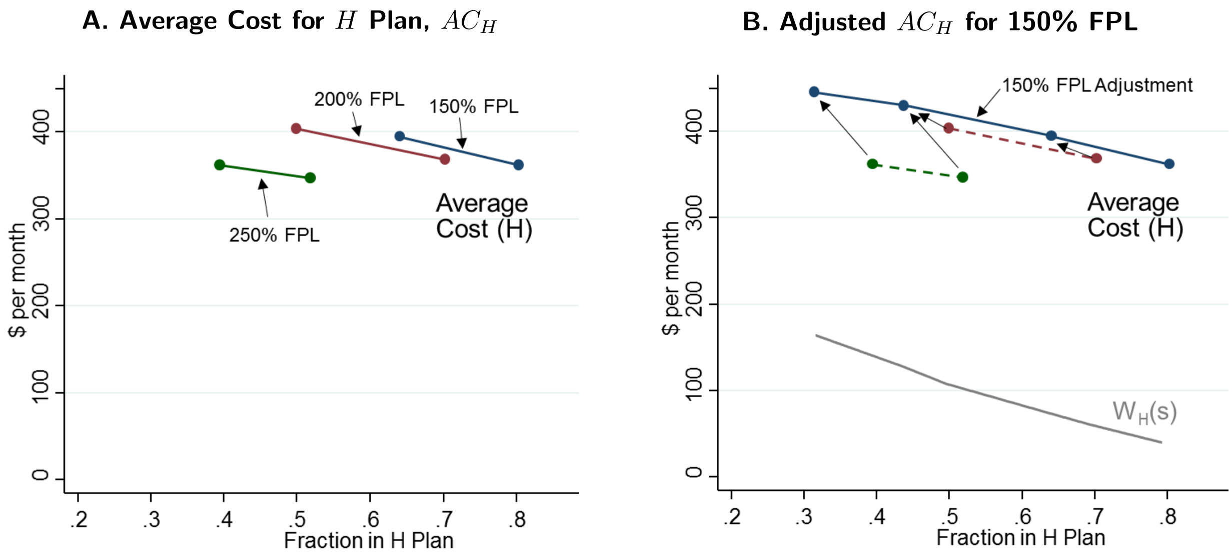

NOTE: These figures show the construction of our average cost curve. Panel A shows the raw average cost points for the H plan, drawn from the RD estimates around each of our three income discontinuities (as summarized in Table 2). Panel B shows how we generate our adjusted AC_H for 150% FPL by translating the 200% and 250% FPL line segments to line up with the 150% FPL segment. This adjusted AC_H curve is what we use for the remainder of our analysis.

where AC(s) was defined in equation (4). This formula integrates over all the individuals who choose the H contract at these prices, which under the vertical model structure is everyone with $s > s^*_{HL}(p_H - p_L)$. Therefore, the share who buy H depends only on $p_H - p_L$.

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:

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

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 $D_H = 1 - s^*_{HL}$). Given estimates of the discontinuity on total costs of H, $\frac{dTC_H}{d\theta}$, and 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, $s^*_{HL}(p)$,

$$C_H(s^*_{HL}) = \frac{\frac{dTC_H}{d\theta}}{\frac{dD_H}{d\theta}} \tag{5}$$

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 $s = s^*_{HL}$. 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 (e.g., Einav et al., 2010a). The key requirement here is that $p_H$and p_L do not change by the same amount at the discontinuities (so that $d(p_H - p_L) \neq 0$.)

Panel B of Figures 5 and 6 showed how shares (D_H) and average costs (or AC_H) changed at the pricing discontinuities. We map these into (5) using the identity that total costs equal average costs times demand: $TC_H = AC_H \cdot D_H$.

The resulting cost curve $C_H(s)$ is shown in Figure 11, along with the previously estimated curves $W_{H,}$, AC_H and W_L. The x-values for these points are defined as 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, $AC_H(s)$. The gap between C_H and AC_H is sizable, especially for 150% and 200% FPL. For instance at 150% FPL, the marginal person who leaves the H plans cost \$236 per month, about 40% less than the average remaining enrollee.

Constructing C_L We cannot estimate C_L in a similar manner as C_H because doing so would require variation in p_L and p_H that is orthogonal to $p_H - p_L$, which we do not have. However, we can again exploit the fact that few people actually purchase the L plan in our data, which implies that the average L enrollee is likely quite similar to the marginal enrollee. In the language of the model, because few people buy L, the set of types s that purchase L falls in a relatively narrow range. For instance, just above 150% FPL, just 6% of the population buys L, and we estimate (see Figure 8 panel A) that the marginal individual who enrollees at that premium is at the 70th percentile of the willingness to pay distribution (i.e. $1 - s = 0.7$) ; thus the 6 percent who buy at that premium span $1 - s \in [0.64, 0.70]$. In other words, when we observe costs in the L plan just above 150% FPL, we observe an average of $C_L(s)$ for individuals in this narrow WTP range. This is analogous to what we can see for C_H – where we infer the average of $C_H(s)$ for a small set of "marginal" individuals who leave the plan when its premium increases.

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 the one threshold for simplicity. In practice, this means that we use our estimates (see Table 2) of the average C_L of \$177 per month for those enrolled just below 150% FPL (where the relevant s range is $1 - s \in [0.80, 0.94]$) and \$241 per month just above 150% FPL (where the relevant s range is $1 - s \in [0.64, 0.70]$).

We plot these two $C_L(s)$ points (locating them at the midpoint of each s range) in Figure 11. Broadly, the implied C_L curve is quite similar to the C_H curve, at least 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. This means that the much lower observed average costs in the L plan (see Table 1) is driven largely by favorable selection – even conditional on risk adjustment – rather than by lower moral hazard. For a given s type, costs under

Figure 11: Willingness to Pay and Cost

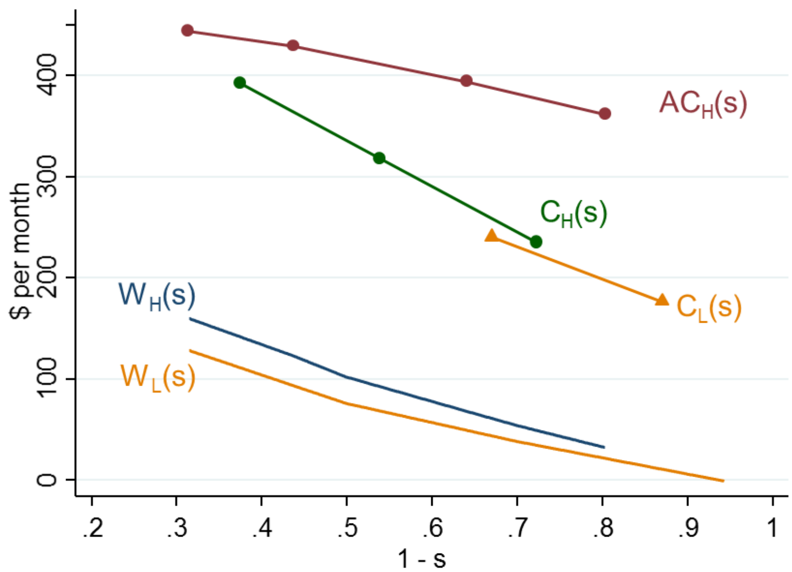

NOTE: The figure compares our estimated average cost, marginal cost, and willingness to pay (WTP) curves for the H plan. The average cost and WTP curves are identical to those shown in Figures 9 and 10. The marginal cost points are constructed using the formula in equation (5) for each pair of AC_H points, with the x-value set as the midpoint between the two points x-values.

the two plans are similar despite different networks.[31]

4.3 Implications for Take-up

Figure 11 displays our key findings for individuals at 150% of FPL. Our price variation spans roughly the 70th to the 6th percentile of the WTP distribution (essentially all but the top 30th percentiles of the WTP distribution). Throughout this range, the $W_H(s)$ curve is substantially below both the average cost curve and the cost curve. The fact that throughout the range of our data we we find $C_H(s) > W_H(s)$ is particularly striking, since $W_H(s)$ and $C_H(s)$ (for a given value of s on the x-axis) represent enrollee WTP and insurer costs *for the same people*.[32] Thus, enrollees' WTP for the H contract is far below the costs that insurers incur in covering them.

For example, at the median of the WTP distribution, $W_H(0.5)$ is \$103, compared to average costs – $AC_H(0.5)$ – of \$419 for those with median or higher WTP, and costs for the median individual $C_H(0.5)$ of \$337. The marginal enrollee is willing to pay less than one-third of the costs he imposes on the insurer. Even at the highest in-sample point we observe (the 70th percentile of the WTP distribution), $W_H(0.7)$ of \$161 per month is still substantially below average costs of insuring the top 30th percentiles of the WTP distribution ($AC_H(0.7) = 445$), as well as costs for the 70th percentile WTP individual – $C_H(0.7) = \$408$. Even if one could eliminate adverse selection and set premiums for

[31]The existence of substantial adverse selection in this market that is not fully offset by risk adjustment is consistent with evidence from Shepard (2016).

[32]By contrast, $AC_H(s)$ represents the average of c_H over all individuals to the left of s.

the marginal enrollee equal to his costs (i.e. $p_H(s) = c_H(s)$), at least 70 percent of individuals would not enroll.

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, as shown in Figure 11: To the right of $1 - s = 0.64$, costs range from \$177 to \$241 per month whereas WTP is less than about \$50 per month. Indeed, the observed C_L points lie above our *maximum* in-sample WTP for L of \$129 per month. If the C_L curve continues to rise with WTP (i.e., moving leftward in the graph) – consistent with adverse selection – then 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 that take-up is likely to be low even for 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 WTP estimates for L in Figure 9 indicate that only 69% of the market would enroll if offered L with a \$39 monthly premium. Likewise, our WTP estimates for H indicate that only 76% would enroll in H if offered H with a monthly enrollee premium of \$42.40, which is a 90% subsidy relative to its insurer price. These estimates have implications for understanding enrollment in the ACA subsidized exchanges. In the ACA's subsidy design, enrollee premiums are even higher than in Massachusetts for people of a given income. For instance, a single individual at 175% FPL would owe about \$83 for the second-cheapest ACA silver plan.[33] Our estimates from W_L suggest that at an enrollee premium of \$83, only about half the market would buy L. And take-up might be even lower than suggested by W_L because the actuarial value of ACA silver plans is lower than the pre-ACA plans in Massachusetts.[34]

Our results suggest that unless enrollee premiums are set substantially below average costs, the market for low-income health insurance would significantly unravel. As seen in Figure 11 W_H lies below AC_H for at least 70% of the distribution; this suggests that without subsidies that lower enrollee premiums below average cost, at most 30 percent of the market would enroll in H. Table **3** summarizes some take-up implications of potential subsidies for a market that offers only H.[35] It shows, for example that with enrollee premiums that are 75% below average costs (i.e. subsidy in excess of \$300 of average costs) 50% of the population would enroll. Premiums would need to be 90% below average costs in order to induce 80% of the market to enroll. Interestingly, the per-enrollee subsidy cost increases by only \$2 as subsidies move from 90 percent to 100 percent of average cost, reflecting the fact that average enrollee costs are declining steeply as healthier individuals are brought into the market.

Figure 12 provides a simple linear extrapolation (in dashed lines) as one way to roughly approximate

[33]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. The enrollee premium of the cheapest silver plan would depend on insurer prices.

[34]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 70%, but cost-sharing subsidies raise this to 94% for enrollees between 100-150% FPL, 87% between 150-200% FPL, and 73% between 200-250% FPL.

[35]Given the assumptions of our vertical model model (see Section 4.1), in a market that offered only H, the individuals who currently buy L would either be uninsured or buy H and the costs they impose on the H plan is accurately reflected in the c_H curve.

Table 3: Implications of Alternative Subsidies for H Plan

Enrollee Premium	Subsidy		Share Insured	Marginal Enrollee		Average Cost
	% of AC	\$/month		WTP	Cost	
\$237	50%	\$237	6%	\$237	\$468	\$475
\$105	75%	\$315	49%	\$105	\$340	\$421
\$37	90%	\$329	79%	\$37	\$210	\$366
\$0	100%	\$331	96%	\$0	\$141	\$331

NOTE: The 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 \$237, \$105, \$37, 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.

WTP and costs out of sample. We construct this by extending the left-most segments of W_H and AC_H linearly to the left (and recalculating c_H using equation (5)).[36] Likewise, we extend the right-most segments of of W_H and A_H out to its value above the right-most value of W_L that we observe (at $1 - s = 94\%$) using a linear extrapolation of the right-most segment of the ΔW_{HL} and AC_H curves.[37] This extrapolation leads to estimates of willingness to pay are still everywhere far below average and marginal 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 individual's WTP lies everywhere below their own 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 the basis of enrollee own costs $C_H(s)$, 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.[38]

[36]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.

[37]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));

[38]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) was 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 variance of costs changes around our pricing discontinuities. While the results are fairly noisy, there is no evidence that the standard deviation of costs is increasing at the price discontinuities

Figure 12: Value and Cost of H – Extrapolation

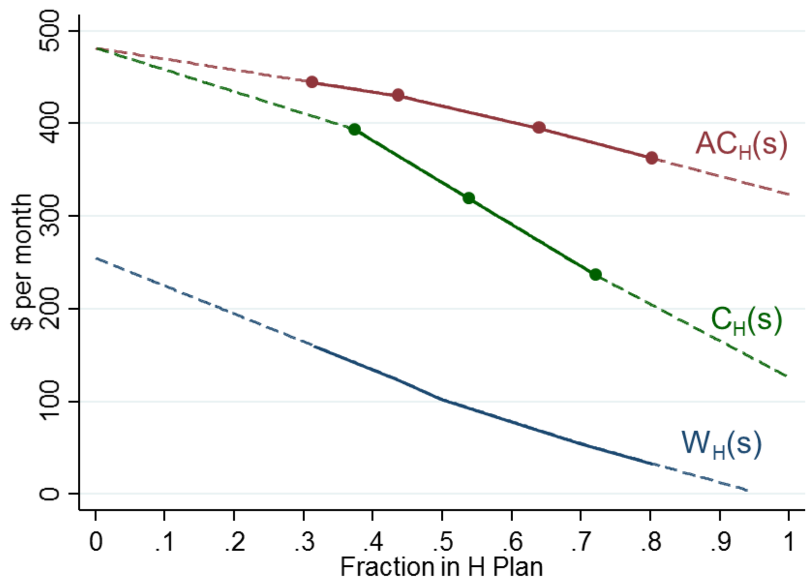

NOTE: The figure shows out-of-sample extrapolations for our willingness to pay (WTP), average cost, and marginal cost curves for the H plan. The solid lines are our in-sample estimates, identical to those shown in Figure 11. The dashed lines are the extrapolations. Both WTP and average cost are extrapolated linearly using the slope of the left-most and right-most line segment. Marginal cost is extrapolated by calculating based on the implied values from the average cost curve, applying the formula for MC in equation (5).

Table 4: 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	100% Subsidy
Baseline Estimates	[0.06, 0.69]	$103	$337	$419	49%	79%	96%
(1) Alternate RD Specifications							
Smaller Bandwidth (25% FPL)	[0.06, 0.71]	$100	$314	$421	48%	78%	97%
Quadratic Functional Form	[0.16, 0.72]	$98	$348	$405	49%	73%	84%
Omit 200% FPL Estimates	[0.06, 0.70]	$97	$344	$421	46%	79%	96%
(2) Alternate Take-up Estimates							
Unscaled ACS Eligible Pop.	[0.44, 0.81]	$24	$186	$354	29%	46%	57%
(3) Accounting for Mandate Penalty							
Use Normalized Premiums	[0.06, 0.69]	$74	$337	$419	32%	68%	100%

NOTE: Table shows sensitivity analysis of our estimates of WTP and costs for H plan and of enrollment in H under various subsidies in a market that only offers H. The top line ("Baseline Estimates") reproduces results from Figure 11 and Table 3. The remainder of the table shows analogous estimates with (1) alternate RD specifications, (2) alternate take-up estimates, and (3) normalizing premiums used in our demand estimates by the expected mandate penalty. For each specification, we report the in-sample range of s values, estimates of WTP_H, AC_H,and $C_H$at the median of the WTP distribution ($s = 0.5$), and the share who purchase H with various percent subsidies of average costs.

4.4 Sensitivity

We explored the sensitivity of our key findings to a number of our empirical implementation choices. Table 4 summarizes some of our findings. Each row represents a single deviation from the baseline specification, as indicated. In all the alternative specifications we consider below, our qualitative results continue to hold: willingness to pay for H is less than both average cost (AC_H) and own cost (C_H) withing the range of our data. Quantitatively, the results suggest that our baseline estimates are at the upper range of estimates of willingness to pay for insurance among low income adults: under some of our alternative specifications, the share who enroll in H at a given subsidy level is substantially lower.

RD Specification

Our baseline RD analysis was a simple linear RD in which we allowed the 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 panel of Table 4shows that the results are quantitatively quite stable across alternative implementation assumptions.

The first two rows show results using a narrower (25% of FPL) bandwidth around the threshold, and results using a quadratic (rather than linear) functional form for the running variable, Inc_b (with

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 consistent with the underlying cost variation in the data.

the baseline bandwidth). The third row shows results from our baseline specification but excluding one of our three thresholds: the 200% of FPL threshold. This threshold is potentially problematic because of two other small changes that occur at this point that we described in Section 2: eligibility expands slightly at 200% of FPL to cover pregnant women and HIV-positive individuals (for whom Medicaid eligibility ceases at 200% of FPL) and copays increase slightly at 200% of FPL, resulting in a decline in actuarial value of the plans from 97% to 95%. Since these discrete changes could affect enrollment independently of the change in enrollee premium, they technically violate our assumption that eligibility and enrollment would be smooth through the income thresholds in the absence of the variation in enrollee premium. Therefore we show results excluding this threshold and just using the variation at 150% and 250% of FPL to construct value and cost curves. We fill in the (now missing) space between the 150% and 250% FPL line segments by extrapolating the 150% FPL segment linearly.

Alternative Estimates of Take-up

The administrative data alone are sufficient to estimate willingness to pay, average cost, and own-type costs for the enrolled population. This is key for many of our results, including the fact that willingness to pay is substantially below average cost and own cost. However, an estimate of the share of eligibles who enroll in CommCare (i.e. take up) is essential for pinning down where in the distribution of willingness to pay for insurance our observed demand changes occur. For example, our baseline estimates suggest that the in-sample demand curve spans the 69th to the 6th percentile of the willingness to pay distribution; this estimate – and likewise the location for the "median" WTP individual – depends on our take-up assumption.

As we discussed in Section 2, in our baseline analysis, we used data from the ACS to estimate that 63 percent of the CommCare eligible population enrolls in CommCare. Here, we show the sensitivity of our results to an alternative take-up estimate that we calculated by dividing the administrative counts of enrollment in CommCare with the ACS estimates of the size of the eligible population. This yields a substantially lower take-up rate of 37%.[39] Under this alternative take up assumption, our in-sample demand curve now spans the 81st to the 44th percentile of the willingness to pay distribution (rather than the 69th to 6th as in our baseline). The results therefore imply even lower take up for a given subsidy. For example, with this alternative take-up rate, we estimate that with a 75 percent subsidy of average costs, only 29 percent of eligible individuals would enroll in H, compared to 49 percent in our baseline analysis.

Accounting for the Mandate Penalty

Our baseline analysis uses the sticker prices for L and H for enrollee premiums p_L and p_H. This implicitly assumes that individuals do not take the expected mandate penalty from remaining uninsured into account in deciding whether or not to buy insurance – an assumption, which as we discuss in Section 2 may be a reasonable one given the institutional environment. Nonetheless, here we report results where we instead assume the individual responds to the normalized enrollee premium; these

[39] Specifically, we estimate an average of 62,096 CommCare enrollees in FY 2011 in the administrative data, compared to an eligible population of 168,042 in the ACS.

normalized premiums, summarized in Table 2 are calculated by subtracting the expected mandate penalty values (shown by the black dots in Figure 2) from the "sticker prices" shown in that same figure.

Accounting for the mandate penalty lowers the effective premiums and therefore implies even lower willingness to pay than our baseline estimates; at the median of the willingness to pay distribution, for example, willingness to pay is about three-quarters of the baseline. Estimates of average costs are of course unchanged. As a result, our estimates now imply that with a 75 percent subsidy of average costs, only about 32 percent would enroll in H.

Inertia in Plan Demand

Our model focuses on demand for H, L, and U as static choices made by eligible individuals. One concern with this approach is that it abstracts from inertia or switching costs, which have shown to be relevant for insurance plan choices over time (Handel, 2013; Ericson, 2014; Polyakova, 2016). Intuitively, people might make flexible, fully informed plan choices when they "enter the market" (i.e., initially become eligible for CommCare) but then remain "stuck in place" over time, even as their prices and/or subsidies change.

This raises several potential concerns with our estimates. One concern is that enrollees' income might change, leading them to move across the RD income threshold over time. For instance, an individual might enter the market with income of 149% of FPL and choose to buy insurance. But if their income increased to 151% FPL – reducing their subsidy and raising their premiums for formal insurance – they might be unaware of the change and not fully react. Institutionally, this type of inattention 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. Moreover, we note that if this type of inertia were important, our estimates would tend to *understate* our estimate of 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 is likely 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 insurance (H or L) to uninsurance at the RD income thresholds. Thus, switching between H and L is less likely to be empirical 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.[40] 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 in theory be problematic.

As one way to investigate the potential empirical importance of inertia for our estimates, we re-

[40]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.

estimated the RD estimates of enrollment counts separately for new enrollees entering the market in 2011 (who by definition made an active choice). The results are shown in Panel B of Appendix Figures 18 and 19; for comparison Panel A shows the results for total enrollment counts. New enrollees comprise about one-fifth of all enrollee-months, suggesting that inertia is not a concern for a non-trivial share of our study population. Moreover, the percent reductions in enrollment at the income discontinuities are 22-33% for new enrollees, similar to the 25-27% reduction in total enrollment at the same thresholds. 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 formal insurance options, H and L. While this is a reasonable model of the CommCare market in 2011, one might be concerned that our results depend on the assumptions underlying this structure. In Appendix C, we show that we can eliminate the vertical assumptions and still obtain *bounds* on WTP for CommCare health insurance (W^{Ins}) – defined as WTP for each individual's most preferred plan. The relaxed model assumes only that consumers make optimizing plan choices. We show how 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 WTP across a range of the population distribution.

Appendix Figure 15 shows the resulting bounds on WTP for CommCare, with our vertical model estimates of W_L and W_H shown for comparison. In practice, our W_L and W_H estimates are quite similar to the lower and upper bounds. 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 also only slightly above W_H. From this exercise, we conclude that our basic estimates of (low) WTP for insurance are robust to weakening the vertical model.

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 relative prices among the H options do not change much at the income thresholds (their prices are nearly identical for all consumers; see Appendix Table 6). Second, given the similarity of $C_H(s)$ and $C_L(s)$ (see Figure 11), it seems unlikely that cost differences among the (much more similar) H plans for a given s type would be large.

5 Discussion and Normative Implications

Our results suggest that willingness to pay for insurance among a low-income population is extremely low. Our estimated willingness to pay curve lies far below the average cost curve, which is about four times larger than WTP. Moreover, this result is not simply the result of adverse selection driving a wedge between average costs and costs for marginal enrollees. Rather, individuals have willingness to pay that is roughly one-third of their *own* expected costs imposed on the insurer. This stylized fact runs counter to a standard assumption in textbook models of insurance demand – that WTP for insurance equals expected costs *plus* a value of risk protection.

Why might this be the case? And, should low levels of take up and willingness to pay raise normative concerns about the value of health insurance subsidies? We briefly discuss both issues here.

5.1 Why Is WTP so Low?

Why is willingness to pay so far below individuals' own costs they impose on the insurance company? Of course, one potential explanation is that low WTP reflects a behavioral bias, such as lack of information about eligibility, inattention, inertia, or misperceptions (Spinnewijn (2017)Baicker et al. (2012)). One could also worry that individuals are liquidity constrained. However, 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.

Our finding above 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. However, we cannot rule out behavioral biases more broadly. Instead, we consider two reasons for why individuals might truly not derive significant value from health insurance relative to the costs they impose on the insurer: moral hazard and uncompensated care. We conclude that the size of uncompensated care for low-income populations provides a plausible explanation for their low WTP.

Moral Hazard

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 marginal cost they would have had to pay for it when uninsured. Quantitatively, however, 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 increases health care spending by about 25% relative to being uninsured (Finkelstein et al., 2012); as discussed, the CommCare products studied here are similar to Medicaid plans in Massachusetts. Moreover, the fact that we find willingness to pay far below own cost for both the more and less generous H and L plans suggest that moral hazard is not a sufficient explanation of low willingness to pay relative to costs.

Uncompensated Care

An increasing body of evidence has highlighted that uninsured individuals do not pay the full cost of their medical care and that formal health insurance reduces the extent of charity or uncompensated care, either through free charity care offered to low-income patients, or implicitly free in the form of unpaid bills and bankruptcy (Finkelstein et al. (2012); Mahoney (2015); Garthwaite et al. (2015); Dobkin et al. (2016); Hu et al. (2016)). Estimates suggest that the uninsured pay about 20% to 35% of their cost of care (Coughlin et al., 2014; Hadley et al., 2008; Finkelstein et al., 2015), which is remarkably similar to our estimated ratio of WTP to marginal costs for the H plan.

In contrast to moral hazard, 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 and marginal cost.

Let $C_U(s)$ denote expected uncompensated care costs that type-s individuals receive if uninsured. Because our data are comprised of the insured population, we do not observe $C_U(s)$, and moreover, C_U is quite difficult to measure even for the uninsured. Nevertheless, we can form a rough approximation of it as follows. Let x denote the share of the uninsured's total health care costs that they pay out of pocket, so that $1 - x$ reflects the fraction of uncompensated care provided by doctors and hospitals. Recall that $c_H(s)$ denotes the cost we observe for a type-s individual in the H plan. Note that this cost is potentially higher than what their costs would be if uninsured because of a moral hazard response. Therefore, we approximate the *total* cost of the uninsured as $\frac{C_H(s)}{1+\phi}$, where ϕ is an estimate of moral hazard. This yields an estimate of the cost of uncompensated care provided to type s, $C_U(s)$, given by the formula:

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

We assume ϕ is 0.25 and x is 0.2 based on results from the Oregon Heath Insurance Experiment (Finkelstein et al., 2012, 2015);[41] thus net costs – $C_H(s) - C_U(s)$ – are only about one-third of gross costs $C_H(s)$.

Figure 13 presents our resulting approximation of the "net cost" curve, from subtracting uncompensated care, i.e., $C_H(s) - C_U(s)$. The value of $C_H(s) - C_U(s)$ provides an estimate of the net increase in costs paid by third parties (either the insurer or uncompensated care providers) when a type-s enrollee shifts from uninsurance to formal insurance in the H plan. Alternatively, $C_U(s)$ can be thought of as a negative externality of uninsurance, making $C_H(s) - C_U(s)$ an estimate of the net social cost of insurance. As shown in the figure, our approximation suggests that uncompensated care is quantitatively important, representing about 2/3 of the (gross) insurer costs of formal insurance. Thus, 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 percentile) willingness to pay is about \$10 higher than our approximation of net costs. In this sense, the provision of uncompensated care provides a plausible

[41] For our illustrative purposes, we assume that the cost associated with the insurance policy is similar to the H policy primarily because we are able to estimate the $C_H(s)$ curve (as opposed to the $C_L(s)$ curve) using our observed price variation. However, as shown above, the $C_L(s)$ is little different from $C_H(s)$ over the range we can observe both, suggesting that this exercise might not be much different if we could estimate and use $C_L(s)$ in the formula.

Figure 13: Marginal Cost Curve Adjusted for Uncompensated Care

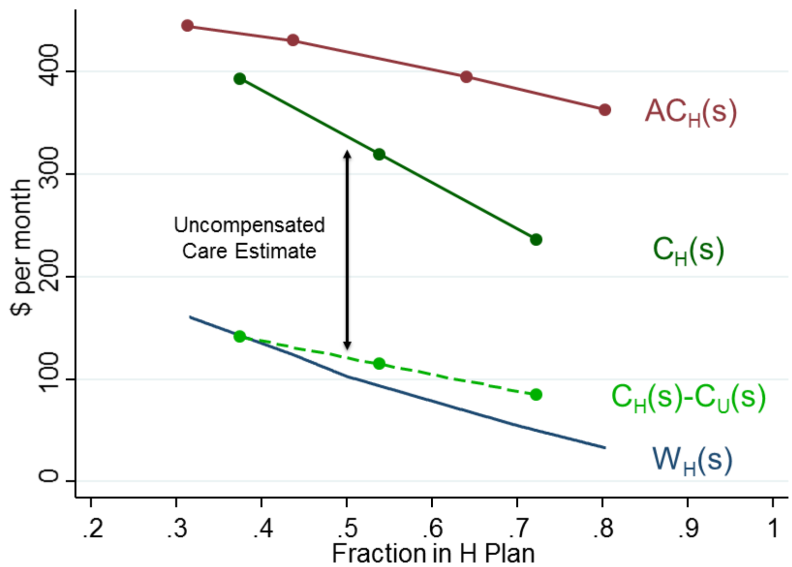

NOTE: The figure shows our estimates of willingness to pay ($W_H(s)$), average cost ($AC_H(s)$), and marginal cost ($C_H(s)$) for the H plan, as in previous figures. In addition, it shows (in the green dashed line) our approximation to the marginal cost curve adjusted for uncompensated care, which is constructed as marginal costs minus estimated uncompensated care (or $C_H(s) - C_U(s)$).

rationalization of why we observe low willingness to pay relative to the marginal cost the individuals would impose on insurers.[42]

A large role for uncompensated care also can 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 contrast with the findings of Hackmann et al. (2015) for higher-individuals. They 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 et al. (2015) who, based on *ex-post* valuation of Medicaid, find it 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 et al. (2015) study adults below 100% of FPL.

In sum, plausible estimates of the size of uncompensated care provided to the low-income population can rationalize the gap between WTP and costs, thereby providing a potential explanation for the low take-up of formal insurance even at high rates of subsidy.

[42]Once again, we have focused on the H contract. As discussed already, willingness to pay for L is slightly below willingness to pay for H, while marginal costs of L may be substantially below marginal cost of H. While willingness to pay for $L$is therefore clearly below c_L, the same type of exercise could suggest that willingness to pay for L is similar to or above the net marginal cost of L.

5.2 Normative Implications

We find evidence of low willingness to pay for formal insurance by low-income individuals. Should this raise concern for the optimality and desirability of health insurance expansions? To address such normative questions, we need to make the additional assumption that willingness to pay is in fact the welfare-relevant metric for judging the welfare of the 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 et al., 2010b). However, it does require assuming away behavioral frictions that may drive a wedge between estimated demand and the value individuals truly derive from the insurance.[43]

Using the model in Section 4, we write the aggregate welfare of to beneficiaries when prices, p_H, are set such that a fraction $1 - s^*$ of the market is enrolled:

$$V\left(1 - s^*\right) = \int_{s^*}^{1} (W_H\left(s\right) - p_H) ds \tag{6}$$

where $p_H = W_H(s^*)$ is the price required for a fraction $1-s^*$ to choose to enroll. The term, $W_H\left(s\right)-p_H$, captures the consumer surplus of an individual s enrolled in the H plan; this surplus equals her willingness to pay $W_H\left(s\right)$ minus the premium she pays, p_H.

Conversely, the cost to the insurer net of the enrollee premiums they collect of enrolling a fraction $1 - s^*$ of the market is given by

$$\chi\left(1 - s^*\right) = \int_{s^*}^{1} (C_H\left(s\right) - p_H) ds$$

where $C_H\left(s\right) - p_H$ denotes the cost to an insurer of enrolling an s type. This reflects the cost to the government of enrolling an s type if we assume that the government is providing H at a price p_H or if the government is subsidizing the H policy in the private market and insurers have zero profit pricing.

We first consider a marginal expansion of an existing insurance subsidy. Suppose a share $1 - s^*$ of the population is enrolled in the H plan and the government considers increasing subsidies to lower the enrollee premium by \$1. Differentiating equation (6) with respect to an increase in the subsidy by $-dp_H$ shows that the welfare impact of a marginal subsidy expansion is given by \$1 per enrollee,

$$\frac{dV}{-dp_H} = V'\left(1 - s^*\right) \frac{d\left(1 - s^*\right)}{-dp_H} = 1 - s^* \tag{7}$$

where $\frac{d(1-s^*)}{-dp_H} = \frac{1}{W'_H(s^*)}$ measures how much the size of the market expands in response to a reduction in premiums. Intuitively, a price of insurance that is lower by \$1 is valued at \$1 for all those currently purchasing insurance, $1 - s^*$.[44] We can write this welfare impact as the product of two terms: the

[43] Another potential concern, recently emphasized by Hendren (2017), is that observed demand may understate the ex-ante value of insurance that would be measured before the individual has learned something about their health type. However, we show in Appendix B, our estimates suggest that even using an *ex-ante* demand measure, willingness to pay would still be substantially below marginal costs.

[44] Although some may choose to purchase insurance because of the price reduction, the envelope theorem ensures that

marginal welfare impact of expanding the market, $V'(1-s^*)$, and the increase in the size of the market from a subsidy increase, $\frac{d(1-s^*)}{-dp_H}$. The greater the steepness of the demand curve (i.e. the larger the value of W_H'), the less the market size will expand in response to a reduction in price.

Conversely, the cost of the subsidy has two components:

$$\chi'(1-s^*)\frac{d(1-s^*)}{-dp_H} = (1-s^*) + \frac{C_H(s^*) - p_H}{W_H'(s^*)} \tag{8}$$

The first component is the "mechanical" cost, $1-s^*$, for inframarginal enrollees. The second component arises because the subsidy expands the market by $\frac{d(1-s^*)}{-dp_H} = \frac{1}{W_H'(s^*)}$ and the marginal enrollees impose a cost of $C_H(s^*) - p_H$. Thus if the cost of the marginal enrollee $C_H(s^*)$ is higher (lower) than their willingness to pay (p_H), 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)).

Hendren (2013) shows that one can construct a simple benefit-cost ratio to measure the marginal value of public funds (MVPF): for every dollar of government expenditure, how much welfare is delivered?[45] In our context, the MVPF per dollar of government expenditure on subsidies is given by the ratio of the benefits in equation (7) to costs in equation (8),

$$MVPF(s^*) = \frac{1}{1 + \frac{C_H(s^*) - p_H}{(1-s^*)W_H'(s^*)}} \tag{9}$$

We consider the impact of a marginal subsidy expansion at $s^* = 0.5$, using our baseline estimates (see top row of Table 4 for a summary) this yields a cost per enrolled individual of $1 + \frac{337-103}{-0.5*239} \approx 2.95$ (the denominator in (9)) compared to a benefit of 1. Put differently, beneficiaries are willing to pay about \$0.3 for every \$1 increase in government spending on subsidies, or a MVPF of about \$0.3. This suggests that subsidies are not welfare increasing relative to a non-distortionary lump-sum transfer that would be valued at \$1 per unit of government expenditure. We obtain similar results for the value of expanding the size of the insurance market through subsides when 30% of the market is enrolled (MVPF = 0.27) and when 70% of the market is enrolled (MVPF = 0.43). Intuitively, the results in Figure 11 indicate that C_H is always above willingness to pay within our sample and therefore marginal subsidy increases anywhere within our sample will limit the value of health insurance subsidies.

Uncompensated Care One key concern with these calculations is that they focus on the cost to the insurer, $C_H(s) - p_H$, as opposed to the net resource cost, $C_H(s) - C_U(s) - p_H$. As noted above, these can differ substantially. Individuals choosing to forego health insurance exercise the ability to utilize

those on the margin are indifferent between being insured and uninsured (i.e. $W_H(s^*) = p_H$)

[45] If one considers two policies: Policy A that delivers \$0.50 of welfare for every \$1 of expenditure and Policy B that delivers \$0.25 of welfare for every \$1 of expenditure, then increasing expenditure on Policy A financed by a reduction in expenditure on Policy B will increase welfare. As shown in Hendren (2014), the relevant benchmark for an MVPF is not 1, but rather the MVPF of other policies that affect similar populations.

uncompensated care; this externality from insurance choice raises a potential Samaritan's dilemma rationale (Buchanan, 1975) for providing health insurance subsidies by using government taxes to internalize the externality imposed on the providers of uncompensated care when individuals choose to remain uninsured.

More generally, one can incorporate the incidence of health insurance subsidies on providers of uncompensated care into the MVPF calculations. In general, one would wish to trace the incidence of the uncompensated care and weight their benefit by a social welfare weight for those providing uncompensated care. Here, because the government is a common provider of uncompensated care (directly and indirectly through Medicaid payments and tax preferences), we can consider a benchmark case in which the government fully recoups the savings on uncompensated care by eliminating subsidies to uncompensated care providers. In this world, the true cost to the government is not given by equation (8), but rather by replacing $C_H(s^*)$ with the net cost, $C_H(s^*) - C_U(s^*)$. Thus, the MVPF becomes $MVPF_{net} = \frac{1}{1+\frac{C_H(s^*)-C_U(s^*)-p_H}{(1-s^*)\left(-W'_H(s^*)\right)}}$.

When we account for the savings from uncompensated care costs, we now find the net MVPF is much closer to 1. For example, we estimate $MVPF_{net}$ of 0.81 when 70% of the market is insured (compared to our estimate of $MVPF$ of 0.43 when we do not account for uncompensated care) and an $MVPF_{net}$ of 1.23 when 30% of the market is insured, compared to our estimate of $MVPF$ of 0.27 when we do not account for uncompensated care). This suggests that if the government is a primary payer of uncompensated care costs, then the marginal value of public funds of providing subsidies is fairly close to 1.

Accounting for Distributional Effects The above calculations suggest that subsidies for health insurance for low-income adults are valued at less than their gross cost, but are valued fairly close to their "net" cost that subtracts an estimate of uncompensated care. This suggests that subsidies would not increase aggregate welfare if individuals had to finance the gross costs of the subsidy by the individuals themselves. Moreover, even after accounting for uncompensated care, the MVPF falls below 1 when a large fraction of the market is insured. Does this mean that healthcare subsidies are not desirable?

Not necessarily. In practice, subsidies for health insurance for low-income people are not financed by these individuals themselves. Those who pay the "Costs" differ from those who receive the "Benefits." This suggests that a normative analysis of subsidies should also consider any value from distributional impacts and that comparing the MVPF to 1 (i.e. is WTP above government costs) is not the relevant benchmark.[46]

There are at least two ways to analyze the distributional consequences of subsidizing health insurance for low income individuals. First, one can apply a social welfare function or weight individual

[46] It also suggests that expanding insurance subsidies might decrease the returns to labor effort; if true, the resulting fiscal externalities from lower labor earnings would further increase the total cost of expanding subsidies. Kaplow (2006) shows that if the tax schedule is optimized and health care satisfies a weak separability condition in the utility function, these costs will offset the social value of redistribution so that one would not actually want to account for any additional value from redistribution beyond the simple cost-benefit test above. Here, we do not impose these assumptions and calculate the social value of redistribution. We proceed under the benchmark assumption that the subsidies do not cause any fiscal externality from changes in labor earnings; any such effects would serve to further raise the costs of the policy.

surplus using a parameterization of social marginal utilities of income that translate individual willingness to pay ($W_H(s)$) into social willingness to pay (Saez and Stantcheva, 2016). To illustrate, suppose the social welfare function is captured using a utilitarian social welfare function over CRRA individual utility functions. A rough calculation from the Consumer Expenditure Survey suggests that the median consumption of our study population is about 60% that of median consumption in the general population in Massachusetts.[47] With a CRRA coefficient of 3 (i.e. marginal utility of $u'(c) = 1/c^3$), the social value of insurance would be about 5 times the recipient's value; this would suggest society should prefer policies with an MVPF of at least 0.20, yielding a threshold crossed by our MVPF estimates for the entire observed distribution. In this sense, one may find additional subsidies desirable even if there is no positive benefits of insurance on providers of uncompensated care.[48]

An alternative benchmark for the MVPF can be may not be whether individuals are willing to pay the costs (i.e. whether the MVPF is greater than 1) but rather whether the MVPF of the health insurance expansion is higher (or lower) than the MVPF of other policies that deliver welfare gains to similar populations. Redistribution involves distortionary costs, so the relevant benchmark for low-income populations may be an MVPF that is less than 1 (Hendren, 2014). For example, Hendren (2013) calculates that EITC expansions can deliver $0.90 of welfare to low-income individuals for every $1 of government resources expended. By this benchmark, health insurance subsidies are less efficient at redistribution assuming the gross costs reflects the cost to the government (so that the MVPF is around 0.3). In contrast, if the government is the beneficiary of uncompensated care (so that the MVPF is around 0.8-1.2 for observed market sizes), it suggests additional health insurance subsidies may be deliver a positive welfare benefit relative to alternative policies like the EITC. If parties other than the government bear the incidence of providing uncompensated care, one would need to identify them and place social welfare weights on them to appropriately account for these transfers within the social welfare framework. Because the ultimate welfare conclusions depend on who bears the economic incidence of the provision of uncompensated care, a deeper understanding of this incidence is an important direction for further work. Nonetheless, our results suggest that large social welfare gains from greater insurance subsidies to low-income individuals relative to equivalent-sized cash transfers would require a significant social value from reducing uncompensated care costs.

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 percent of the low-income eligible population, we find that willingness to pay for insurance is far below the average cost curve – what it would cost insurers to provide coverage to all who would enroll if the premium were

[47] 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.

[48] More concave social welfare functions would further increase social value, while more (less) concave individual utility functions would also increase (decrease) social value. For example, maintaining a utilitarian social welfare function but with CRRA coefficient of 1 (i.e. log utility), the social value of insurance would be only 1.7 times the recipient's value, which would not be above C_H within our sample.

set equal to that WTP. 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 only about one-third of own costs; thus even if insurers could offer actuarially fair, type-specific prices, at least 70 percent of the market would be uncovered.

From a positive economics perspective, our results point to substantial challenges in generating universal coverage via partially subsidized insurance programs like the ACA's exchanges. For example, we estimate that subsidizing insurer prices by 90% would lead only about three-quarters of potential enrollees to buy insurance. Subsidizing costs by 90% would generate only 80% coverage. More generally, our results help predict and explain low take-up of even heavily subsidized insurance for low-income adults. The enrollee premiums in Massachusetts are significantly lower than those for similar low-income people in the ACA's exchanges. For example, for an individual earning 175% of poverty, the ACA subsidies make the second-cheapest silver plan cost 5.2% of income (or $83 per month for an individual, using the FPL in 2011) – substantially higher than the $39 per month (or 2.5% of income) for the cheapest plan in Massachusetts. Our estimates suggest that at an $83 premium, only about half the eligible population would enroll in the least generous plan in the Massachusetts market.

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 does 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*, it might have, if adverse selection were the primary source of lack of insurance. Our results therefore suggest that adverse selection, while it exists in this market, does not provide a compelling rationale for setting subsidies to encourage most of the market to enroll. In contrast, we provide suggestive evidence of the role of uncompensated care and discuss how it may provide a normative rationale for insurance subsidies.

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. Technical report.

Buchanan, J. (1975). The samaritan's dilemma.

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. *Henry J Kaiser Family Foundation*.

Cutler, D. M. and J. Gruber (1996). The effect of Medicaid expansions on public insurance, private insurance, and redistribution. *American Economic Review 86(2)*, 378–383.

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.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2016). The economic consequences of hospital admissions. Technical report, National Bureau of Economic Research.

Ehrlich, I. and G. S. Becker (1972). Market insurance, self-insurance, and self-protection. *Journal of Political Economy 80(4)*, 623–648.

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 (2010a). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics 125*(3), 877–921.

Einav, L., A. Finkelstein, and J. Levin (2010b). 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. Technical report, National Bureau of Economic Research.

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. *Working Paper*.

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. (2013). The policy elasticity. *NBER Working Paper Series 19177.*

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

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. Technical report, National Bureau of Economic Research.

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

Kaplow, L. (2006). Public goods and the distribution of income. *European Economic Review 50*(7), 1627–1660.

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.

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

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.

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

Shepard, M. (2016). Hospital network competition and adverse selection: Evidence from the Massachusetts health insurance exchange.

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 to

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, FY 2011, spans July 2010-June 2011. 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 analogously 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. We discuss each of these restrictions in turn. The top panel of table 5 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 the ACS's person weight).

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 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.[49] Instead, we

[49] We cannot even directly measure Medicaid enrollment; the ACS does not distinguish between Medicaid and Comm-

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.[50]

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 14 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 14.

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. 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 $EA(s) = (1-s)(C(s) - W(s) - sW'(s))\gamma\frac{1-s}{2}D'(s)$ where $\gamma$is the coefficient of ab-

Care (both are coded as "Medicaid/other public insurance").

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

Figure 14: Estimate of CommCare Eligible Population from ACS Data

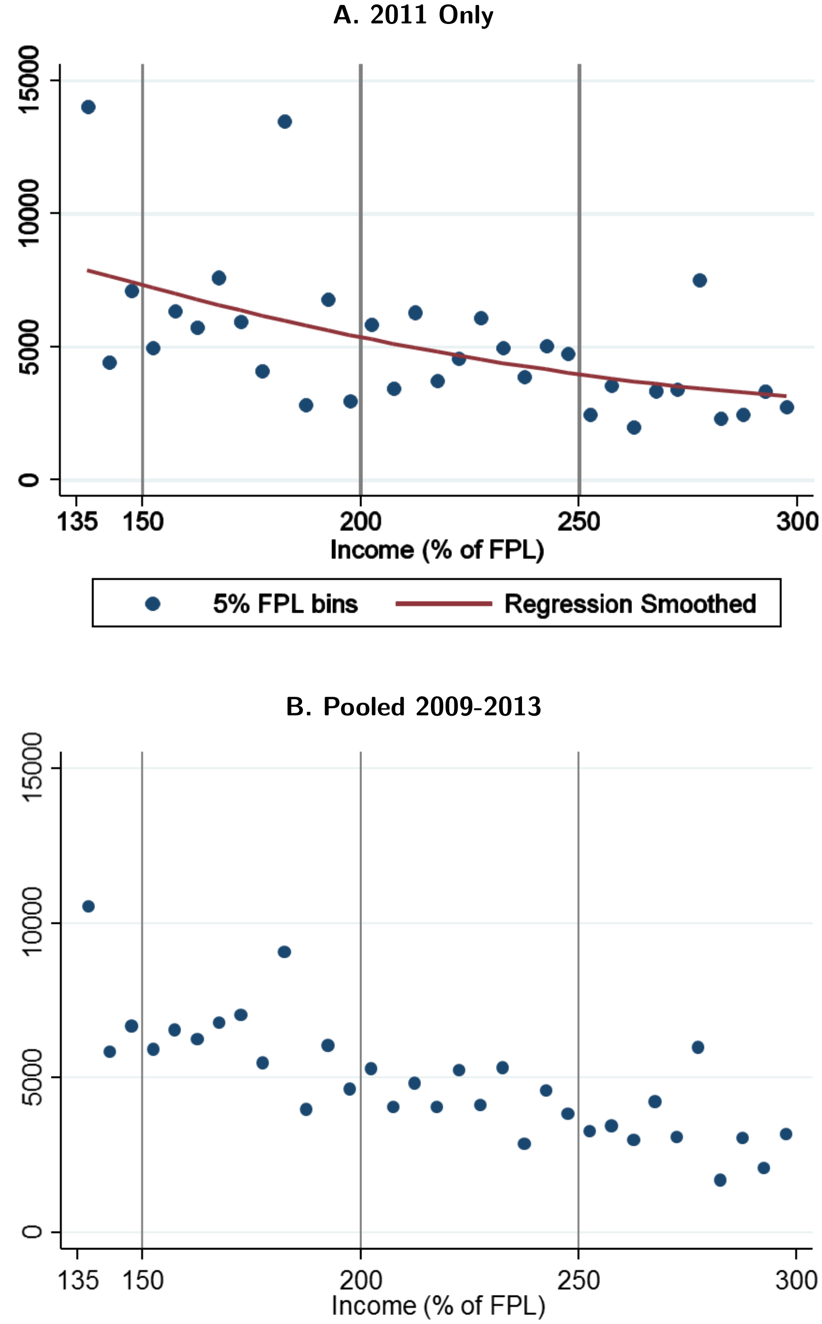

NOTE: The graph shows our smoothed estimates of the CommCare-eligible population from the ACS data. The blue dots are raw estimates of the annual eligible population size (weighting by the ACS "person weight" to generate a population estimate) by 5% of FPL bin. Because these data are relatively noisy – 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 red curve. We use this smoothed eligible population estimate for our estimates of market shares and take-up of CommCare.

Table 5: ACS Sample Construction

Sample / Exclusion	ACS Sample Size # Dropped	% Dropped	# Remaining	Est. Population # 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).

solute risk aversion. 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 $10,000 of annual consumption). Our estimates suggest that even using an *ex-ante* demand measure, willingness to pay would still be substantially below marginal costs. 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 $EA\,(0.5) = 0.5\,(330 - 100 - 0.5 * (-200))\,0.0005 * \frac{0.5}{2}\,(-200) = \$5$ higher ($105 vs $100) than the marginal welfare impact implied by observed demand. This is quite 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 C. 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 public insurance" from this setting, defined as the WTP for each consumer's most preferred plan. This is similar to how we would interpret WTP for Medicaid, where consumers can choose any plan for free. Thus, the WTP for insurance for consumer i is:

$$\text{WTP for Insurance:} \qquad W_i^{Ins} = \max_{j \neq U} \{W_{ij}\}$$

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 this value. 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

$$\text{Choose } U: \qquad W_i^{Ins} \leq \max_{j \neq U} \{p_{ij}\} \equiv p_i^{max} \tag{10}$$

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:[51]

$$\text{Choose } j \neq U: \qquad W_i^{Ins} \geq \min_{j \neq U} \{p_{ij}\} \equiv p_i^{min} \tag{11}$$

We can now map these bounds into bounds on a WTP curve for public insurance. Let s be an index that orders people according to increasing 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 sth quantile of WTP). Suppose that at a given vector of prices (which lets us drop the i subscripts), we observe that s^* share of people choose U, while $1 - s^*$ choose formal insurance. For the marginal type s^*, both conditions (11) and (10) hold, so we can say that:

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

We can use this fact, along with our variation in prices, 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 $\{p^{min}, p^{max}\}$ (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.

Figure 15 shows the resulting estimates. The lower and upper bounds on WTP for insurance are shown (in dashed lines), with our W_L and W_H curves from the vertical model (in solid red and green lines) 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

[51] NOTE: It seems like we could probably get a tighter lower bound using the fact that $W_{ij} \geq p_{ij}$ for the j each individual chose. But I am not sure how to deal with the fact that the j's people choose varies across the discontinuity so have not pursued this.

Figure 15: Bounds on WTP for CommCare Insurance

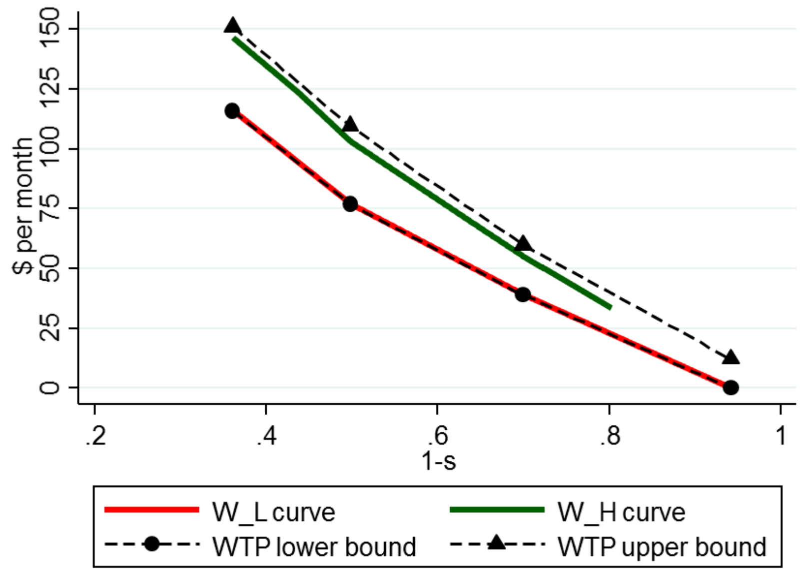

NOTE: The graph shows our estimated bounds on WTP for CommCare, whose construction is described in the text of Appendix C. The lower and upper bounds are shown in black dashed lines (with the point estimates shown in circles and triangles, respectively). For comparison, the graph shows our estimates of W_L (solid red line) and W_H (solid green line) that use the vertical model assumptions.

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 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.

Appendix D. Additional Results

Table 6: Premiums by CommCare Plan

		Income Group (% of FPL)			
CommCare Plan	**H / L**	**100-150%**	**150-200%**	**200-250%**	**250-300%**
BMC	H	$11	$57	$105	$146
Fallon	H	$12	$60	$110	$151
NHP	H	$12	$60	$110	$151
Network	H	$10	$57	$105	$146
CeltiCare	L	$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.

Figure 16: RD for Share in L Plan

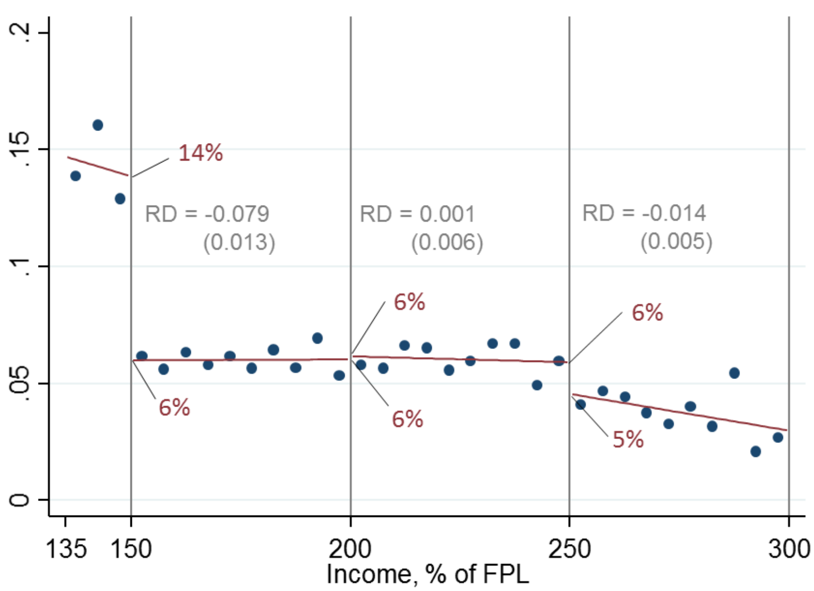

NOTE: The graph shows RD estimates for shares in the L plan, analogous to those shown for any insurance and H in Figure 5 in the text.

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

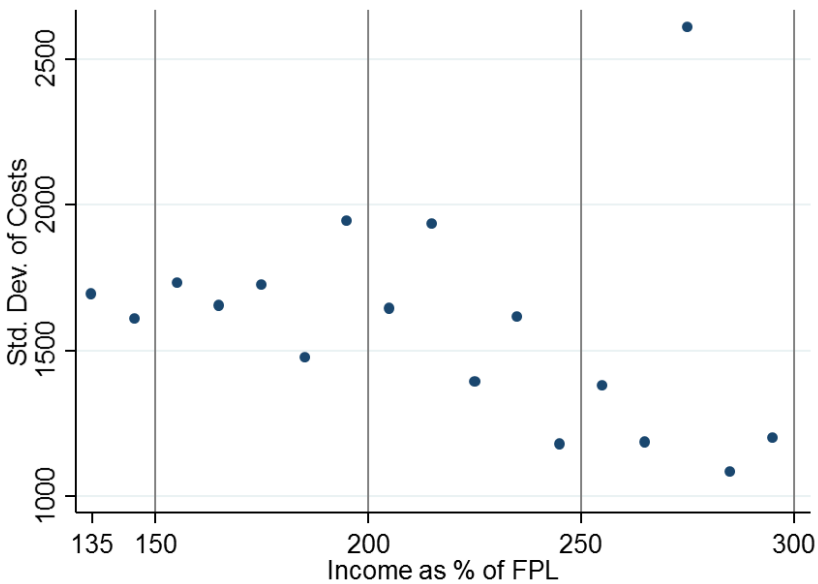

NOTE: The graph shows the standard deviation of insurer costs, by 10% of FPL bin. The standard deviation is calculated across individuals in the data for 2011, using each individuals average insurer-paid cost per month enrolled. As we discuss in the text, there is little evidence that the standard deviation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

Figure 18: Enrollment Counts in CommCare, by Income

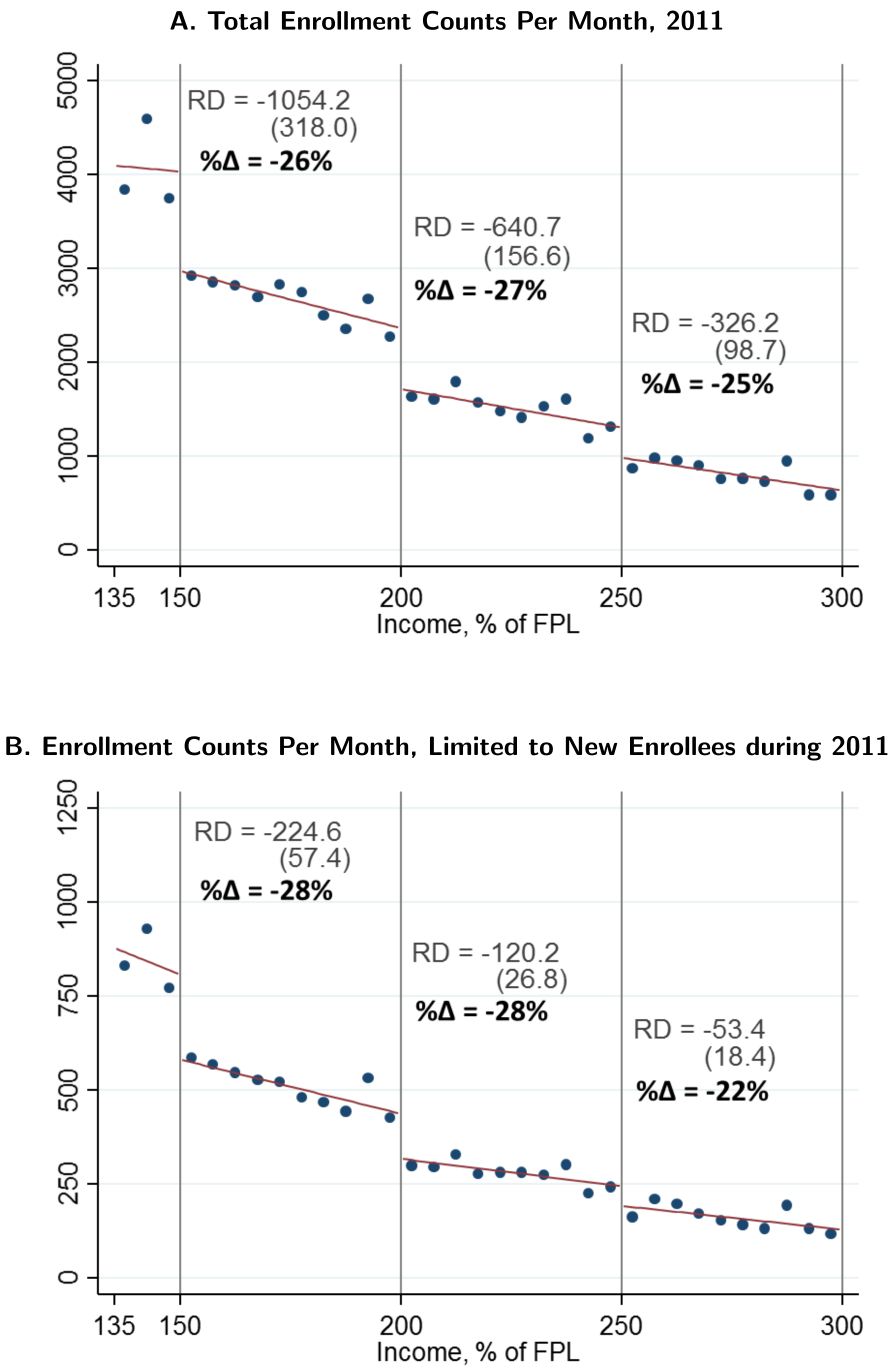

NOTE: The graph shows our baseline RD specification applied to total enrollment counts per month in 2011 (panel A) and enrollment per month limited to new enrollees during 2011 (panel B). The percent reduction in new enrollees is 22-28% at each income discontinuity, quite similar to the 25-27% reduction in total enrollment at the same thresholds. This suggests that inertia is unlikely to be biasing downward our estimates of how much demand falls as premiums rise.

Figure 19: Enrollment Counts in H Plan, by Income

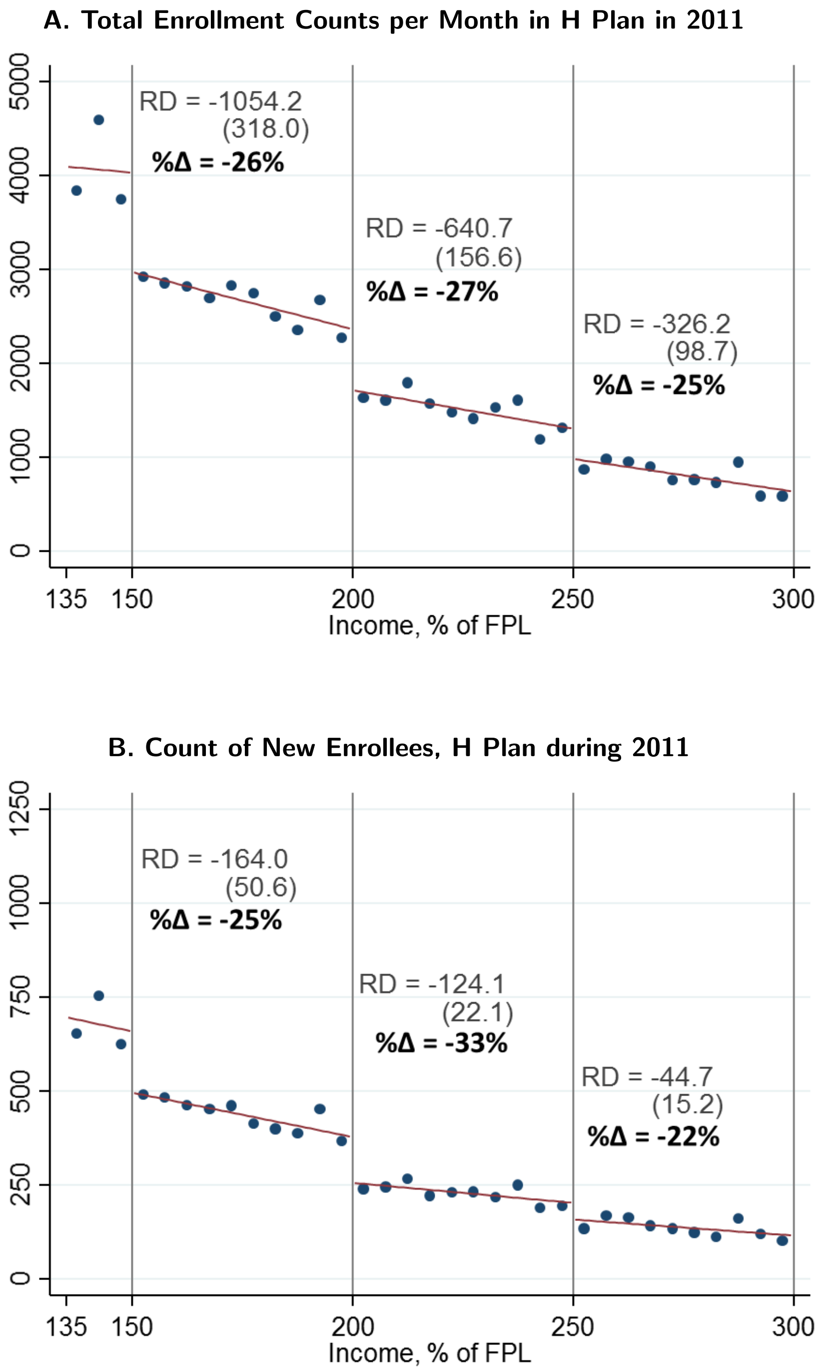

NOTE: These graphs are identical to Figure 18 but applied to counts of enrollees in the H plan only. See the note to Figure 18 for more information.