Moral Hazard Induced Unraveling*

Cameron M. Ellis† Meghan I. Esson‡ Eli Liebman§

August 22, 2023

Abstract

We identify and quantify a new form of welfare loss in insurance markets. We first show theoretically that moral hazard from subsidies for cost-sharing combined with community rating mimics adverse selection and can unravel insurance markets. To quantify the potential welfare loss, we use exogenous variation in the number of subsidized enrollees on the ACA exchanges. We find that subsidy-induced moral hazard led to higher premiums, which has lowered enrollment among the unsubsidized by 7.8 percentage points. We estimate the welfare costs of this "moral hazard induced unraveling" to be around 25% of the welfare loss from existing adverse selection.

JEL Codes: **D82, D60, G22, I13, I18**

Keywords: **Market Failure, Moral Hazard, Insurance, Asymmetric Information, Unraveling**

---

*The authors have no relevant funding sources to disclose. The authors would like to thank Scott Atkinson, Jim Carson, Evan Eastman, Johannes Jasperson, Xiao Lin, Margaux Luflade, Anita Mukherjee, Peter Newberry, Richard Peter, Mark Ragin, Casey Rothschild, Tice Sirmans, Meghan Skira, Bill Vogt, Jackie Volkman-Wise, Sabine Wende, Juan Zhang, and seminar participants at The Southern Economics Association, The American Risk and Insurance Association, The CEAR/MRIC Behavioral Insurance Workshop, The Western Risk and Insurance Association, The Risk Theory Society, The University of Georgia, Temple University, The Federal Reserve Bank of Chicago, Emory University, The University of Iowa, St. John's University, The University of New Hampshire, Florida State University, and Peking University.

†**Contact Author** University of Iowa, Iowa City, IA, 52242, cameron-ellis@uiowa.edu

‡University of Iowa, Iowa City, IA, 52242, meghan-esson@uiowa.edu

§University of Georgia, Athens, GA, 30602, eli.liebman@uga.edu

1 Introduction:

Moral hazard and adverse selection represent two of the most studied market failures in insurance markets. While both are understood to lead to the underprovision of generous insurance (Pauly, 1968; Akerlof, 1970; Rothschild and Stiglitz, 1976; Holmström, 1979; Einav and Finkelstein, 2018), the common wisdom for how to address these market forces differs significantly. Adverse selection is often characterized by externalities from high-risk participants, which can unravel markets (Akerlof, 1970; Rothschild and Stiglitz, 1976). Externalities are a neoclassical justification for policy intervention, and risk adjustment is used in many markets to try to address these externalities. Moral hazard usually is not considered as strong of a case for policy intervention, unless the government possesses a significant informational advantage over the private sector (Cutler and Zeckhauser, 2000; Gaynor *et al.*, 2000; Einav *et al.*, 2010).[1]

In this paper, we show how moral hazard can create externalities that can lead to unraveling in markets with community rating and subsidies for cost-sharing.[2] While other empirical work has explored the potential unraveling effect of adverse selection (e.g., Cutler and Reber, 1998; Buchmueller and DiNardo, 2002; Shepard, 2022), we are the first to show it can also arise through moral hazard.

We begin by showing theoretically how what we term "moral hazard induced unraveling" (MHIU) mimics (and is amplified by) adverse selection, and can induce a spiral similar to Akerlof (1970) that lowers enrollment and potentially unravels the market. Then, to provide context for the size of potential welfare losses, we show that MHIU is occurring on the Affordable Care Act Exchanges (ACA) exchanges where various institutional details allow us to plausibly identify these externalities and estimate welfare. For unsubsidized consumers on the ACA exchanges, the deadweight loss is not trivial – we find that enrollment is 7.8 percentage points lower than in a counterfactual without MHIU, and the welfare loss from MHIU is roughly one-quarter of the welfare loss caused by adverse selection.

The mechanism for MHIU is intuitive: (1) more generous insurance (i.e., lower cost-sharing or "out-of-pocket" (OOP) costs) causally increases utilization for the subset of the population receiving the subsidies (moral hazard); (2) due to community rating, the increased utilization by the subsidized enrollees raises premiums for everyone; (3) higher premiums lead the unsubsidized

[1] For example, Cutler and Zeckhauser (2000) state that: "Moral hazard is a significant concern in insurance policies, but it is not one that necessarily argues for government intervention. Government insurance policies, after all, may engender just as much moral hazard as private insurance policies. There is a rationale for the government to be involved in goods subject to moral hazard only if the government is better able to monitor or punish moral hazard than the private sector. This is not obviously the case in medical care."

[2] Community rating, which means that all consumers are charged the same premium, regardless of prior claims history or other observables, is a common feature in many U.S. health care markets including Medicare Advantage, the Affordable Care Act Exchanges, and Medicare Part D. In health insurance markets, premiums may be allowed to vary by age, income, or family status, but not by the presence of pre-existing conditions.

enrollees to exit the market; (4) which, in a spiral, increases premiums even further to account for these cheaper enrollees dropping out.

In addition to a more general theoretical model, we also propose a graphical framework, similar to Einav *et al.* (2010), for evaluating the market outcomes and welfare effects of MHIU. We show that even when there is no heterogeneity in health status, and therefore no adverse selection, MHIU can lead to a downward sloping average cost curve.[3] Further, we show that when there is adverse selection, the welfare costs of MHIU can be very large relative to the number of individuals who lose insurance. With more adverse selection, marginal enrollees have a larger gap between their willingness to pay for insurance and their expected cost to insure (i.e., a larger risk premium). When MHIU prices these individuals out of the market, it is pricing out consumers with larger risk premia. This divergence between welfare effects and enrollment effects highlights the importance of empirical work on this topic.

We then use the ACA Health Insurance Exchanges (HIX) to causally test our theory. There are a number of features, such as sharp discontinuities in means-tested out-of-pocket (OOP) subsidies, means-tested premium subsidies (without sharp discontinuities), and the staggered timing of Medicaid expansion, which we use to identify MHIU. The means-testing for the HIX's OOP subsidies is based on income relative to the federal poverty line (FPL). In 2018, 150% of the FPL for a single person was $18,090, yet the average deductible of a silver plan on the HIX was $4,375.[4] At this income level, policyholders with such a large deductible are practically uninsured. To address this, the ACA requires insurers to provide subsidies that reduce the amount of cost-sharing for lower-income consumers. These subsidies are large. Figure 1 shows how these subsidies affect the actuarial value of plans offered to consumers at different income levels. The first column (and first row) shows that a silver plan without subsidies has an actuarial value of 70% – the enrollee pays roughly $30 for each $100 of healthcare expenditures, on average. Individuals with incomes up to 150% FPL (the last column) are eligible for OOP cost subsidies amounting to a 94% actuarial value, which means these individuals pay only $6 for each $100 of healthcare expenditures. The bottom four rows of Figure 1 provide an example of how an insurer might alter cost sharing to meet those actuarial values.

Eligibility for the most generous OOP subsidies features a sharp discontinuity at 150% FPL. In 2018, when an individual with an annual income of $18,000 purchased a silver plan, they were

[3] As discussed in Einav *et al.* (2010), classical adverse selection is characterized by a downward sloping average cost curve. As premiums rise, the healthiest enrollees leave the market first, so average costs are increasing in premiums or decreasing in enrollment. When the average cost curve slopes downward, average costs are above marginal costs, there the equilibrium outcome leads to under-insurance as the equilibrium price will be above the marginal cost.

[4] On the HIXs, there are five possible tiers of plans, which are identified based on metal levels, from which individuals can select and purchase private health insurance. The metal levels are categorized based on their actuarial value, which is the expected share of healthcare spending the plan covers: catastrophic, bronze (60%), silver (70%), gold (80%), and platinum (90%). Silver plans are considered the standard, and most subsidies are tied to silver plans. Sprung and Anderson (2018) find that 80-90% of those in the 100-200% FPL range select a silver plan.

Figure 1: Example of Cost Sharing Reduction (CSR) Subsidy Levels:

	Standard Silver – No CSR	CSR Plan for 201-250% FPL	CSR Plan for 151-200% FPL	CSR Plan for up to 150% FPL
Actuarial Value	70% AV	73% AV	87% AV	94% AV
Deductible (Individual)	$7,150	$4,500	$800	$250
Maximum OOP Limit (Individual)	$7,350	$5,700	$1,700	$550
Inpatient hospital	30% (after deductible)	30% (after deductible)	10% (after deductible)	10% (after deductible)
Physician visit	$70	$30	$10	$5

Note: This figure shows how the different income levels correspond to different actuarial values. Insurers are allowed to use any combination of reducing copays, coinsurance, and deductibles, to reach a higher actuarial value. This figure provides examples of how an insurer might change the characteristics of the plan to meet the actuarial value.
Source: Health Reform Beyond the Basics

eligible for a silver plan twice as generous as one purchased by an individual with an income of $18,200 (94% actuarial value vs. 87% actuarial value). The average cost for this population, along with the RAND Health Insurance elasticity, implies that the 50% reduction in out-of-pocket prices would lead to a $721 increase in overall costs. That is, this low income group is more expensive, in large part due to moral hazard from the OOP subsidies.[5]

We empirically show how this subsidy-induced moral hazard imposes an externality on unsubsidized consumers. We exploit the timing of Medicaid expansions across U.S. states as plausibly exogenous variation in the aggregate moral hazard costs on each HIX. When a state expands Medicaid, those below 138% of the FPL are no longer eligible for subsidies as they become eligible for Medicaid. These individuals typically leave the HIX (where they are no longer subsidized) and enroll in Medicaid instead. Prior to Medicaid expansion, 40% of enrollees on the exchange were eligible for the most generous subsidies. Medicaid expansion cuts it in half. The size of the cost differences for these enrollees and the share of them that leave the HIX after Medicaid expansion suggests that Medicaid expansion has the potential to have a sizeable impact on the HIX's risk pools.

Indeed, we find that the Medicaid expansion is associated with a 12% reduction in premiums, and a 4% decrease in the uninsured rate for the unsubsidized population (i.e., those above 400% FPL who should not be impacted by Medicaid expansion directly). The magnitudes of our esti-

[5] As we discuss later, we use the sharp discontinuity to test for moral hazard using the MEPS data. We find a large effect, but the data are noisy, so we conservatively stick to the RAND elasticities.

mates suggest that, in the context of the ACA, moral hazard induced unraveling is economically meaningful: a 4% decrease in the uninsured rate for those above 400% FPL would lead to 120,000 fewer uninsured people if all states expanded Medicaid (compared to if none had). Furthermore, as Medicaid only reduced the share of highly subsidized enrollees by half, an extrapolation of our estimates suggest that if 100% of these subsidized enrollees left the exchanges the share of the unsubsidized population purchasing insurance would rise by 7.8 percentage points.

To isolate MHIU more precisely and quantify the welfare impacts, we calibrate our theoretical model. We use the Medicaid expansion as a cost-shock, which provides plausibly exogenous variation in premiums, to estimate demand for insurance for unsubsidized consumers. We then combine this demand function with the OOP price elasticity estimates from the RAND Health Insurance Experiment and the slope of the average cost curve coefficient from Einav *et al.* (2010) to quantify the welfare loss associated with MHIU. In our baseline case, we find that adverse selection reduces enrollment by 17.2 percentage points while MHIU reduces it by 2.8 percentage points. Adverse selection and MHIU combine to induce a deadweight loss of $175 per person – $1.75 billion in total for the potential market – with MHIU contributing one-fourth and adverse selection contributing the remaining three-fourths.

We then examine the sensitivity of this estimate to different assumptions on the level of moral hazard, adverse selection, insurer markups, and a health/wealth gradient. To vary the amount of moral hazard, we allow for OOP price elasticities that are larger (such as in Ellis *et al.* (2017) and Brot-Goldberg *et al.* (2017)) or smaller (such as in Lavetti *et al.* (2023)) than the RAND estimates. We find that decreasing the amount of moral hazard reduces the welfare loss of MHIU. To allow for markups, we assume that the Medical Loss Ratio regulation, which limits insurers to 15% profits, binds. We find that allowing for markups increases the welfare loss of MHIU. Finally, we examine how changing the amount of adverse selection (by changing the slope parameter of the average cost curve) changes equilibrium outcomes. Increasing the degree of adverse selection has a dramatic impact on the welfare loss from MHIU alone. The welfare loss due to MHIU alone from our highest adverse selection assumption (where the slope coefficient is twice that of Einav *et al.* (2010)) is 17 times larger than the loss assuming no adverse selection, though the enrollment losses are only 1.5 higher. The differing relative magnitudes of the results on enrollment and welfare occur because each consumer priced out by MHIU has a higher potential surplus than those priced out by adverse selection. Regardless of our assumptions, we find that the welfare loss from MHIU is not trivial relative to adverse selection. The lowest estimate we find is that MHIU leads to a welfare loss that is 18% of the size of adverse selection.

Finally, as a mechanism check for moral hazard, we use the sharp discontinuity in out-of-pocket costs to directly test for moral hazard on the exchanges. Using the Medical Expenditure Panel Survey (MEPS) data, we find that those with the most generous OOP subsidies consume approximately \$1,700 more healthcare than their (still highly subsidized) counterparts immediately above 150% of the FPL, despite an annual income difference of only a few hundred dollars. Our estimate is consistent with a price elasticity of -0.47, which lies between the estimates of Manning *et al.* (1987), who examine variation in coinsurance rates, and Brot-Goldberg *et al.* (2017), who examine variation in deductibles. This is reassuring since HIX insurers often alter both to achieve target actuarial values. However, given the data limitations of the MEPS, we use the more conservative RAND estimates as our base case throughout the paper. However, we do provide evidence of the robustness of the MEPS results in Online Appendix C.

Our study contributes to the literature in two primary ways. First, we document a novel form of welfare loss from moral hazard. A classical result from the literature is that moral hazard from reduced cost sharing leads to the over-consumption of health care. Further, this over-consumption can lead to the under-consumption of insurance (Pauly, 1968; Holmström, 1979; Einav and Finkelstein, 2018). We show that, in settings with community rating, there is also an externality on other consumers who do not receive such generous insurance plans. This externality mimics the effect of adverse selection and can potentially unravel the market. While other empirical work has explored the potential unraveling effect of adverse selection (e.g., Cutler and Reber, 1998; Buchmueller and DiNardo, 2002; Shepard, 2022), we are the first to show it can also arise through moral hazard.

Second, we contribute to the further understanding of the interaction of the various "legs" of the ACA (Handel and Kolstad, 2021). We show that three of the law's most studied policies – means-tested subsidies, community rating, and Medicaid expansion – combine to have a large and unexpected impact. Additionally, our work confirms the results of Sen and DeLeire (2018) and Peng (2017), who also find that Medicaid expansion lowered premiums and put forth the hypothesis that this is due to either the health/wealth gradient or adverse selection. We instead offer an alternative and intuitive rationale for why premiums decrease so much following Medicaid expansion by connecting with, and confirming through a separate identification channel, the results of Lavetti *et al.* (2023) that the OOP subsidies lead to moral hazard.

Our work also has important policy implications. Akerlof (1970) and others highlight the role of policy intervention in adverse selection – where markets can cease to exist if the selection is too intense. MHIU mimics adverse selection, leading to downward-sloping cost curves and externalities across individuals. Akin to risk adjustment, reimbursing insurers for the moral hazard costs is a potential solution and a "reverse" natural experiment has already occurred. Prior to 2018, the government reimbursed insurers for the direct costs of CSR subsidies (e.g., if the copay was reduced from \$25 to \$5, the government would reimburse the insurer \$20). However, this reimbursement would not fully reimburse an insurer for extra services used because the copay was lower. In 2018, the Trump administration stopped these reimbursements. Insurers dramatically raised premiums, which led to higher premium subsidies and was projected to *increase* overall government expenditures (Congressional Budget Office, 2017).

2 Institutional Details

2.1 Health Insurance Exchanges:

Among other regulations, the ACA mandated that beginning in 2014 all states were required to have a centralized individual marketplace called a Health Insurance Exchange (HIX). The HIX contains the vast majority of the individual market for health insurance, with 11.75 million enrollees in 2018. On the HIXs there are five possible tiers of plans, which are identified based on metal levels, from which individuals can select and purchase private health insurance. The metal levels are categorized based on their actuarial value, which is the expected share of healthcare spending the plan covers: catastrophic, bronze (60%), silver (70%), gold (80%), and platinum (90%).

Plans are required to cover certain essential health benefits and have a maximum out-of-pocket expenditure no higher than $6,350 for individuals or $12,700 for families. Otherwise, insurers have great flexibility in designing their policies, so long as the insurers are within 2% of the targeted actuarial-value. Premiums are required to be community-rated, meaning that insurers must charge the same premium to all consumers in a rating area except that insurers can vary premiums by age (3-1 limit), tobacco use (1.5-1 limit), and family composition. Rating areas are determined by each state. Most states have defined their rating areas based on counties. Aside from these restrictions, insurers are free to set their own initial premiums, but any rise in premiums greater than 10% is subject to approval by a state board.

Cost Sharing and Premium Subsidies: Any consumer, regardless of income, can purchase insurance from the HIX marketplaces. However, the ACA provides financial assistance to lower-income consumers to purchase health insurance on the HIXs. To better understand the rationale for the subsidies, consider a 40-year-old single individual with a 2018 income of $16,040, which is 150% of the FPL. Without any subsidies, the average deductible for a silver plan in 2019 was $4,375 (Fehr *et al.*, 2019) and the average annual premium for a benchmark silver plan was $5,736 with substantial geographic variation, ranging from $3,912 in Minnesota to $10,380 in Wyoming (Kaiser Family Foundation, 2019). Such high deductibles and premiums mean that the consumer would have to pay more than half their income before they are eligible for reduced out-of-pocket costs through coinsurance or copays.

Insurers who participate in the HIXs are required to offer at least one silver plan, and each silver plan must have versions with reduced cost sharing (i.e., more generous coverage) for lower-income enrollees. The CSR subsidies can lower the amount the consumer pays for deductibles, coinsurance, copay, and the out-of-pocket maximum, and insurers have flexibility with regards to how they achieve the lower cost. The difference between the standard version and the reduced cost sharing version is the CSR subsidy. During our sample period, insurers were compensated by the federal government for the direct cost of these reductions. For example, if the unsubsidized version of the plan has a $25 copay and the subsidized version has a $5 copay, then the government would reimburse $20 per visit. However, the government does not reimburse for what we term the moral hazard cost. If the reduced copay leads to increased utilization, that extra cost is not reimbursed and is therefore partially borne by the insurer.

A traditional silver level plan has an actuarial value of 70%. If the purchaser's income is from 200% to 250% of the FPL, the actuarial value of that plan jumps to 73%; from 150% to 200%, it increases to 87%; from 100% to 150%, up to 94%. The plans are the same as the traditional silver version in every dimension other than cost sharing. CSRs are only available to consumers who purchase silver level plans and are based solely on income at the time of purchase (DeLeire *et al.*, 2017).[6]

HIX enrollees whose incomes are in the 100-400% FPL range are also eligible for Advanced Premium Tax Credits (APTCs) which reduce premium costs. Under the APTC, the subsidy amount is set by the second lowest cost silver plan's (SLCSP) premium and the consumer's income level.[7] An individual at 150% of the FPL would pay no more than 4.03% of their income (or $646) for the SLCSP regardless of the premium an unsubsidized consumer would pay (Centers for Medicare and Medicaid Services, 2019). The difference in the premium for the SLCSP and the premium cap amount is the subsidy amount which can be used to purchase any plan. Therefore, there may be some variation in the menu of premiums offered across rating areas.[8] It is important to note that those with premium subsidies do not face marginal changes in the premiums the plans set (for the SLCSP), premiums are set by income level, the subsidy received varies with the price of the plan.

[6] When purchasing insurance through the HIXs, the consumer is required to enter their income before seeing plan options, and the consumer only sees the CSR plans which correspond to their income level. That is, consumers can not "accidentally" choose the wrong CSR plan. Given that the subsidy is only available to those who purchase a silver plan, it is not surprising that silver plans have a market share of roughly 80 to 90 percent in the 100-200% FPL range (Sprung and Anderson, 2018).

[7] The percentage of income for the required contribution varies smoothly with income: there are no discontinuities. The amount varies from 2.01% of income at 100% FPL to 9.56% at 400% FPL. If the SLCSP's premium is lower than the expected contribution, a person below 400% FPL may not receive a subsidy. This would typically affect those close to the 400% FPL cutoff.

[8] To receive the CSR subsidy the consumer must choose a silver plan, which for lower-income consumers reduces the incentive to select a cheaper bronze plan. See Drake and Anderson (2020) for a discussion of free bronze plans.

2.2 Medicaid Expansion:

The ACA mandated that states alter Medicaid eligibility requirements by (1) raising the income ceiling for eligibility from 100% to 138% of the FPL and (2) allowing non-disabled, childless adults with income below this line access to Medicaid; potentially extending coverage to more than 20 million Americans (Holahan *et al.*, 2012). While the majority of the ACA held up to judicial scrutiny, the Supreme Court ruled the mandatory aspect of the Medicaid expansion was unconstitutional. The ruling allowed states to "opt-out" of the Medicaid expansion and keep their pre-ACA enrollment criteria.

As of 2017, 16 states had opted-out of the Medicaid expansion of the ACA, citing budgetary constraints or costs as their main point of opposition. Figure 2 shows the timeline of each state's decisions to expand Medicaid.[9]

Figure 2: Timeline of Medicaid Expansion:

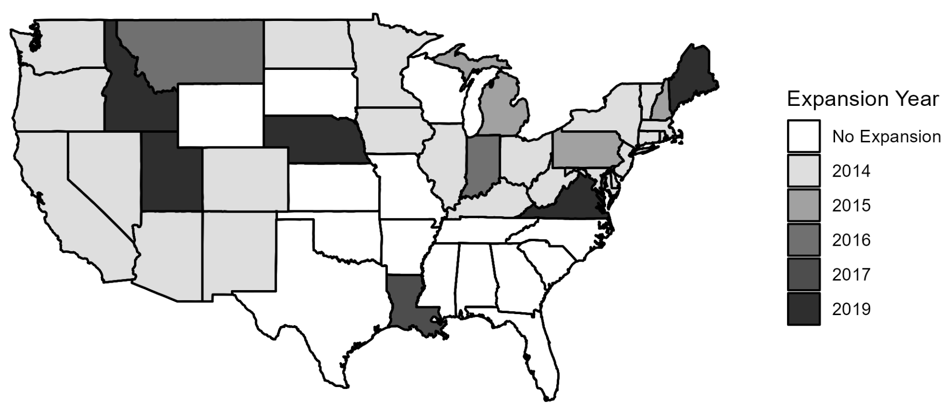

Note: This figure shows the timeline of ACA Medicaid expansion in the US.

When a state expands Medicaid, previously ineligible individuals (i.e., non-disabled, childless adults) below 138% of the FPL are eligible for Medicaid, which makes them ineligible for subsidies on the HIX. The consequence of these cutoffs is that in states where Medicaid has not expanded, individuals from 100% to 150% of the FPL are eligible for the most generous CSRs.[10] In states that have expanded, those from 138% to 150% of the FPL are eligible for the subsidies. Therefore,

[9] A few states received waivers for the proposed expansion by providing a plan for how to increase coverage. For example, rather than including those newly eligible in Medicaid, Arkansas uses the federal funds to subsidize purchases on the Arkansas Health Insurance Exchange. For our analysis, we follow Sen and DeLeire (2018) and treat Arkansas as an opt-out state because the lower general health population is partaking in the private market. Omitting Arkansas or treating it as an expansionary state does not qualitatively change our results (or the significance).

[10] Those who are under 100% of the FPL in non-expansion states are ineligible for both the APTC and CSRs. They are also not required to pay the shared responsibility payment also called individual mandate penalty. This leads to a coverage gap as they may also be ineligible for Medicaid (Rosenbaum and Wilensky, 2020).

when a state expands Medicaid, we should expect individuals from 100% to 138% of the FPL to switch from highly subsidized exchange plans (which affects the risk pool on the HIX) to Medicaid coverage.

3 Theory

3.1 Graphical Intuition

We begin with a graphical example of moral hazard induced unraveling using the framework of Einav and Finkelstein (2011). In Sections 3.1 through 3.4, we abstract from adverse selection to show how MHIU works, then introduce and discuss adverse selection in Section 3.5.

Figure 3, Panel (a) provides a stylized graphical representation of a market with MHIU. There are two types of consumers: subsidized (S) and unsubsidized (U). Subsidized consumers are subsidized on two margins, premiums and out-of-pocket costs. Because of premium subsidies, subsidized consumers pay a flat premium determined exogenously by the government. This makes their demand curve vertical, since they are immune to premium variation by the insurer. Unsubsidized consumers have a downward sloping demand curve.

To differentiate between adverse selection and MHIU, we begin by assuming away adverse selection. In particular, we assume all consumers would have the same health care costs if they faced the same out of pocket costs. Therefore, the marginal cost curve is flat within a group. However, subsidized consumers also receive subsidies which lower their OOP costs, hence moral hazard shifts their marginal cost curve upwards by a uniform amount, as shown by the difference between MCˆS and MCˆU in panel (a) of Figure 3. This makes the overall marginal cost curve a step function.

Average costs are just the weighted average of the subsidized and unsubsidized types' costs, and match marginal costs if subsidized consumers are the only ones in the market. At low enough premium levels, unsubsidized consumers start entering the market and, as premiums decline, more unsubsidized consumers enter. Since unsubsidized consumers are cheaper to insure, the average cost curve slopes downward.

The focus of this paper is on the externality to the unsubsidized types. Panel (b) of Figure 3 shows the same graph zoomed in on the unsubsidized population. Again, we have a constant MCˆU curve. Since the subsidized consumers are also in the market, the AC curve is downward sloping and a weighted average of the costs from each type.

Now, we derive the efficient and equilibrium outcomes with and without MHIU. First, consider the case with no subsidized consumers. If there are no subsidized consumers, then the average cost

Figure 3: Graphical Representation of Moral Hazard Induced Unraveling

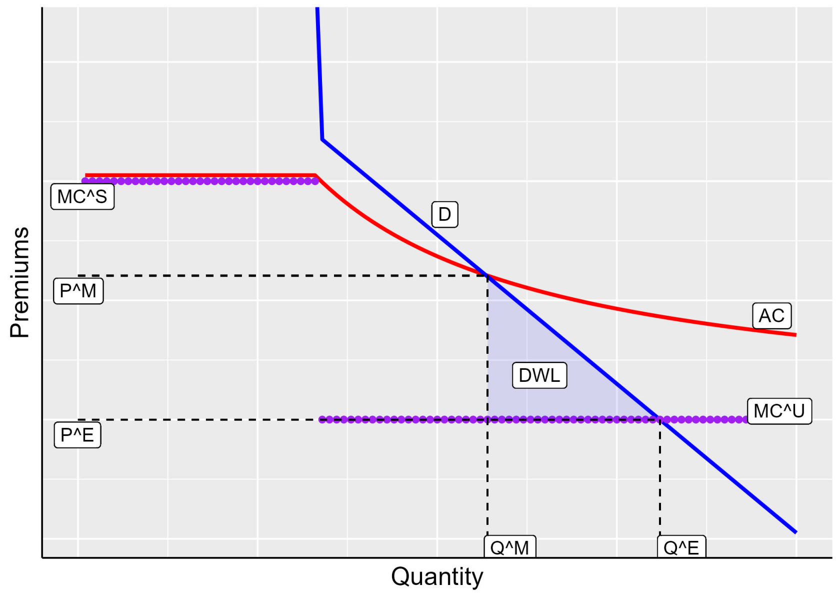

(a) Equilibrium for Entire Population

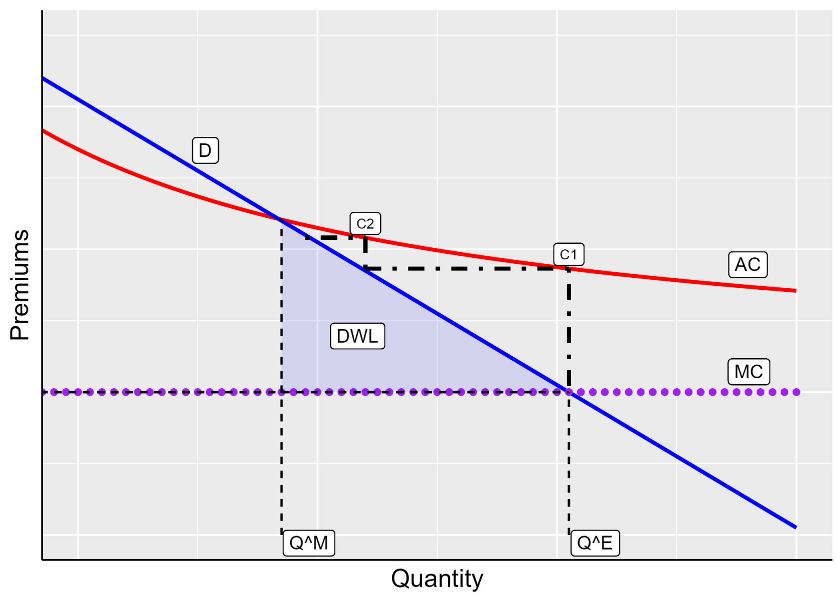

(b) Only the Unsubsidized Population

Note: Panel (a) shows a market with all individuals in the market. Subsidized consumers purchase at any gross premium, because their premiums are subsidized and they have an average and marginal cost of MCˆS. As the gross premium falls unsubsidized types enter, reducing the average cost and increasing the number of insured. Panel (b) shows the same graph, but zoomed in on the unsubsidized population.

curve is constant and equals the MCˆU curve. The efficient allocation is when premiums equal marginal cost, which is at quantity QˆE. As discussed in Einav and Finkelstein (2011), willingness to pay is the sum of risk premia and expected costs. At quantity QˆE, all consumers with positive risk premia purchase insurance, as the premium is equal to their expected cost.[11] The zero-profit equilibrium premium is where the demand and average cost curve intersect. Without MHIU (and without adverse selection) the equilibrium outcome and efficient outcome are the same, QˆE.

If subsidized consumers are added into the market, the efficient outcome is no longer an equilibrium. The subsidized consumers increase the average cost of the insurance pool to $C1$, which is now higher than the efficient premium, PˆE . The average cost $C1$ is not an equilibrium because some unsubsidized consumers are not willing to pay $C1$ and would leave the market. However, reducing the number of unsubsidized consumers means that average costs rise more, unraveling the market to a point $C2$. One can imagine this unraveling spiral repeating until either the market fully unravels or a new equilibrium is reached, which we show in point QˆM. The shaded triangle labeled DWL is the welfare loss of moral hazard induced unraveling. For these consumers, their risk premia meant that their willingness to pay was above their marginal cost. However, they do not purchase insurance because the equilibrium price is higher than their willingness to pay.

This unraveling occurs because the average cost curve is downward sloping. It requires marginal consumers to drop out to raise premiums to create the feedback effects to go from $C1$ to $C2$ to QˆM. In Einav *et al.* (2010), adverse selection was the cause of a downward sloping average cost curve while moral hazard led to a parallel shift in the average and marginal cost curves. In our case, MHIU also creates a downward sloping AC curve, leading to potential unraveling which mimics the welfare loss of adverse selection, even when there is no heterogeneity in health status. Below, we show how adverse selection and MHIU interact, but first we formalize these results allowing for more general demand and cost curves.

3.2 Setup and Notation

Demand and Cost Consider a market where perfectly competitive risk neutral firms offer an identical insurance plan to risk averse individuals, who have a binary choice of purchasing insurance or not. Let there be two types of consumers, L and H, who face low and high out-of-pocket prices: O_L and O_H. In the ACA exchanges, this corresponds to low and high-income consumers where, due to the means-tested OOP subsidies, low-income consumers face lower out-of-pocket prices. We assume that out-of-pocket prices are set exogenously. Under the ACA, the actuarial value of a plan

[11] In Einav and Finkelstein (2011), they note that if all consumers are risk averse and there are no other frictions the efficient outcome may be a corner solution where all individuals purchase insurance and willingness to pay is greater than marginal cost for all consumers due to the risk premium. Graphing the outcome this way would not change the main implications of this analysis.

is set by law, so this assumption equates to consumers not lowering their income to qualify for more subsidies.[12]

Let $M_i \in \mathbb{N}$ denote the number of potential enrollees of type i, $i \in \{L, H\}$. M_i is exogenous and it will be the focus of our comparative statics, because this is the source of exogenous variation in our empirics. We denote demand by $D_i(p) \in [0, 1]$. Demand is the share of potential enrollees of type i who buy insurance, for a given level of premium. $D_i(p) \cdot M_i$ is the number of enrollees of type i. Recall that in the textbook case, demand for insurance has two parts, a consumers risk (or expected cost) and their risk premium due to their risk average. A consumer may have higher willingness to pay for insurance if they expect to cost more or because they are more risk averse and hence have a larger risk premium.

We assume that the *expected* cost of insuring a patient, denoted $c(O_i)$, depends *only* on the patient's out-of-pocket price and that moral hazard exists: lower out-of-pocket prices cause consumers to seek more treatment, and thus costs rise as out-of-pocket prices fall. This assumes away adverse selection, which helps us show how MHIU can mimic adverse selection. We relax this assumption in Section 3.5.

Supply and Equilibrium We assume that by regulation there is community rating, so firms can only charge one premium in the market. We define the one (gross) premium in the market as p. However, for ease of exposition (and to mirror the ACA), Type L and Type H consumers face different net premiums due to premium subsidies. Type H consumers do not receive premium subsidies and pay $p \in P > 0$. Type L face an exogenously set premium, p_L, where the insurer receives the difference through a government subsidy, $S \equiv p - p_L$ if $p > p_L$ and equal to 0 otherwise.[13]

Equation 1 defines $\bar{c}(p)$ as the average cost across the two types. It is simply a weighted average of the average costs of the L and H types. Average cost is a function of premiums because it changes the share of H types in the market. For the L types, demand is a function of p_L which is constant in premiums, so L type demand is constant in gross premiums.

$$\bar{c}(p) = \frac{D_L(p_L) \cdot M_L \cdot c(O_L) + D_H(p) \cdot M_H \cdot c(O_H)}{D_L(p_L) \cdot M_L + D_H(p) \cdot M_H} \tag{1}$$

[12] This appears to hold in our data as there is no evidence of bunching at the 150% FPL threshold (see Appendix C.5).

[13] In the ACA exchanges, there is also a third group who face high out-of-pocket prices but subsidized premiums. For ease of exposition, we ignore them. Inclusion of this group does not alter the comparative statics, but will mute their magnitude.

As is common in the literature, we consider a symmetric, zero profit equilibrium and, similar to Einav *et al.* (2010) and Akerlof (1970), we take product quality as exogenously set.[14]

3.3 Assumptions for General Analysis

First, we assume that demand is (a) continuous and differentiable, (b) downward sloping, and (c) positive.

Assumption A1: Demand Curve

a) $D_i \in C^1(P, [0, 1])$

b) $D_i'(p) < 0$

c) $D_L(p_L) > 0$ and $D_H(p) > 0$

Next, we assume that patients with lower out-of-pocket prices have higher expected costs, i.e. there is moral hazard. To orient the reader to the empirics, Section 4.3 shows that $c(O_L) - c(O_H) > 0$.[15]

Assumption A2: Moral Hazard $c(O_L) > c(O_H)$ if $O_L < O_H$.

We also assume that average costs, $\bar{c}$, do not rise too quickly in premiums. Raising p by \$1 will lead to fewer high-income (low-cost) types. The loss of high-income types will in turn increase average costs. Assumption A3 requires that average costs do not increase by more than \$1.[16]

Assumption A3: $\frac{\partial \bar{c}}{\partial p} < 1$

Finally, we limit ourselves to considering a symmetric, zero profit equilibrium. In order to prove the equilibrium exists via a fixed point theorem, we also require that P be convex and compact.

Assumption A4: P is convex and compact

[14] An extension of our model, in the style of Azevedo and Gottlieb (2017), allowing insurers to endogenously select what kind of cost sharing (deductibles, co-pays, etc.) to change would be an excellent avenue for future research.

[15] Lavetti *et al.* (2023) also supports this assumption.

[16] This is a condition on the functional form of demand and the difference in costs between high- and low-income types. We show the exact formula for $\frac{\partial \bar{c}}{\partial p}$ in Lemma 1 in Appendix A.

3.4 Results for General Analysis

Due to the prominence of non-existence results in unraveling insurance models, we begin by showing that an equilibrium exists in our model. All proofs are left to Appendix A.

Proposition 1: Existence of an Equilibrium Premium, p^*

A zero profit equilibrium implies that all insurers price at average cost. Average costs are simply the weighted average cost of consumers in each group which is positive from Assumption A1c:

$$\bar{c}(p) = \frac{D_L(p_L) \cdot M_L \cdot c(O_L) + D_H(p) \cdot M_H \cdot c(O_H)}{D_L(p_L) \cdot M_L + D_H(p) \cdot M_H} \tag{2}$$

Under assumptions A1 and A4, there exists a $p^* \in P$ such that $\bar{c}(p^*) = p^*$.

Our first testable prediction is to show that adding more low out-of-pocket price (type L) consumers to the eligible population will cause (unsubsidized) premiums to rise. Our empirical analog in Section 4.1 shows that (1) low out-of-pocket price consumers are removed from the market by the Medicaid expansion; and that (2) premiums fall due to this reduction in M_L.

Result 1: Adding subsidized consumers leads to higher equilibrium premiums: $\frac{\partial p^*}{\partial M_L} > 0$

The intuition for the proof is simple: adding higher cost consumers to the risk pool increases average costs and therefore premiums.

Next we show how MHIU can lead to unraveling. That is, equilibrium premium changes ($\frac{\partial p^*}{\partial M_L}$) are *larger* than changes in costs directly due to the initial cost shock due to the change in M_L ($\frac{\partial \bar{c}}{\partial M_L}$).

Result 2: Unraveling: $\frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L} > 0$

As premiums rise, fewer high-income consumers enroll. Our empirical analog to this result is shown in Section 4.1.3, which shows that Medicaid expansion (which lowers premiums) leads to more unsubsidized consumers purchasing insurance and fewer uninsured. Having fewer high cost enrollees amplifies the effect of cost increases: as fewer high-income types enroll, average costs rise even more. This suggests that moral hazard can start a spiral under community rating. Higher premiums lead to fewer high-income (low-cost) consumers, which raises premiums again. Then the cycle is repeated.[17] Moral hazard leads to welfare loss from overconsumption of healthcare. When combined with community rating, moral hazard creates an externality that creates a spiral similar in spirit to an adverse selection death spiral, though we have assumed away heterogeneity in health status.

3.5 Incorporating Adverse Selection

What happens if there is adverse selection in the market? As in Einav *et al.* (2010), when there is adverse selection, the least costly individuals to insure also have the lowest willingness to pay, so they will be the first to exit the market as premiums rise. This implies that the average (and marginal) costs are a function of premiums. Now, consider a set of potential enrollees who vary in their costs, both due to risk type and out-of-pocket costs. We define $ac_H(p, O_H)$, as the average cost for unsubsidized consumers who choose to enroll, where enrollment is a function of premiums. Likewise, $ac_L(p_L, O_L)$ is the average cost for the subsidized consumers. We define adverse selection locally as average costs for a population rising when premiums rise.

Definition: Adverse selection among the unsubsidized population $\frac{\partial ac_H(p,O_H)}{\partial p} > 0$

Notice that this is capturing the average cost changes within the unsubsidized population, which were previously assumed away.

When adverse selection is worse, we show that unraveling due to MHIU is worse. We define adverse selection being worse to mean the localized slope of the average cost curve is steeper (the derivative is larger). Formally, we show that the difference between the ultimate equilibrium price and the initial cost shock (due to MHIU) is larger when adverse selection is worse.[18]

Result 3: $\frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L}$ is larger when adverse selection is worse.

The proof and formal statement of the result is left to the appendix. The key part of the proof is that adding high-cost subsidized enrollees increases average costs. When average costs increase, there are two reasons that the equilibrium premium rises. First, more low-cost unsubsidized consumers exit, which was true under MHIU. Second, the unsubsidized consumers who exit are lower cost

[17] The market may not fully unravel so long as $D_H(c(O_L)) > 0$ because some high-income types will still enroll keeping average costs below $c(O_L)$.

[18] Our result relies on small changes as we are taking derivatives. In this context, $\frac{\partial ac_H(p,O_H)}{\partial p}$ is a scalar corresponding to the rate of change of the average cost curve locally around the equilibrium. In this case, more adverse selection means the slope is steeper. In Section 3.6, this corresponds to changing δ, the slope of the average cost curve, when we assume linear average costs for the unsubsidized.

than the average unsubsidized consumer who remains, because of adverse selection. This increases $ac_H(p, O_H)$ even more. Hence the impact on equilibrium premiums is even larger than it would have been if only the first mechanism were taken into account. That is, MHIU and Akerlof (1970) style unraveling complement each other.

Figure 4 presents the graphical model of MHIU when there is adverse selection. Now, the AC and MC curves without MHIU are now downward sloping because reducing premiums means more people will enroll in insurance, and if there is adverse selection these marginal consumers are healthier than the average consumer.

Figure 4: Graphical Representation of Moral Hazard Induced Unraveling, with Adverse Selection

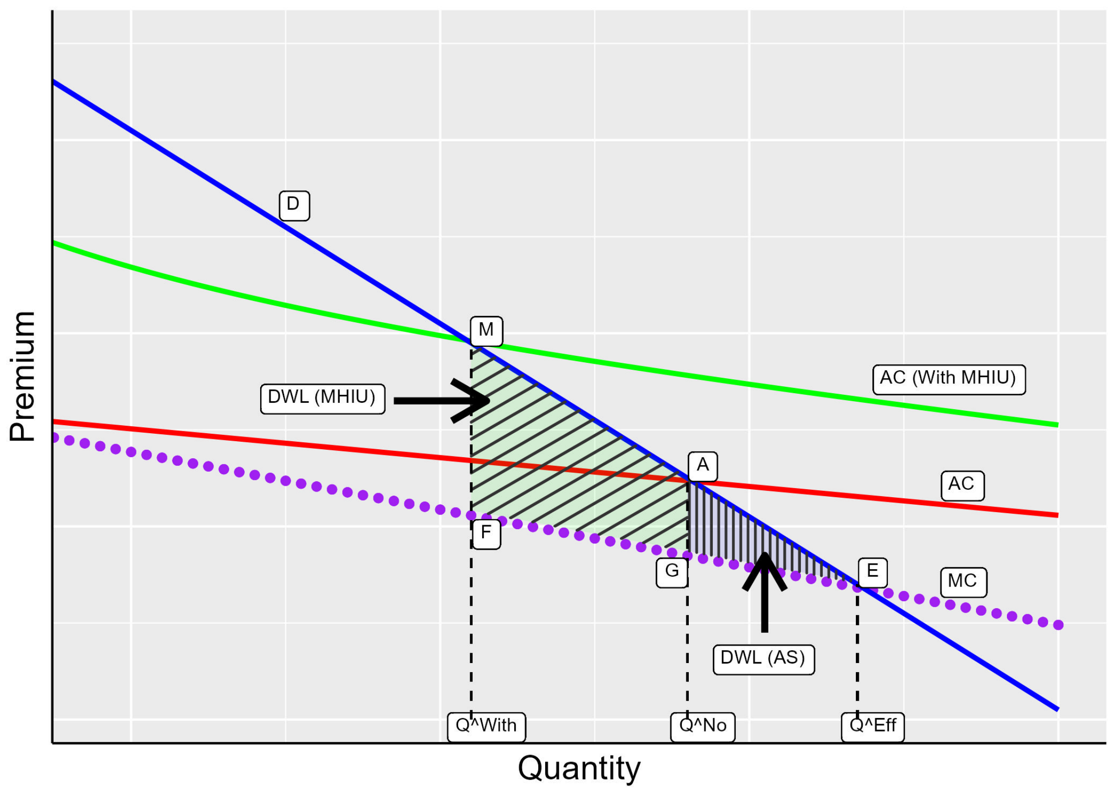

Note: This figure shows what happens when there is both MHIU and adverse selection. Where the demand curve intersects the marginal cost curve is the efficient outcome (Point E). With adverse selection, but not MHIU, the equilibrium is where the demand curve intersects the average cost curve without MHIU, Point A. When there is MHIU, the equilibrium is where the demand curve intersects the average cost curve with MHIU, Point M.

The efficient outcome is where premium equals marginal cost, which occurs at Point E. Even with no MHIU, this is now not an equilibrium as charging this premium would be unprofitable as the average cost curve is above the marginal cost curve. The equilibrium outcome with only adverse selection is point A, where premium equals average cost. Therefore, the vertically barred triangle labeled DWL (AS) is the welfare loss from classical adverse selection alone.

Adding in MHIU shifts the average cost curve up and steepens the slope. This leads to higher premiums and lower enrollment, potentially at a new equilibrium Point M. The welfare loss from MHIU alone is now the diagonally barred trapezoid MAGF. Notably, when there is more adverse selection, the expected surplus (willingness to pay minus expected cost) of those priced out of the market by MHIU is larger than with no adverse selection. Graphically, this is shown by comparing the trapezoid MAGF in Figure 4 to DWL in Figure 3. The reason for this is that inframarginal consumers have larger risk premia when there is more adverse selection. To see this, recall that the welfare loss of adverse selection is that some risk averse consumers whose cost is below their willingness to pay (i.e., those who have positive risk premia) do not purchase insurance. This also means that inframarginal enrollees have a larger gap between their cost and their willingness to pay, as shown by the vertical distance GA in Figure 4. Therefore, those priced out of the market by MHIU may correspond to more welfare loss than when there is no adverse selection. This means that the welfare costs of MHIU can differ significantly depending on the slopes of the demand and cost curves. Therefore, the welfare importance of MHIU and its importance relative to adverse selection is ultimately an empirical question.

3.6 Model for Quantifying Welfare

We now propose an empirical framework for estimating the welfare loss of MHIU on the unsubsidized population by building on the graphical model from the previous section.[19] Assuming linearity, we define the demand and cost curves as follows:

$$D_H(P) = \alpha + \beta * P \tag{3}$$

$$AC(P) = \gamma + \delta * P + \sigma * \mu \tag{4}$$

$$MC_H(P) = \frac{\alpha * \delta}{\beta} + \gamma + 2 * \delta * P \tag{5}$$

where D_H is the demand for the unsubsidized consumers, P is the premium, AC is the average cost curve, and MC_H is the marginal cost curve, which is derived in Einav *et al.* (2010). δ is the slope of the average cost curve and represents traditional adverse selection, i.e. healthier consumers dropping out of the market when premiums rise. The main difference between this model and Einav *et al.* (2010) is σ and μ, which represent the share of the population that has subsidies and the difference in costs between the subsidized and unsubsidized population, respectively.[20]

[19] Importantly, this framework is only meant to estimate the spillover on the unsubsidized population and not the total welfare impact of the subsidies. There are many other potential benefits that the subsidies have, such as alleviating liquidity constraints, that may outweigh the negative impacts we find (e.g. Ericson and Sydnor, 2018).

[20] One can rewrite average costs as:

The equilibrium price and quantity are when $P = AC(P)$. With MHIU this is $P^{With} = \frac{\gamma+\sigma\mu}{1-\delta}$. Without MHIU, the $\sigma\mu$ term drops out. This is the additional cost due to MHIU directly, which is scaled by the amount of adverse selection (δ) in the linear model. That is, we can see directly how MHIU raises costs and how adverse selection and MHIU interact (results 1-3). Using these equations, we can calculate the areas of interest from Figure 4. The welfare loss due to adverse selection only is the triangle AGE:

$$\begin{aligned}\Delta_{AGE} &= \frac{1}{2}(Q^{Eff} - Q^{No}) \cdot (P^{No} - MC(P^{No})) \\ &= \frac{-\delta^2}{2(1-2\delta)\beta}\left(\alpha + \frac{\beta\gamma}{1-\delta}\right)^2\end{aligned} \tag{6}$$

The welfare loss of AS and MHIU is the triangle MFE.

$$\begin{aligned}\Delta_{MFE} &= \frac{1}{2}(Q^{Eff} - Q^{With}) \cdot (P^{With} - MC(P^{With})) \\ &= \frac{-\delta^2}{2(1-2\delta)\beta}\left(\alpha + \frac{\beta\gamma}{1-\delta} - \frac{(1-2\delta)\beta\sigma\mu}{\delta(1-\delta)}\right)^2\end{aligned} \tag{7}$$

And finally, the welfare loss due to MHIU is trapezoid MFGA:

$$\begin{aligned}\Delta_{MFGA} &= \Delta_{MFE} - \Delta_{AGE} \\ &= \frac{\delta\sigma\mu}{(1-\delta)}\left(\alpha + \frac{\beta\gamma}{1-\delta} - \frac{(1-2\delta)\beta\sigma\mu}{2\delta(1-\delta)}\right)\end{aligned} \tag{8}$$

4 Empirics

The remainder of the paper is devoted to estimating the magnitude of MHIU. We briefly summarize the methodology for estimating these parameters.

Figure 5 shows a graphical interpretation of our empirical model to estimate the model parameters. The potential unraveling of the system is clear by the cyclic nature of points (b), (c), (d), and (e): subsidized consumers raise average costs, which raises premiums, which lowers demand from unsubsidized consumers, which mechanically increases the percentage of subsidized consumers in the market. The cycle then repeats until either an equilibrium is reached or the market fully unravels and there are no unsubsidized consumers left.

$$\bar{c}(p) = ac_H(O_H, p) + (ac_L(O_L) - ac_H(O_H, p))\left[\frac{D_L(p_L)\cdot M_L}{D_L(p_L)\cdot M_L + D_H(p)\cdot M_H}\right]$$

We define $\sigma(p)$ as the term in brackets and $\mu(p) = ac_L(O_L) - ac_H(O_H, p)$. Then $ac_H(O_H, p) = \gamma + \delta * p$, which is consistent with the model in Einav *et al.* (2010) if we are just measuring the unsubsidized population. To reduce notation, we omit the functional notation of $\sigma(p)$ and $\mu(p)$.

Figure 5: Graphical Representation of Marginal Effects

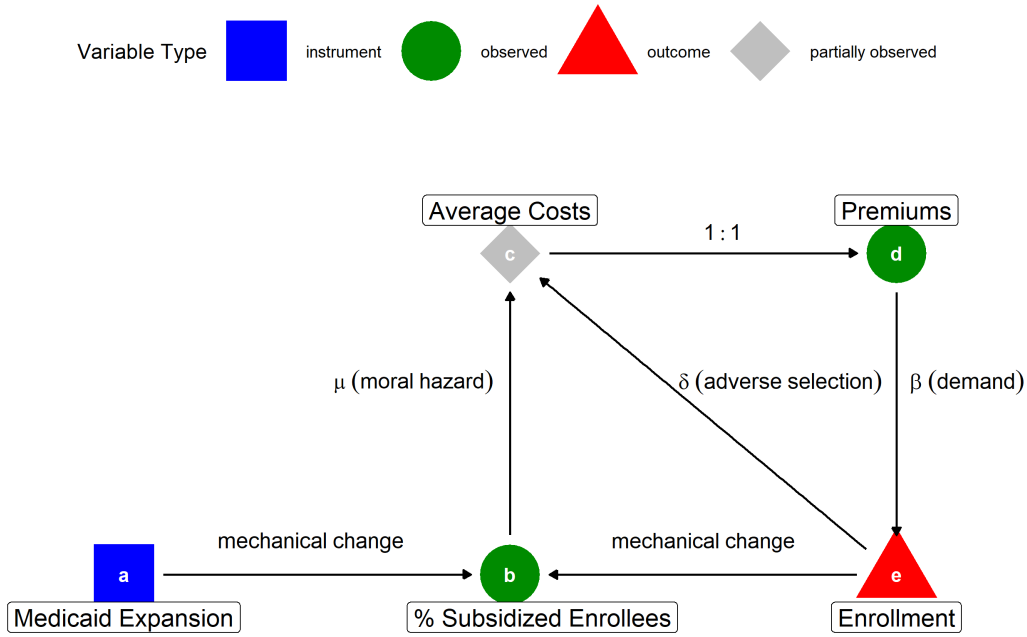

Note: This figure shows a graphical interpretation of the marginal effects of our empirical model. This figure is similar in style to a directed acyclic graph, but is cyclic. Point (c) is only observed in the dataset used as a mechanism check.

For our main estimation, we assume estimates of μ and δ from the literature. We use estimates of the elasticity of healthcare from the RAND Health Insurance Experiment, Lavetti *et al.* (2023), Ellis *et al.* (2017) and Brot-Goldberg *et al.* (2017) to derive μ. We also develop our own estimate of μ in Section 4.3 using the sharp discontinuity in CSR subsidies at 150% FPL on the ACA exchanges. While our estimates are in line with the literature, they are noisy and on the higher end, so we conservatively use the RAND Health Insurance Experiment elasticity estimate as our base case throughout the paper. We use the estimates in Einav *et al.* (2010) to determine δ, which tells us the slope of the average cost curve, or (e) to (c) in Figure 5. We test this assumption by assuming $\delta = 0$ (no adverse selection), δ is half as big, and twice as big as Einav *et al.* (2010).

We use the timing of Medicaid expansion as plausibly exogenous variation in the number of subsidized enrollees. Using data before and after Medicaid expansion, we can measure the mechanical change in the share of subsidized enrollees, which represents going from (a) to (b) in Figure 5, and identifies σ. Since Medicaid expansion reduces the share of high cost, subsidized enrollees, we can use Medicaid expansion as a plausibly exogenous cost shock to identify demand parameters, α and β. Finally, using these parameters, we back out γ.

4.1 Demand and Welfare Estimation

When a state expands Medicaid, a large number of highly subsidized consumers are mechanically removed from the HIXs. If these consumers are significantly more expensive than unsubsidized consumers then, due to community rating, removing them from the exchanges may lower average costs and premiums for unsubsidized consumers (i.e., incomes above 400% FPL). Because Medicaid expansion should only impact those with incomes below 138% FPL, we can determine the impacts of moral hazard as a spillover on the unsubsidized population.

In this section, we estimate a difference-in-differences model using the timing of Medicaid expansion as our treatment variable. These estimates are both used to estimate demand for our structural model, while also providing reduced form evidence of the role of MHIU. We caveat that while we think moral hazard plays a large role in the reduced form effects, there may be other reasons low-income enrollees are more costly than high income consumers (such as the health-wealth gradient). One advantage of the structural model is that it allows us to isolate the moral hazard mechanism from other factors (which we include in various specifications of the structural model).

In our difference-in-differences estimates the treatment group consists of states that expanded Medicaid after 2014 (so they have a pre-expansion period). The control group is states that never expanded Medicaid. We provide evidence that Medicaid expansion reduced premiums and increased enrollment for the unsubsidized population. Notably, in Online Appendix B, we show that these results are unique to the 400+% FPL population, who, because of premium subsidies, are the only consumers who face marginal premium changes.

4.1.1 Data:

HIX Enrollment Data (HIX OEP): We use the HIX Open Enrollment Period Public Use Files (HIX OEP) to measure enrollment on the HIX exchanges. The HIX OEP is either compiled by CMS using HealthCare.gov data or reported to CMS by state-based exchanges. The unit of observation in the raw data is rating area and year. We observe the number of enrollees in each rating area across different disaggregations, including age bins, metal levels, and, our focus, income. Our main sample consists of the 17 states that are on the exchanges, did not expand Medicaid in 2014 or 2015, and reported to the HIX OEP.[21] These 17 states account for 249 rating areas, which gives use 747 rating area-year observations.

American Community Survey (ACS): We use the American Community Survey (ACS) from 2012-2018 to measure insurance-related outcomes and rating area population sizes by income group.

[21] The HIX OEP does not include rating area level data for 2014 therefore 2015 expanders must be excluded. For a detailed description of the data, please refer to Online Appendix B.

The unit of observation is a person, and our sample consists of individuals who are below age 65, have incomes of at least 400% FPL, and either purchase insurance on the exchanges or go uninsured. We additionally draw demographic controls data from the ACS. The ACS data allows us to include 2015 (and in some cases 2014) expanders and therefore we have 20 (28) states that account for 275 (336) rating areas.[22]

Health Insurance Exchange Compare (HIX Compare): Data on premiums are from the HIX Compare dataset from the Robert Wood Johnson Foundation. The HIX Compare data contain premiums for consumers age 27 at the rating area, plan, carrier, and metal level for plans in every state and DC in plan years 2014 through 2017 and for the states who used the federal marketplace in 2014. Following the literature, we aggregate to rating-area level and use the second-lowest-price silver plan as our premium measure.

Table 1 presents summary statistics for all rating areas and separately for rating areas that expanded Medicaid both pre and post-expansion. Panel A shows the share of the HIX enrollment in each income group. In all our sample rating areas, 40% of those enrolled on the HIX had incomes in the 100-150% FPL range. A considerable amount of the risk pool is very expensive to insure. Comparing the late expanding states before and after Medicaid expansion, we see that the share of the enrollment that was 100% to 150% FPL drops by half, from 37% to 19%. In addition, those above 400% FPL (or below 100% FPL) who are not eligible for premium subsidies are a small share of enrollment on the exchanges, at about 9%.

[22]For a full discussion of the years available for each variable of analysis, please see Online Appendix B.

Table 1: Summary Statistics: Enrollment and Premiums

Panel A: HIX Enrollment (HIX OEP Data)						
Variable	Full Sample:		Expand - Pre		Expand - Post:	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Share of HIX 100-150% FPL	0.40	0.12	0.37	0.089	0.19	0.056
Share of HIX 150-200% FPL	0.23	0.034	0.22	0.022	0.26	0.014
Share of HIX 200-250% FPL	0.13	0.032	0.13	0.022	0.16	0.012
Share of HIX 250-400% FPL	0.15	0.053	0.18	0.045	0.24	0.029
N	747		37		50	
Panel B: Insurance Status (ACS Data)						
Variable	Full Sample:		Expand - Pre		Expand - Post:	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Prob. of 400+% FPL ACS Pop Uninsured	0.036	0.19	0.038	0.19	0.028	0.17
Prob. of 400+% FPL ACS Pop Employer	0.87	0.33	0.89	0.31	0.89	0.31
N	3,595,818		663,235		1,190,125	
Prob. of 400+% FPL ACS Pop Direct	0.77	0.42	0.74	0.44	0.80	0.40
N	213,208		20,405		33,870	
Panel C: Silver Plan Premiums (HIX Compare Data)						
Variable	Full Sample:		Expand - Pre		Expand - Post:	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Silver Plan Premiums	3,398.96	635.42	3,301.50	592.99	3,260.58	641.30
N	59,013		4,320		9,183	

Note: Panel A uses HIX OEP data. The unit of observation is a rating area-year. The years of analysis are 2015-2017 and the expansion states are those that expanded in 2016 or 2017. Panel B uses data from the ACS and is weighted using the ACS person weights and the unit of observation is individual-year. For panel B, different years/Medicaid expansion states are utilized based on data availability. The uninsured rate and employer-sponsored rate utilize the years 2012-2018 and the expansion states are those that expanded between 2014-2017. This is because there is data prior to 2014 for these variables. Given that the definition of direct purchase differs prior to 2014, the years used are 2014-2017 and expansion states are those states that expanded 2015-2017. The potential market for the direct purchase population is individuals 400+% FPL that do not have employer sponsored insurance, and who purchase directly with the insurance company or are uninsured. Panel C uses premium data from HIX Compare. The unit of observation is a rating area, year and plan-carrier. The years of analysis are 2014-2017 and the expansion states are those that expanded between 2015-2017.

Panel B gives other insurance metrics constructed in the ACS data. Comparing before and after expansion, the uninsured rate falls from 3.8% to 2.8% – a 26% decline. While we see big shifts between pre and post expansion in expanding states, most of the measures are similar for the all states sample and in the late expanding states prior Medicaid expansion, suggesting our expanding states are observably similar to non-expansion states prior to expansion.

4.1.2 Methods:

We exploit the differential timing of Medicaid expansion as plausibly exogenous variation in the risk pool on the HIXs in a difference-in-differences design. If there is a violation of the stable treatment effect assumption with differential treatment timing, the traditional two-way fixed effects style difference-in-differences will be biased due to identification from the comparison of switching units with already treated units.[23] Recent developments in the literature have established methods to address this limitation.[24] A common theme for these new methodologies is to exclude the problematic comparison of switching units with already treated units in identifying variation. We use the imputed difference-in-differences approach of Borusyak *et al.* (2021), which has the benefit of allowing for fixed effects that are nested within the treatment group level. This method consists of four steps. First, fit a model of the dependent variable on the untreated unit-years.[25] Next, use the fitted model from step 1 to predict the counterfactual (potential not-treated) outcome for the treated unit-years ($\widehat{y^0_{st}}$). Then, take the difference between the counterfactual predicted outcome acquired in step 2 and the *actual* outcome. This yields the unit-time specific treatment effect. Finally, regress the unit-time specific treatment effects from step 3 on the treatment variable. We conduct this analysis of the effect of Medicaid expansion on: (1) the share of total HIX enrollment with incomes below 150% FPL, (2) HIX premiums, (3) HIX enrollment by individuals, and (4) the uninsured rate for individuals with incomes above 400% FPL. Formally, we estimate:

$$y^0_{st} = x'_{st}\alpha + \theta_s + \tau_t + \epsilon_{st} \tag{9}$$

$$y^1_{st} - \widehat{y^0_{st}} = \beta_1 Medicaid\_Expansion_{st} + \nu_{st} \tag{10}$$

where equation (9) is estimated using only the non-treated units, and equation (10) is estimated using all of the units. x_{st} is a vector of state-level (or plan level in the case of premiums) controls, and θ_s and τ_t are state (or rating-area in the case of dependent variables (1) and (2)) and year fixed effects. $Medicaid\_Expansion_{st}$ is a binary treatment variable for when state s implements Medicaid expansion in year t. Last, ϵ_{st} and ν_{st} are the mean-zero error terms.

While effects on premiums and enrollment would provide *premia facie* evidence of welfare loss, as discussed above, measuring the magnitude of welfare loss requires knowing the slope of the demand curve. To estimate demand, we use Medicaid expansion as an instrument for the effect of premiums on enrollment in the 400% FPL population. Results from the difference-in-differences on premiums provide a first-stage relevance test.

We construct the IV estimate of demand through a Wald estimator using the coefficients from equation (10). Specifically, we divide our difference-in-differences estimate when enrollment is the

[23]See Goodman-Bacon (2021) for discussions of the drawbacks of two-way fixed effects with differential timing.

[24]For example, Callaway and Sant'Anna (2020), Athey *et al.* (2021), Borusyak *et al.* (2021), Gardner (2021), and Jakiela (2021).

[25]Because we use two-way fixed effects as our imputation model, we are also using Gardner (2021), which is a special case of Borusyak *et al.* (2021).

outcome by the estimate when premiums are the outcome.[26] Because our treatment is applied at the state level, standard errors for all models, including the IV, are clustered by state (Abadie *et al.*, 2023).

We also check for the existence of pre-trends. Because the HIX did not exist prior to the ACA, there is only one pre-period for enrollment on the HIX, premiums, and the direct purchase measure (where the ACA regulations were applied consistently). Therefore, we focus our pre-trends check on the uninsured rate for which we observe more pre-expansion years. Figure B.1 in Appendix section B.3 presents the results and shows that we find no evidence of pre-trends.

4.1.3 Results:

Table 2: Difference-in-Differences Estimation Results

	(1)	(2)	(3)	(4)	(5)
	Log Share of HIX 100-150% FPL	Premiums (000s):	P(HIX Purchase) (% point):	Unsubsidized Uninsured Rate (% point):	Demand IV
Estimated ATT	-0.472^{***}	-0.371^{***}	2.916^{***}	-0.126^{*}	-0.079^{***}
90% Confidence Interval	[-0.511, -0.413]	[-0.505, -0.212]	[1.25, 4.336]	[-0.281, -0.02]	[-0.128, -0.05]
Implied Intercept	-	-	-	-	0.983^{***}
90% Confidence Interval					[0.888, 1.145]
Person Controls?	No	No	Yes	Yes	-
Plan Controls?	No	Yes	No	No	-
Year Fixed Effects?	Yes	Yes	Yes	Yes	-
Rating-Area Fixed Effects?	Yes	Yes	No	No	-
State Fixed Effects?	No	No	Yes	Yes	-
Pre-Expansion Sample Mean	0.415	3.09	74	3.6	-
Implied Post-Expansion Mean	0.259	2.72	76.9	3.5	-
Implied No-MHIU Mean	0	2.1	81.8	3.3	-
Observations	747	59,013	213,208	3,595,818	-

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table shows the results of the difference-in-differences estimation as well as corresponding sample means for (1) the share of total HIX enrollment with incomes below 150% FPL, (2) premiums on the HIXs, (3) The probability of direct purchasing insurance by individuals 400+% FPL without employer sponsored insurance, and (4) the uninsured rate for individuals with incomes above 400% FPL. Column (5) shows the results of an IV using columns (2) and (3) (divided by 100) to estimate demand (i.e., β and α). Standard errors for all models are block bootstrapped at the state level. Sample means in column (1) are not logged. Column (3) considers the 400+% FPL population (i.e., unsubsidized population) who are either uninsured or purchase on the exchanges (labeled as direct purchase in the ACS). Column (4) is the entire 400+% population. Column (4) is the only column to utilize states that expanded Medicaid in 2014 (as there is pre-data) and uses ACS data from 2012-2018. Column (1) uses data from 2015-2017. Column (2) and (3) use data from 2014-2017. The implied post-expansion and No-MHIU means are linear extrapolations of the regression coefficients on the sample means.

[26] If one were to run a two-stage least squares regression of premiums on enrollment, the corresponding Wald estimator would be:

$$\beta = \frac{E[Enr^1_{st} - Enr^0_{st}|M=1] - E[Enr^1_{st} - Enr^0_{st}|M=0]}{E[Prem^1_{st} - Prem^0_{st}|M=1] - E[Prem^1_{st} - Prem^0_{st}|M=0]} \tag{11}$$

where Enr refers to enrollment measures and $Prem$ refers to premium measures and M an indicator for Medicaid expansion. Others, such as Duflo (2001), have used this instrumented difference-in-differences design. The advantage of this methodology is we can better account for differential timing in treatment. We have also run the standard IV using two-stage least squares and two-way fixed effects and the results are nearly identical.

Table 2 presents the results for the impact of Medicaid expansion on the share of total HIX enrollment with incomes below 150% FPL, premiums, HIX enrollment by individuals, the uninsured rate for the unsubsidized population, and the corresponding demand estimates. Column (1) of Table 2 shows how the share HIX enrollment that is in the 100-150% FPL range changes as Medicaid expands. This is mechanical as Medicaid expansion removes those between 100-138% FPL from the exchanges. Indeed, we find that Medicaid expansion reduces the share of the 100-150% FPL population by 47%, from a baseline of roughly 40%.[27]

Column (2) of Table 2 shows that Medicaid expansion (and thus the reduction of the share of HIX enrollment by those most expensive to insure) is associated with an average premium decline of $374, a 12% reduction relative to the pre-expansion mean. This is an economically meaningful change in premiums and represents the first-stage in our IV below. Our results are within the range found in this literature. Our findings are smaller than the mainline findings in Peng (2017) who uses two-way fixed effects difference-in-differences and synthetic control methods, and larger than the findings in Sen and DeLeire (2018) who uses a state boundary discontinuity approach.

Columns (3) and (4) indicate that following Medicaid expansion total direct purchasing among those without employer sponsored insurance and above 400% FPL increased by 2.9 percentage points while the uninsured rate for the 400+% FPL population *decreased* by 0.126 percentage points (3.5 percent) even though this population is not eligible for Medicaid. In 2014, there were 3 million uninsured individuals above 400+% FPL in all states. A 4% reduction in the uninsured rate would lead to approximately 120,000 fewer individuals above 400+% FPL being uninsured.[28] Furthermore, Medicaid only reduced the share of highly subsidized enrollees by half, so these reduced form estimates may not capture the full impact of MHIU. The row which reports the "Implied No MHIU mean" shows the linear extrapolation of the coefficient estimate to there being no highly-subsidized enrollees. The estimate on the share purchasing, 81.8, suggests that direct purchasing would rise by 7.8 percentage points among the unsubsidized population if there was no MHIU (81.8-74).

Columns (3) and (4) suggest an economically meaningful increase in insurance among the 400% FPL population after Medicaid expansion. However, as discussed, the magnitude of the welfare impact of this is not easily quantified without understanding the shape of the demand and cost curves. Column (5) combines the estimates of Columns (2) and (3) to estimate the demand curve, which we use in our welfare estimates, and presents our estimates for β and α. β is derived from a

[27]Given that in Section 4.3 we find these consumers below 150% FPL are roughly $1,700 more expensive to insure than those directly above it, this represents an economically meaningful change in the risk pool and cost savings for insurers.

[28]While more welfare relevant, one concern is that the uninsured rate may be capturing changes in other markets, besides the on- or off-exchange markets where premiums are changing. In Appendix Section B.4, we reassuringly find no impact on the employer-sponsored market.

Wald estimator from Equation 11.[29] α is then imputed using the pre-expansion sample means and the slope β. Our estimate for β implies that for every \$1000 increase in annual premiums, demand is reduced by 7.8 percentage points. This estimation implies an insurance elasticity of -0.33 for this population. This measure is consistent with other estimates of premium elasticities.[30]

Our primary identifying assumption for the demand estimation IV is that the *only* direct impact of Medicaid expansion on the unsubsidized population is through the premium channel. To check this assumption, we re-run our enrollment analyses using the population between 150% FPL and 400% FPL, where enrollees receive premium subsidies – so they are not subject to marginal changes in premiums – and are also not directly impacted by Medicaid expansion. If Medicaid expansion were impacting enrollment through a channel other than premiums, one would expect it to impact enrollment for those between 150% FPL and 400% FPL as well. Second, to check if there is something about Medicaid expansion that contemporaneously impacted desire for insurance among the 400+% population, we check the employer-sponsored insurance market for the 400+% FPL population. Results are in Online Appendix Section B.4. We find no effect in either placebo group.

In Online Appendix Section B.4, we estimate specification curves to test the robustness of our results. Specification curves graph the coefficient estimates and confidence intervals for numerous specifications, and allow the researcher to test robustness to more choices than they would be able to in a table. We test the 88-112 different specifications (depending on the regression), which vary by the sample of states we include and the controls we include. Nearly all estimates are statistically significant and the results are quite stable across specifications.[31]

4.2 Welfare Calculations

Now that we have estimated our demand curve, we can estimate the welfare loss from MHIU. Recall that our demand and cost system are defined as:

$$D_H(P) = \alpha + \beta * P$$
$$AC(P) = \gamma + \delta * P + \sigma * \mu$$

Our estimates of β and α, the slope and intercept of the demand curve, respectively, are $-.000078$ and 0.98. These are the coefficients in Column (5) of in Table 2. σ is 0.415, which is the pre-

[29] Column (5) is derived by taking the estimate of Column (3) (divided by 100 to change back from percentages) divided by Column (2).

[30] For example, Cutler and Reber (1998) finds an extensive margin insurance elasticity of 1, and 2 for across plan elasticity. Chan and Gruber (2010) and Royalty and Solomon (1999) find elasticities below 1.

[31] We exclude 2018 for the direct purchase results as the Trump administration made multiple changes to policies which impact this market (and may impact states differentially). In Appendix Section B.4, we examine the effect of including 2018. We find no meaningful difference when 2018 is included.

expansion mean share of the HIX who are 100-150% FPL (Table 2). The difference in costs for the subsidized consumers, μ, is derived from elasticity estimates in the literature. We also provide our own estimate of μ in Section 4.3, where our results closely align with the elasticities found in Ellis *et al.* (2017) and Brot-Goldberg *et al.* (2017). However, for our base case, we use the more conservative RAND Health Insurance Experiment elasticity of -0.2, which implies that subsidized individuals are $721 more costly than unsubsidized individuals in our setting.[32] We follow Einav *et al.* (2010) and set $\delta = 0.155$. Finally, assuming that $P = AC$, we back out γ directly from the average cost curve as all other terms in the AC curve are now known.[33]

Given these parameters, we can construct counterfactual equilibrium and welfare estimates using equations (6) through (8). Table 3 shows the welfare impacts of MHIU under various sets of parameters. We start with the base case using the estimates from the previous sections. We find that overall enrollment declines by 20 percentage points due to AS and MHIU relative to efficient pricing. Given that 74% of this population purchases insurance (column (3) of Table 2), this would imply that nearly everyone would be insured in the efficient scenario, which is consistent with all consumers being risk averse and is the base case in Einav and Finkelstein (2011). Of the 20 percentage points of uninsurance, 3 percentage points is due to MHIU. This suggests that adverse selection leads to about 7 times more of the total uninsurance than MHIU, but MHIU is still large in magnitude.

These reductions in enrollment correspond to a welfare loss of $175 for the average unsubsidized individual. Of that, $130 is due to AS, and $46 is due to MHIU. AS causes a much larger decline in enrollment in our model, but the welfare impacts are of comparable magnitudes (MHIU accounts for 26 percent of the welfare loss and 14 percent of the enrollment loss). This is because the consumer surplus for those priced out by MHIU, but not AS, is higher than for those priced out by AS (as shown in Figure 4).

One advantage of using a structural model is we can also isolate the role of moral hazard from the health-wealth gradient. While the RAND elasticity and the size of the OOP price reductions imply a large cost increase for low income consumers, if low income enrollees are sicker, on average, then they may be more costly for that reason as well. Our difference-in-differences design would pick up both effects. However, our structural model uses only the demand elasticity (which does not require us to take a stand on the mechanism) and assumptions about differences in average costs, where we can isolate these mechanisms. Using only the RAND elasticity, the implied premium difference

[32] As defined in Footnote 20, technically σ and μ are functions of market premiums. In our counterfactuals, we always go from the observed outcome to an outcome with no subsidized consumers ($\sigma = 0$) so the curvature of these functions does not impact our results.

[33] Our goal with these counterfactuals is to test how sensitive our estimates of the importance of MHIU are, as we change parameter estimates. Therefore, we hold fixed the observed equilibrium for each counterfactual, but calculate the No MHIU equilibrium and efficient outcome using different assumptions for the parameter estimates. Likewise, we allow γ to change as other parameters change, to ensure our model is consistent with the observed equilibrium.

Table 3: Welfare Estimates

		(1)	(2)	(3)	(4)	(5)	(6)
		Enrollment loss due to MHIU	Enrollment loss due to MHIU and AS	Welfare loss (Δ_{MFE})	Loss due to AS (Δ_{AGE})	Loss due to MHIU (Δ_{MFGA})	Share of Welfare Loss due to MHIU
1	RAND Elasticity (Base Case)	0.028	0.2	175	130	46	0.26
2	Lavetti et al Elasticity	0.018	0.188	155	126	28	0.184
3	Ellis et al Elasticity	0.062	0.241	254	141	113	0.445
4	Brot-Goldberg et al Elasticity	0.08	0.244	281	132	149	0.53
5	No Adverse Selection	0.024	0.024	4	0	4	1
6	1/2 as Much Adverse Selection	0.026	0.096	49	26	23	0.464
7	2x More Adverse Selection	0.034	0.244	390	318	72	0.185
8	If markups are 15%	0.034	0.244	270	201	69	0.255
9	Health/Wealth Gradient	0.028	0.244	264	207	57	0.215

Note: This table shows welfare estimates from our theoretical model and various changes to parameters. AS is adverse selection, MHIU is moral hazard induced unraveling. All numbers are positive in the sense that losses are positive. For example, MHIU lead to 0.028 percentage points lower enrollment in our base case. Rows 2-4 change the out-of-pocket cost elasticity parameter, μ, to those reported in Lavetti *et al.* (2023), Ellis *et al.* (2017), and Brot-Goldberg *et al.* (2017). These elasticities are -0.13, -0.44, and -0.57, respectively. Rows 5-7 use different slopes of the average cost curve, δ. In our base case we use $\delta = .155$ from Einav *et al.* (2010). These rows use 0, 0.0755, and 0.31, respectively. Row 8 allows for markups such that the equilibrium condition is $.85P = AC$. The efficient outcome is left unchanged, so the markup is accounted from in the welfare loss and enrollment loss to adverse selection. Row 9 allows for low income consumers to be sicker by adding $1,000 in cost to low income consumers. Some parameter values cause the estimated efficient outcome to go negative, we bound premiums at 0 and quantities at 1, which is consistent with some consumers using no health care.

from the equilibrium with and without MHIU is $349. While Medicaid expansion did not remove all low income individuals, this suggests that the RAND elasticities (and classical adverse selection) can explain much of the difference in premiums that we see in the difference-in-differences design.

To see how different mechanisms impact our results, we vary the amount of moral hazard, adverse selection, and allow for market power. We first explore how our estimates are impacted by moral hazard. Rather than the RAND estimate, we use the out-of-pocket cost elasticities from Lavetti *et al.* (2023), Ellis *et al.* (2017), and Brot-Goldberg *et al.* (2017), which are -0.13, -0.44, and -0.57, respectively.

As shown in rows 2 through 4, the welfare loss due to MHIU falls considerably when there is less moral hazard. Using the numbers from Lavetti *et al.* (2023), the welfare loss due to MHIU falls to 28, which is still 18% of the total welfare loss. Using Brot-Goldberg *et al.* (2017) elasticity, the welfare loss of adverse selection itself is a little larger (because γ is decreasing in the elasticity), but the loss due to MHIU is much higher.[34] This is not surprising as more elastic subsidized consumers means that the externality on unsubsidized consumers is larger. Using this range of estimates, we

[34]The estimates in Rows 3 and 4 of Table 3 align with the elasticity that we will find in section 4.3. These rows find enrollment losses due to MHIU between 6 and 8 percentage points. This is consistent with our finding of a 3 percentage point increase in direct purchasing in our reduced form analysis, where the share of the 100-150% FPL population on the HIX only declined by about half. Indeed, the implied No MHIU mean in Table 2, which is a linear extrapolation of the reduced form estimate, and suggests that enrollment would drop by 7.8 percentage points (81.8-74). This suggests that adverse selection leads to about 2.7 times more of the total uninsurance than MHIU, but MHIU is still large in magnitude.

see the enrollment losses due to MHIU vary from 2 percentage points to 8 percentage points, and the share of welfare loss due to MHIU varies from 18 percent to 53 percent. While this is a wide range, even at the lower end these results suggest economically meaningful changes in enrollment and welfare loss due to MHIU.

Rows 5 through 7 explore what would happen if there was no adverse selection ($\delta = 0$), if we cut the slope of the average cost curve in half ($\delta = 0.0775$), and if we doubled the slope of the average cost curve ($\delta = 0.31$). With no adverse selection, the MHIU now accounts for 100% of the welfare loss. However, the estimated welfare loss due to MHIU is much smaller, just $4 per person. As discussed in Section 3.5, the smaller welfare loss is because marginal consumers account for less surplus as their risk premium is smaller. When $\delta = 0.31$, the welfare loss of MHIU is 18 times as large as when $\delta = 0$. These results show that even with relatively similar amounts of enrollment differences due to MHIU, the welfare impact of MHIU can differ tremendously depending on how much adverse selection is present.

Row 8 explores loosening the assumption of a zero profit equilibrium. To account for Medical Loss Ratio regulation, which caps insurer profits at 15 percent, we instead assume that profit is such that .85 P = AC in each equilibrium, though we do not change the efficient outcome. In the table, the welfare loss due to the markup is included in the loss due to adverse selection (column 4), because the efficient outcome has no markup. In this case, the welfare loss is slightly higher and increases from 175 in the base case to 270 because of the markup. Welfare losses for MHIU increase to 69, from 46. Even with allocating the welfare loss of markups to adverse selection, we see that in relative terms markups have large welfare impacts when interacted with MHIU. This is because markups price out more consumers leaving the marginal (to MHIU) consumer as having more per-person surplus.

Row 9 explores the possibility that lower income enrollees may be more expensive because they are less healthy, on average. To examine how this would impact our results, we compute a counterfactual where we allow for the average cost of low income consumers to be scaled up by $1,000 in addition to allowing for $721 in moral hazard.[35] That is, absent moral hazard, we assume that low income consumers are $1,000 more expensive than the average high income consumer due to the health/wealth gradient and $721 more expensive due to moral hazard, given the RAND elasticity.[36] In this counterfactual, the welfare loss of adverse selection and the health-wealth gradient is $207, which is larger than the baseline result of $130 because the $1,000 spillover from the 40% of low

[35] We chose $1,000 because this implies a price difference that is consistent with our difference-in-differences estimates. In practice, one could rescale the $721 to include moral hazard on top of the $1,000 price difference. We chose not to do this to keep our estimates comparable across columns, therefore our estimates are slightly smaller than what would happen if we also inflated the moral hazard cost proportionally.

[36] In row 9, the efficient allocation does not include either the moral hazard or $1,000 cost, the equilibrium without MHIU includes only the $1,000 cost, and the equilibrium with MHIU includes both the $721 and $1,000 cost.

income consumers in the market is not accounted for in the baseline. The enrollment loss due to MHIU is similar in both cases, but the welfare loss due to MHIU is larger with the health/wealth gradient at $57 than the baseline of $46. This is because these inframarginal consumers have larger risk premia when other mechanisms force more consumers out of the market. While both sources of welfare loss grow, the share of welfare loss due to MHIU remains above 20%. That is, even if one assumes that these consumers are higher cost because of both the health-wealth gradient and moral hazard – with the health-wealth gradient being larger in magnitude – we still find that the moral hazard impacts implied by the RAND elasticity are economically meaningful.

To summarize, we see that enrollment and welfare losses due to MHIU are large and vary a lot depending on model parameters and, importantly, on the amount of adverse selection. While adverse selection is one of the most studied topics in health economics, the importance of MHIU relative to AS in the unsubsidized market is not trivial. Our findings suggest that MHIU accounts for at least 18% of the total welfare loss in all the cases we examine and can range much higher.

4.3 Mechanisms Check

In this section, we provide evidence that the CSR subsidies were leading to moral hazard on the ACA exchanges. There is a large literature finding evidence of moral hazard in many contexts, however, to our knowledge, Lavetti *et al.* (2023) is the only other paper which explores the moral hazard due to the CSR subsidies. We view this section as being complementary to their paper – we use different data, a different geographic setting, and different identification strategies, yet find qualitatively similar results. We briefly discuss our methods, results, and robustness checks here.

To estimate the level of moral hazard, we exploit a natural experiment to examine how changes in out-of-pocket prices impact healthcare utilization and expenditures. Individuals with incomes between 100-150% FPL are eligible for CSR subsidies amounting to a 94% actuarial value plan, which means these individuals pay only $6 for each $100 of healthcare expenditures. Individuals in the 150-200% FPL range receive subsidies amounting to an 87% actuarial value and thus pay $13 for each $100 of healthcare expenditures. Someone at 152% of the FPL is responsible for roughly double what an enrollee at 148% of the FPL would be expected to cover out-of-pocket. We examine how this plausibly exogenous doubling in out-of-pocket costs impacts healthcare utilization, which we argue is evidence of moral hazard. Because the existence of moral hazard is vital to our model, we also explore potential confounders – adverse selection and the health/wealth gradient – in Section 4.3.3 and find no evidence of either.

Figure 6: Regression Discontinuity Plots

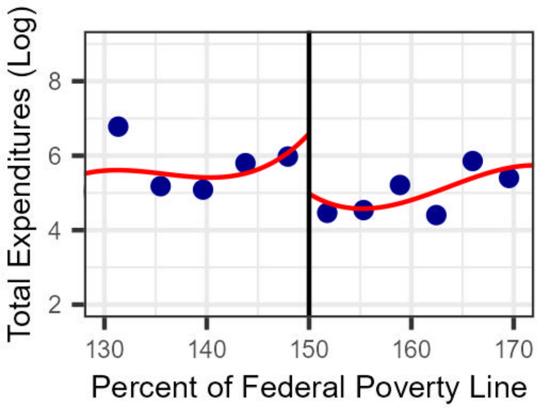

(a) Total Expenditures

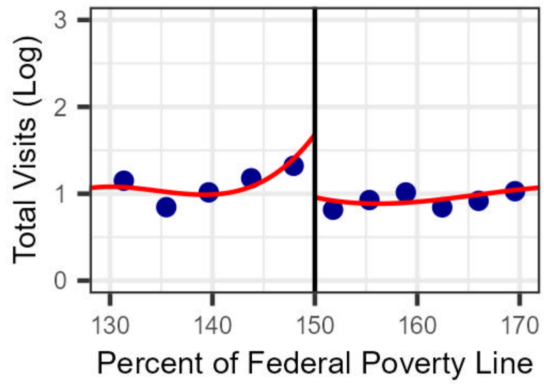

(b) Total Visits

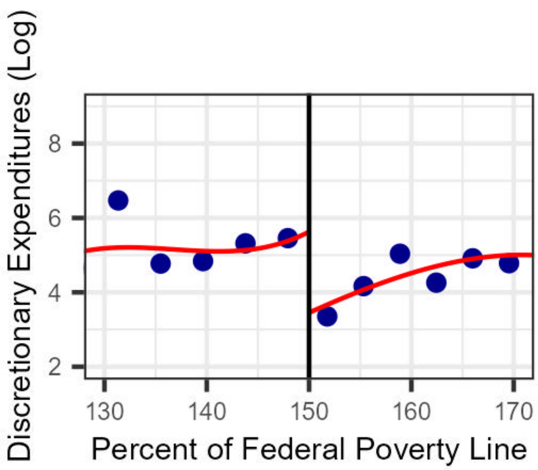

(c) Discretionary Expenditures

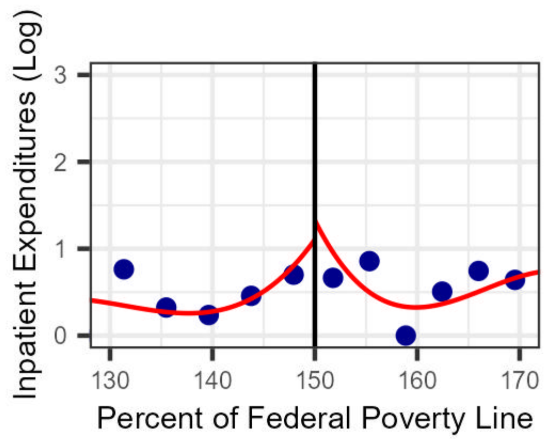

(d) Inpatient Expenditures (placebo)

Note: This figure plots the log for total CSR eligible expenditures and visits for bins of consumers in the health insurance exchanges based on their income relative to the federal poverty line. Consumers below 150 are eligible for CSR subsidies amounting to a 94% actuarial value plan and those above are eligible for an 87% actuarial value plan. Each dot represents a bin of approximately 20 consumers.

4.3.1 Methods and Results for Moral Hazard Exploration:

We use the Household Component (HC) of the Medical Expenditure Panel Survey (MEPS). The HC-MEPS is a two-year panel survey of health expenditures, utilization, health status, demographics, income, and health insurance coverage from individual households in a nationally-representative sample.[37] We use the 2014-2018 HC-MEPS data, which contains 158,909 individual-level observations in the under-65 sample. Our analysis focuses on the 2,604 individuals with HIX coverage and expenditure/health data, which corresponds to 1,709 families.[38] We also use the Johns Hopkins Adjusted Clinical Group (ACG) System to control for health status, which is often used in the setting of premium rates as part of the underwriting process.[39] Because ACG scores are both highly predictive of health expenditures and exogenous to subsidy status, we use them to formally limit the potential impact of adverse selection through an Oster (2019) test.

Figure 6 shows the log (plus one) of total healthcare expenditures and visits, discretionary expenditures, and inpatient expenditures for consumers binned by their FPL for those individuals in the health insurance exchanges.[40] Panel (c) shows a visual discontinuity in the more discretionary categories of health expenditures while there is no discontinuity in our placebo check of inpatient expenditures in Panel (d). When combining discretionary and inpatient expenditures, panels (a) and (b), the visual evidence is more muted. Summary statistics and further discussion of the data can be found in Online Appendix C.1.

We exploit the sharp discontinuity in the CSR subsidies to show lowering the out-of-pocket price leads to more consumption of healthcare. The CSRs are based on the individual's income level, which we exploit using a regression discontinuity design:

$$y_{it} = f(income_{it}, X_{it}) + Over150_{it} + e_{it} \tag{12}$$

Where y_{it} is health expenditures or visits for individual i in year t; $income_{it}$ is income as a percentage of the federal poverty line (FPL); X_{it} is a vector of controls including health status, sex, marital status, race, age, and region and year dummies; $Over150_{it}$ is a binary indicator for if the consumer's income level is greater than 150% of the relevant (for that year) FPL; and $f()$ is a local-polynomial spline with uniform kernel weighting. Because the MEPS income measures are self-reported and not directly equivalent to the modified adjusted gross income used by the ACA, there is potential

[37] Healthcare expenditure and income data are self-reported, though MEPS verifies the household survey responses for healthcare expenditures and visits by contacting the healthcare providers of the respondents. Misreported incomes may lead to measurement error in our data. As long as the misreporting of income is not correlated with health status this misreporting should have an attenuating bias.

[38] Due to our unbalanced panel, 892 individuals with HIX coverage appear in our dataset twice meaning the total number of observations is 3,496.

[39] The ACG software uses a mixture of clinical knowledge and claims experience (i.e., demographics (we use age and gender, the ACG system does not account for income) and diagnosis codes related to doctors visits, hospitalizations, and prescriptions) to predict various outcomes like the probability of being a high user of care, the probability of hospitalization, and predicted total expenditure. We use the predicted total expenditure as a control because the three are highly correlated. For more information about the ACG system please refer to https://www.johnshopkinssolutions.com/solution/acgsystem/.

[40] Because health expenditures are right-skewed, we use log plus 1, which allows us to retain zero-valued observations. Discretionary expenditures are non-inpatient spending excluding healthcare that is not eligible for CSRs (i.e., excluding dental, vision, and other medical equipment). Inpatient expenditures are excluded from discretionary expenditures as they are used as a placebo test.

for measurement error very close to the cut-point. To address this, our preferred specification uses a "donut RD" where observations in the 149-151% FPL range are not included.

Our main identifying assumption is that there is no endogenous sorting into treatment; specifically, consumers are not manipulating their reported income to qualify for additional subsidies. We check for this in Appendix C.5 and find no evidence of sorting. Standard errors for all models are clustered at the family level.

The results for estimating equation (12) for total, inpatient, and discretionary healthcare expenditures and visits are presented in Table 4. The bandwidths were calculated using the Calonico *et al.* (2014) method and the number of observations in the bandwidth are shown. Columns (1) and (3) display the results controlling for gender, marital status, race, age, educational attainment, family size, census region, and year. Columns (2) and (4) display the results controlling for the demographics described above and health status (i.e., total predicted healthcare costs). Panels (A), (B), and (C) present the results for total, inpatient, and discretionary healthcare expenditures and visits, respectively.

Panel (A), Column (1) of Table 4 shows that individuals in the health insurance exchanges who are eligible for the most generous CSR subsidies spend more per year on healthcare than those individuals just above 150% of the FPL. This equates to a demand elasticity of -0.47 and implies increased expenditures of approximately $1,710 per year. This elasticity is larger than that in the RAND health experiment, who examine coinsurance variation, but well within the range found in Brot-Goldberg *et al.* (2017), who examine deductible variation. In our setting, both copays and deductibles are typically changed to achieve higher actuarial value. Our estimate is also larger than the effect found in Lavetti *et al.* (2023), though our confidence intervals overlap. Our estimate is in line with Ellis *et al.* (2017), who examine healthcare demand elasticity by service type and find an overall elasticity of -0.44.

Column (4) indicates that individuals below 150% FPL visit a healthcare professional once more per year than those above 150% FPL. Panel (B) of Table 4 shows that we fail to find a statistically significant effect for more generous CSR subsidies on both inpatient expenditures and visits for HIX enrollees, consistent with the intuition that more severe ailments should be less discretionary. Panel (C) indicates that our results from Panel (A) are driven largely by the more "discretionary" health expenditure types.

4.3.2 Balance Tests, Robustness Checks, and Other Mechanisms:

In Online Appendix Sections C.2, C.3, C.4, and C.5, we perform a number of robustness checks and validation exercises, which we summarize briefly here.

Table 4: Regression Discontinuity Models for Healthcare Expenditures and Utilization

	ln(Expenditures):		ln(Visits):	
	(1)	(2)	(3)	(4)
Total Healthcare (A):				
Over 150% FPL	-0.796^{*}	-0.826^{*}	-0.247^{*}	-0.256^{**}
	(0.481)	(0.433)	(0.149)	(0.128)
Implied Change (μ)	1712.109	1754.253	1.103	1.141
Implied Elasticity	−0.470	−0.482	−0.187	−0.194
Placebo - Inpatient Healthcare (B):				
Over 150% FPL	0.369	0.353	< 0.001	< 0.001
	(0.306)	(0.294)	(0.012)	(0.012)
Implied Change (μ)	−271.885	−258.364	−0.000	0.000
Implied Elasticity	0.382	0.363	0.000	−0.000
Discretionary Healthcare (C):				
Over 150% FPL	-0.984^{**}	-1.013^{**}	−0.209	-0.219^{*}
	(0.481)	(0.433)	(0.143)	(0.122)
Implied Change (μ)	1415.229	1440.041	0.780	0.812
Implied Elasticity	−0.537	−0.546	−0.162	−0.169
Bandwidth	37.998	37.998	37.998	37.998
Number of Observations (LATE)	716	716	716	716
Year FEs?	Yes	Yes	Yes	Yes
Census Region FEs?	Yes	Yes	Yes	Yes
Demographic Controls?	Yes	Yes	Yes	Yes
Health Controls?	No	Yes	No	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table shows the results of the regression discontinuities for the log of total spending and visits on healthcare. Panel (A) shows all healthcare expenditures and visits. As a placebo test, inpatient expenditures and visits, which capture emergency room visits, are also shown in Panel (B). Panel (C) depicts non-inpatient spending and visits that are eligible for CSRs (i.e., excluding dental, vision, other medical equipment, and inpatient expenditures as these are a placebo test). The sample is for individuals younger than 65 with insurance from the ACA marketplaces. The bandwidths were calculated using the Calonico *et al.* (2014) method and the kernel is uniform (as recommended in Lee and Lemieux (2010)). The additional covariates are predicted health expenditures as well as sex, marital status, race, age, educational attainment, family size, census region, and year. Standard errors for all models are clustered at the family level. The average expenditures for those below 150% FPL in our bandwidth is $3,120.19.

In Section C.2, we include a specification curve which shows that our results are robust to different combinations modeling choices including (1) the use and size of the donut; (2) the size of the RD bandwidth; (3) the type of RD kernel; (4) the transformation of the expenditure variable (logs versus inverse hyperbolic sine); (5) the type of control variables; and (6) the type of fixed effects. We try 2,016 different combinations using these modeling choices. Every point estimate is below zero with 80% of these estimates being statistically below zero.

In Section C.3, we also do a more formal examination of the balance of predetermined covariates above and below the threshold. If there are no systematic differences across the treatment eligibility threshold (i.e., CSR eligibility cannot be manipulated) then no discontinuity should exist for predetermined covariates (i.e. age, sex, race, education level, etc.). For all of the covariates, the point estimates are statistically insignificant and, in most cases, quite small indicating that there is no evidence of a discontinuity for these covariates.

In Section C.4, we examine the point estimates of regressions which use different cutpoints. If our results are being driven solely by the treatment, we would expect estimates away from the true thresholds to be insignificant and close to zero. Our results show a distinct U-shape with the largest coefficient estimates near the true threshold.

Finally, in Section C.5, we examine if individuals are able to manipulate the running variable. For our analysis, this would imply that individuals' are manipulating their income by working fewer hours, fewer jobs, etc. If this is occurring, it will violate the continuity assumption of the running variable, which is required for identification in a regression discontinuity design. To determine if this is the case, we run the McCrary (2008) sorting test, which checks for a discontinuity in the density of the data along the running variable and we fail to reject the null hypothesis that there is no sorting into treatment assignment.

4.3.3 Discussion of Alternative Mechanisms

We next argue that our results are due to moral hazard alone, as we find little evidence of two other potential alternative explanations for this result: the health/wealth gradient and adverse selection.

Health/Wealth Gradient: Our empirical strategy of a regression discontinuity, which compares people with similar incomes, is meant to address the concern that lower-income people are generally sicker than higher-income people. In the limit, the average person at 149% FPL should have a similar health status to the average person at 151% FPL. To test this assumption empirically, Table C3 in Appendix Section C.7 presents a falsification test using the employer-sponsored insurance (ESI) market. As with Table 4, we examine discretionary, inpatient, and total expenditures/visits. If there is a discontinuity in population health status across the 150% FPL threshold, we would expect to see a discontinuity in the ESI population as well. The point estimates are all much smaller in magnitude and none of the estimates in Table C3 are statistically significant.

Adverse Selection: One concern in our setting is that enrollment is endogenous: sicker people may be more likely to enroll in insurance when that insurance is more generous. There are three potential margins for adverse selection. First, people across the threshold could differentially sort into more generous plans. This margin is limited because the CSR subsidies are only available for silver plans and, within our bandwidth, subsidized silver plans are more generous than gold plans on either side of the threshold.[41] The second margin is potential "cross-market" selection between

[41] Indeed, Sprung and Anderson (2018) find that 80-90% of those in the 100-200% FPL range select a silver plan.

Medicaid and the HIXs. Recent work by Clemens (2015) and Holmes (2022) has documented this kind of behavior, where sicker consumers are more likely to use Medicaid due to its extremely low cost sharing. We believe this is unlikely to be driving our result. Federal regulations require that those below 138% FPL in states that have expanded Medicaid are no longer eligible for subsidies. Thus, it would require wide-scale falsification of incomes for this to explain our results.

The third margin is consumers opting out of insurance altogether. As detailed in Online Appendix Section C.6, selection along this third margin would actually attenuate our estimates, not inflate them away from zero. In much of the literature on moral hazard and adverse selection (e.g., Chiappori *et al.* (1998)), economists compare the costs of those selecting between more generous or less generous insurance. This leads moral hazard and adverse selection to impact costs in the same direction. In contrast, we compare the costs of those who purchase insurance (rather than going uninsured) across two different levels of plan generosity, where the generosity level is exogenously determined by the consumer's income. With more generous coverage more individuals will select into having insurance than the less generous plan. If there is adverse selection, the marginal (and average) enrollee for the more generous plan will be healthier than the marginal (and average) enrollee for the less generous plan. Therefore, in our setting, selection of consumers opting out of insurance altogether should attenuate our estimates as the average enrollee in the more generous plan should be healthier.

To address adverse selection concerns through the other two channels, we first note that the McCrary (2008) sorting test (in Appendix C.5) is a direct test of selection in our setting, and we fail to reject the null hypothesis of zero selection. As further robustness, we control for measures of observable health status from the ACG algorithm. We expect this measure to be correlated with health status more broadly, and therefore capture differences in health status across the threshold, if present. Columns (2) and (4) in Table 4 control for health status, while Columns (1) and (3) do not. The difference between the point estimates between these columns in Panel A are not economically or statistically significant. That controlling for health does not change the estimates in Table 4 suggests that the out of pocket price elasticities are more important than differences in enrollee health status.

We formalize this intuition through Oster's (2019) test for omitted variable bias. The idea of Oster's (2019) test is to bound the potential impact of selection based on the potential predictive power of the unobserved selection relative to an observed covariate. Table C4 presents the results of this test for CSR eligible expenditures excluding inpatient expenditures. The first row depicts our estimate, standard error (clustered at the family level), t-value, and R^2 values used for the test. The latter rows depict the bounds on our estimate due to various levels of potential selection. We find that if the unobserved selection (that is orthogonal to ACG score) is four times as predictive of health expenditures as ACG scores, our estimate would drop from -0.984 to -0.375 and remain statistically significant.

5 Conclusions

The literature on moral hazard has focused on the over-consumption of health care and the under-consumption of generous insurance (Pauly, 1968; Holmström, 1979; Einav and Finkelstein, 2018). In this paper, we show that when there are community rating regulations, moral hazard leads to another source of welfare loss: an externality on others whose insurance is not subsidized. We show how this externality mimics the effect of adverse selection and can induce a type of Akerlof (1970) spiral that leads to lower enrollment and the potential unraveling of the market.

Two components of the Affordable Care Act, the Health Insurance Exchanges and the Medicaid Expansion, provide a natural experiment to test this theory. These components provide (1) exogenous variation in the number of consumers who have generous insurance, and (2) creates a distinct subgroup of consumers whose moral hazard has spillovers to other consumer groups while remaining unaffected by their own moral hazard. These spillovers lead to economically meaningful reductions in enrollment, and welfare losses that are not trivial relative to the welfare losses due to adverse selection.

Moral hazard induced unraveling does have an obvious policy response – directly reimbursing the insurers for the moral hazard costs. Interestingly, a "reverse" natural experiment has already occurred and seems to show it is possible. In late 2017, the Trump administration halted the existing CSR subsidy reimbursements to insurers. That is, prior to the Trump administration ruling, the government would reimburse insurers for the CSR subsidies provided to low-income consumers. For example, when the CSR subsidies lowered a copay from $25 to $5 the government would reimburse insurers $20. The Trump administration removed this CSR subsidy reimbursement to insurers. The Congressional Budget Office (CBO) estimated that, because premium subsidies would rise, the canceled CSR subsidy reimbursement would actually increase government expenditures by, "$6 billion in 2018, $21 billion in 2020, and $26 billion in 2026," (Congressional Budget Office, 2017). The logic behind the CBO finding depends on policies that instruct insurers on how to respond to the loss of CSR subsidy reimbursements (as demonstrated in Rao and Nowak (2019)).[42]

[42] As the CSR subsidies are only available for silver plans, many insurers raised the price of silver plans only – which is referred to as "silver loading." There was some policy discussion around forcing insurers to spread their higher costs to all metal levels, known as "broad loading." Getting rid of CSR reimbursements (as the Trump administration did) caused premium subsidies to rise in lock-step with the second lowest silver plan premiums. If consumers in 200-400% FPL were buying the second lowest cost silver plan, they would be unaffected. However, if there is "silver loading" then this policy would mostly affect the cost of silver plans. If these consumers were purchasing a bronze or gold plan – whose premiums rose by less than the second lowest cost silver plan – then these consumers would see their net premiums fall even if their gross premiums are rising. This would lead to higher enrollment among the 200-400%

Our findings are even more relevant today than during our sample period, as this is an area where there is recent policy activity. First, there is debate on if insurers are to even be reimbursed for the direct costs of the subsidies, which would exacerbate the already substantial cost differential. Second, the recent American Rescue Plan has increased incentives for states to expand Medicaid and removed the income threshold for receiving premium subsidies, though this is only in effect for two years. Additionally, in the time of COVID-19, 31 million people filed for unemployment (and lost their health insurance). Kaiser estimated that 31% were eligible for subsidies on the exchanges (KFF, 2020). Their enrollment could lead to negative spillovers that price out the few unsubsidized people who remain on the exchanges. While our results are economically meaningful in their own right, several, much larger, insurance markets also feature this kind of "managed competition" that could lead to moral hazard induced unraveling. Quantifying the impact in other insurance markets is an excellent avenue for further research.

FPL population, lower enrollment for those above 400% FPL, and higher government expenditure, which is what the CBO found.

Bibliography

Abadie, A., Athey, S., Imbens, G. W. and Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, **138** (1), 1–35.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, **84** (3), 488–500.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G. and Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, pp. 1–15.

Azevedo, E. M. and Gottlieb, D. (2017). Perfect competition in markets with adverse selection. *Econometrica*, **85** (1), 67–105.

Bellemare, M. F. and Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, **82** (1), 50–61.

Borusyak, K., Jaravel, X. and Spiess, J. (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.

Brot-Goldberg, Z. C., Chandra, A., Handel, B. R. and Kolstad, J. T. (2017). What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics. *The Quarterly Journal of Economics*, **132** (3), 1261–1318.

Buchmueller, T. and DiNardo, J. (2002). Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania, and Connecticut. *The American Economic Review*, **92** (1), 280–294.

Callaway, B. and Sant'Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.

Calonico, S., Cattaneo, M. D. and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, **82** (6), 2295–2326.

Cattaneo, M. D., Idrobo, N. and Titiunik, R. (2019). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.

Centers for Medicare and Medicaid Services (2019). Aptc and csr basics. https://marketplace.cms.gov/technical-assistance-resources/aptc-csr-basics.pdf.

Chan, D. and Gruber, J. (2010). How sensitive are low income families to health plan prices? *American Economic Review*, **100** (2), 292–96.

Chiappori, P.-A., Durand, F. and Geoffard, P.-Y. (1998). Moral hazard and the demand for physician services: First lessons from a french natural experiment. *European Economic Review*, **42** (3-5), 499–511.

— and Salanié, B. (2000). Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy*, **108** (1), 56–78.

Cinelli, C. and Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **82** (1), 39–67.

Clemens, J. (2015). Regulatory Redistribution in the Market for Health Insurance. *American Economic Journal: Applied Economics*, **7** (2), 109–134.

Congressional Budget Office (2017). The effects of terminating payments for cost-sharing reductions.

Cutler, D. M. and Reber, S. J. (1998). Paying for health insurance: The trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, **113** (2), 433–466.

— and Zeckhauser, R. J. (2000). The anatomy of health insurance. In *Handbook of Health Economics*, vol. 1, Elsevier, pp. 563–643.

DeLeire, T., Chappel, A., Finegold, K. and Gee, E. (2017). Do individuals respond to cost-sharing subsidies in their selections of marketplace health insurance plans? *Journal of Health Economics*, **56**, 71–86.

Drake, C. and Anderson, D. M. (2020). Terminating cost-sharing reduction subsidy payments: The impact of marketplace zero-dollar premium plans on enrollment. *Health Affairs*, **39** (1), 41–49.

Duflo, E. (2001). Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *American Economic Review*, **91** (4), 795–813.

Einav, L. and Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, **25** (1), 115–38.

— and — (2018). Moral hazard in health insurance: What we know and how we know it. *Journal of the European Economic Association*, **16** (4), 957–982.

—, — and Cullen, M. R. (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *The Quarterly Journal of Economics*, **125** (3), 877–921.

Ellis, R. P., Martins, B. and Zhu, W. (2017). Health Care Demand Elasticities by Type of Service. *Journal of Health Economics*, **55**, 232–243.

Ericson, K. M. and Sydnor, J. R. (2018). *Liquidity constraints and the value of insurance*. Tech. rep., National Bureau of Economic Research.

Fehr, R., Cox, C., Levitt, L. and Claxton, G. (2019). How affordable are 2019 aca premiums for middle-income people? https://www.kff.org/health-reform/issue-brief/how-affordable-are-2019-aca-premiums-for-middle-income-people/.

Finkelstein, A. and Poterba, J. (2014). Testing for asymmetric information using "unused observables" in insurance markets: Evidence from the uk annuity market. *Journal of Risk and Insurance*, **81** (4), 709–734.

Gardner, J. (2021). Two-stage differences in differences. *Working Paper*.

Gaynor, M., Haas-Wilson, D. and Vogt, W. B. (2000). Are invisible hands good hands? moral hazard, competition, and the second-best in health care markets. *Journal of Political Economy*, **108** (5), 992–1005.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, **225** (2), 254–277.

Handel, B. R. and Kolstad, J. T. (2021). *The Affordable Care Act After a Decade: Industrial Organization of the Insurance Exchanges*. Tech. rep., National Bureau of Economic Research.

Holahan, J., Buettgens, M., Carroll, C. and Dorn, S. (2012). The Cost and Coverage Implications of the ACA Medicaid Expansion: National and State-by-State Analysis. https://www.kff.org/health-reform/report/the-cost-and-coverage-implications-of-the/.

Holmes, J. (2022). Cross-Market Selection In U.S. Health Insurance Markets . *Working Paper*.

Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, **10** (1), 74–91.

Jakiela, P. (2021). Simple diagnostics for two-way fixed effects. *arXiv preprint arXiv:2103.13229*.

Kaiser Family Foundation (2019). Marketplace average benchmark premiums. https://www.kff.org/health-reform/state-indicator/marketplace-average-benchmark-premiums.

KFF (2020). As unemployment skyrockets, kff estimates more than 20 million people losing job-based health coverage will become eligible for aca coverage through medicaid or marketplace tax credits — kff. https://www.kff.org/coronavirus-covid-19/press-release/as-unemployment-skyrockets-kff-estimates-more-than-20-million-people-losing-job-based-health-coverage-will-become-eligible-for-aca-coverage-through-medicaid-or-marketplace-tax-credits/.

Lavetti, K., DeLeire, T. and Ziebarth, N. R. (2023). How do low-income enrollees in the affordable care act marketplaces respond to cost-sharing? *Journal of Risk and Insurance*, **90** (1), 155–183.

Lee, D. S. and Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, **48** (2), 281–355.

Manning, W. G., Newhouse, J. P., Duan, N., Keeler, E. B. and Leibowitz, A. (1987). Health insurance and the Demand for Medical Care: Evidence from a Randomized Experiment. *The American Economic Review*, **77** (3), 251–277.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, **142** (2), 698–714.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, **37** (2), 187–204.

Pauly, M. V. (1968). The economics of moral hazard: Comment. *The American Economic Review*, **58** (3), 531–537.

Peng, L. (2017). How Does Medicaid Expansion Affect Premiums in the Health Insurance Marketplaces? New Evidence from Late Adoption in Pennsylvania and Indiana. *American Journal of Health Economics*, **3** (4), 550–576.

Rao, P. and Nowak, S. A. (2019). Effects of alternative insurer responses to discontinued federal cost-sharing reduction payments.

Rosenbaum, S. and Wilensky, G. (2020). Closing the medicaid coverage gap: Options for reform. *Health Affairs*, **39** (3), 514–518.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, **90** (4), 629–649.

Royalty, A. B. and Solomon, N. (1999). Health plan choice: Price elasticities in a managed competition setting. *Journal of Human Resources*, **34** (1), 1–41.

Sen, A. P. and DeLeire, T. (2018). How does expansion of public health insurance affect risk pools and premiums in the market for private health insurance? Evidence from Medicaid and the Affordable Care Act marketplaces. *Health Economics*, **27** (12), 1877–1903.

Shepard, M. (2022). Hospital network competition and adverse selection: evidence from the massachusetts health insurance exchange. *American Economic Review*, **112** (2), 578–615.

Simonsohn, U., Simmons, J. P. and Nelson, L. D. (2020). Specification curve analysis. *Nature Human Behaviour*, **4** (11), 1208–1214.

Sprung, A. and Anderson, D. (2018). Mining the silver lode. *Health Affairs Blog*.

Appendix A Proofs

Proposition 1: Existence of an Equilibrium Premium
A zero profit equilibrium implies that all insurers price at average cost. Average costs are simply the weighted average cost of consumers in each group which is positive from assumption A1c:

$$\bar{c}(p) = \frac{D_L(p_L) \cdot M_L \cdot c(O_L) + D_H(p) \cdot M_H \cdot c(O_H)}{D_L(p_L) \cdot M_L + D_H(p) \cdot M_H} \tag{A1}$$

Under assumptions A1 and A4, there exists a $p^* \in P$ such that $\bar{c}(p^*) = p^*$.

Proof. Because average cost, $\bar{c}$, is a function of p, we show that there exists a fixed point. Given the continuity of demand (Assumption A1), $\bar{c}(p)$ is continuous in p. By Assumption A4, P is compact and convex. Therefore, by Brouwer's Fixed Point Theorem a p^* exists such that $\bar{c}(p^*) = p^*$.

□

Lemma 1:

$$0 < \frac{\partial \bar{c}}{\partial p} = -\left(c(O_L) - c(O_H)\right) \cdot \left[\frac{D_L(p_L) \cdot M_L}{\left[D_L(p_L) \cdot M_L + D_H(p) \cdot M_H\right]^2}\right] \cdot D'_H(p) \cdot M_H < 1$$

Proof. The equation is direct from differentiating Equation A1. This is easier to see when rewritten as below:

$$\bar{c}(p) = c(O_H) + \left(c(O_L) - c(O_H)\right)\left[\frac{D_L(p_L) \cdot M_L}{D_L(p_L) \cdot M_L + D_H(p) \cdot M_H}\right]$$

By Assumption A1, $D'_H(p) < 0$. If Assumption A2 holds all other terms are positive (except the negative sign). Therefore, $0 < \frac{\partial \bar{c}}{\partial p}$. By Assumption A3, $\frac{\partial \bar{c}}{\partial p} < 1$.

□

Lemma 2:

Under assumption A3:

$$\frac{\partial p^*}{\partial M_L} = \frac{\frac{\partial \bar{c}}{\partial M_L}}{1 - \frac{\partial \bar{c}}{\partial p}} \tag{A2}$$

Proof. Define $G = p - \bar{c}(p)$. As required by the implicit function theorem, $\frac{\partial G}{\partial p} \neq 0$. This is immediate from Assumption A3 and the definition of G. Also, G is continuous as demand is continuous and Equation A1.

By the implicit function theorem:

$$\frac{\partial p^*}{\partial M_L} = \frac{-\frac{\partial G}{\partial M_L}}{\frac{\partial G}{\partial p}|_{p=p^*}} = \frac{\frac{\partial \bar{c}}{\partial M_L}}{1 - \frac{\partial \bar{c}}{\partial p}}$$

□

Lemma 3:

$$\frac{\partial \bar{c}}{\partial M_L} = \left[c(O_L) - c(O_H)\right] \cdot \frac{\left[D_L(p_L) \cdot D_H(p) \cdot M_H\right]}{\left[D_L(p_L) \cdot M_L + D_H(p) \cdot M_H\right]^2} > 0 \tag{A3}$$

Proof. The equation is direct by differentiating Equation A1. By Assumption A2 (moral hazard), the term in brackets is positive. By Assumption A1, demand is positive. By Assumption A1, the denominator is positive, hence $\frac{\partial \bar{c}}{\partial M_L} \geq 0$.

□

Result 1: $\frac{\partial p^*}{\partial M_L} > 0$

Proof. By Lemma 2, we can rewrite $\frac{\partial p^*}{\partial M_L}$ as shown in Equation A2. Lemma 3 proves that the numerator of Equation A2 is positive. The denominator of Equation A2 is positive as a consequence of Assumption A3.

□

Result 2: $\frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L} > 0$

Proof. Following Equation A2, we just need to show that $0 < \frac{\partial \bar{c}}{\partial p} < 1$. This is shown in Lemma 1.

□

The are some additional modifications required for result 3, as we have modified the definition of average and marginal costs, and are adding adverse selection. Because $ac(\cdot)$ is a function of premiums, we need to modify Assumption A2 (moral hazard) to account for this.

Assumption A2a: Moral Hazard $ac_i(p, O_H) < ac_i(p, O_L)$ for any p, if $O_L < O_H$.

Assumption A2a states that for a given population, if out-of-pocket costs are lower, then their average costs are higher, regardless of the premium. Because we need a comparison across groups, this condition is too restrictive. Instead, we need an comparison of average costs across groups.

Assumption A2b: Moral Hazard $ac_H(p, O_H) < ac_L(p_L, O_L)$ for any p, if $O_L < O_H$.

Assumption A2b states that low income consumers are more expensive than high income consumers at any premium. Because in our empirical setting the difference in out of pocket costs is large enough that we think moral hazard is driving this relationship, though other mechansisms, like the health/wealth gradient, could play a factor as lower income people are generally sicker than high income people.

In addition, we need to define what it means to worsen the degree of adverse selection. Our goal is to see how a ceteris paribus change in selection impacts MHIU, which we are exploring with small changes in M_L. If one thinks of adverse selection as changing the slope of the average cost curve, then for a ceteris paribus change in adverse selection one would need to pin down a point to rotate the average cost curve around. This is a complicated question as small changes in adverse selection can lead to large changes in equilibrium outcomes. While interesting, we are focused on how adverse selection impacts MHIU, on the margin. Hence, we hold objects fixed near the equilibrium and rotate the average cost curve around this equilibrium point.

As we are focusing on small changes, we can define adverse selection as $\frac{\partial ac_H(p, O_H)}{\partial p}$ which is a scalar corresponding to the slope of the average cost curve around price p. In this case, more adverse selection means the slope is steeper: $\frac{\partial ac_H(p, O_H)}{\partial p}$ is a larger scalar. In Section 3.6, this corresponds to changing δ, when we assume linear costs for the unsubsidized.

Result 3: $\frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L}$ is larger when $\frac{\partial ac_H(p,O_H)}{\partial p}$ is larger, holding other objects fixed.

Proof. By Lemma 2, we can rewrite $\frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L}$ as:

$$\left[\frac{\frac{\partial \bar{c}}{\partial p}}{1 - \frac{\partial \bar{c}}{\partial p}}\right] \cdot \frac{\partial \bar{c}}{\partial M_L}$$

For simplification, let $t = \frac{\partial ac_H(p,O_H)}{\partial p}$, $X(t) = \frac{\partial \bar{c}}{\partial p}$, and $Y = \frac{\partial \bar{c}}{\partial M_L}$. Then, let

$$A(t) = \frac{\partial p^*}{\partial M_L} - \frac{\partial \bar{c}}{\partial M_L} = \frac{X(t)}{1 - X(t)} \cdot Y \tag{A4}$$

In this formulation, we need to show that:

$$\frac{\partial A(t)}{\partial t} > 0$$

Because Y does not depend on t, by the Chain Rule it follows that:

$$\frac{\partial A(t)}{\partial t} = \frac{\partial A(t)}{\partial X(t)} \cdot \frac{\partial X(t)}{\partial t} \tag{A5}$$

We need to show that both of the terms in Equation A5 are positive. From Equation A4, $\frac{\partial A(t)}{\partial X(t)} = \frac{Y}{[1-X(t)]^2}$, which is positive if $X(t) \neq 1$ and $Y > 0$. By Assumption A3, $X(t) \neq 1$. Lemma 3 proves that $Y > 0$.[43]

Now consider the second term in Equation A5. Differentiating Equation A1 with respect to p (but using the cost function implied by Assumption A2b) yields:

$$\begin{aligned}\frac{\partial \bar{c}}{\partial p} = \frac{\partial ac_H(p,O_H)}{\partial p} \cdot \left[1 - \frac{D_L(p_L) \cdot M_L}{D_L(p_L) \cdot M_L + D_H(p) \cdot M_H}\right] \\ - \left(ac_L(p_L,O_L) - ac_H(p,O_H)\right) \cdot \left[\frac{D_L(p_L) \cdot M_L}{\left[D_L(p_L) \cdot M_L + D_H(p) \cdot M_H\right]^2}\right] \cdot D'_H(p) \cdot M_H\end{aligned} \tag{A6}$$

The second line of this equation is the same as the impact due to MHIU (see Lemma 1). Because $X(t)$ was defined as $\frac{\partial \bar{c}}{\partial p}$, we need to differentate equation A6 with respect to $\frac{\partial ac_H(p,O_H)}{\partial p}$ to get $\frac{\partial X(t)}{\partial t}$. From equation A6, it should be clear that $\frac{\partial X(t)}{\partial t}$ will be positive so long as long as the term in brackets on the first line is positive, which is assured by Assumption A1.

□

[43]Lemma 3 was proved using Assumption A2, but it can be easily shown to hold under Assumption A2b as well.

Online Appendix B Demand Estimation Robustness

B.1 Summary of Enrollment Measures:

B.2 Data and Sample Construction Details

The HIX OEP data give us a direct measure of enrollment on the exchanges, but it has some drawbacks in availability. First, it does not provide rating area level data in 2014. Therefore, for analyses that use the HIX enrollment as the dependent variable, we drop the three states that expanded in 2015 – New Hampshire, Pennsylvania, Michigan – since we do not observe a pre-period. This sample has 17 states and 249 rating areas. Also, in 2017 the HIX OEP combines the 400+% FPL enrollment and the less than 100% FPL enrollment. For analyses that use the 400+% FPL HIX enrollment, we are limited to 2015-2016. We measure the uninsured, direct purchase, and employer-sponsored insurance rate with the ACS data. One advantage of the ACS data is that we have a longer sample period, so we can include all of our late expanding sample states. For the uninsured rate and employer-sponsored rate, we use 2012-2018 data and include states that expanded Medicaid in 2014 as we have pre-period data.[44] We do not include earlier years for the direct purchase measure since the pre-2014 direct purchase market is very different.

To define our sample, we begin with the 37 states, plus Washington D.C., that reported to the HIX OEP database from 2015-2017. Then, we drop Alaska and Hawaii because their FPL levels are not defined consistently with the other states. We drop an additional eight states because their rating areas are not defined by counties, which complicates the mapping from counties to rating areas in the ACS data. For our main analysis, we also drop the eight states who meet these criteria but expanded Medicaid in 2014. However, we include these states in robustness checks. The states in our sample are Arkansas, Florida, Georgia, Indiana, Iowa, Kansas, Louisiana, Maine, Michigan, Mississippi, Missouri, Montana, New Hampshire, North Carolina, Pennsylvania, South Carolina, South Dakota, Tennessee, Utah, and Wisconsin. The eight states included in our sample when we include those that expanded in 2014 are Arizona, Delaware, Illinois, New Jersey, Nevada, Ohio, Oregon, and West Virginia.

We use the IPUMS data portal to construct the ACS data using the one-year sample files. We do this because the federal poverity level the income to poverty ratio in the ACS does not correspond to how the federal poverty level is measured for subsidy eligibility. The State Health Access Data Assistance Center (SHADAC) has a correction for this, but it is only available in the one year sample files. For more details, please refer to https://www.shadac.org/publications/using-shadac-health-insurance-unit-hiu-and-federal-poverty-guideline-fpg-microdata.

Our ACS sample consists of the under-65 population. We also drop those who are enrolled in Medicare or Medicaid, since these individuals are not eligible for subsidies on the exchanges. Our primary dependent variables use the ACS data focusing on the 400+% FPL level, where Medicaid enrollment is very uncommon. The ACS provides geographic information at the Public Use Microdata Area (PUMA) level, which is not nested within rating areas. To account for this, we expand the data to the individual county level. We reweight individual-counties by the probability their PUMA is in that given county, then sum that individual's weight to the individual-rating area level – where an individual can be in multiple rating areas. To avoid confusion, we discuss our unit of observation as an individual, though in cases where an individual is in a PUMA which corresponds to multiple rating areas, we count that individual in each rating area, but then reweight

[44] The states that expanded Medicaid in 2014 are Arizona, Delaware, Illinois, New Jersey, Nevada, Ohio, Oregon, and West Virginia. States that expanded Medicaid prior to 2014 are excluded from all analyses.

them. For example, if 75 percent of the population in a given PUMA lives in county A and 25 percent in county B, for each individual in that PUMA, we assign them to both counties with 75 percent of their weight going to county A and 25 percent going to county B, then sum enrollment across counties to the rating area level. The county to PUMA weights were accessed from the Missouri Census data center here: http://mcdc.missouri.edu/applications/geocorr2014.html.

B.3 Parallel Trends:

Our identification in Subsection 4.1.2 relies on the assumption of parallel trends. That is, in the absence of treatment our treated units would have continued on the same trajectory as prior to treatment. Additionally, parallel trends requires that the control units are a good counterfactual for our treated units. Since parallel trends cannot be directly observed as the counterfactual outcome of treated units not being treated cannot be observed, we test for parallel trends by observing the pre-trends. To do this, we use the imputed difference-in-differences method described in Subsection 4.1.2 including leads and lags. Figure B.1 presents the results from the event study to examine the parallel trends utilized in our difference-in-differences analysis. Given that there is no data for HIX enrollment/premiums prior to 2014, we examine these parallel trends for the uninsured rate for the 400+% FPL population. Prior to Medicaid expansion (which occurs in "Time to Treatment" year 0), we see that expanding states are not statistically different than non-expanding states. Additionally, we see no evidence of a pre-trend in the uninsured rate for the unsubsidized population. After expansion, there is a noticeable dip in the uninsured rate for expansion states, which becomes statistically significant at the 5% level 4 years after expansion. This indicates that it takes time for (1) premiums to decrease enough for the unsubsidized population to enroll and (2) the dissemination of information about the HIXs and the lower premiums.

Figure B.1: Event Study for Uninsured Rate for the 400+% FPL Population

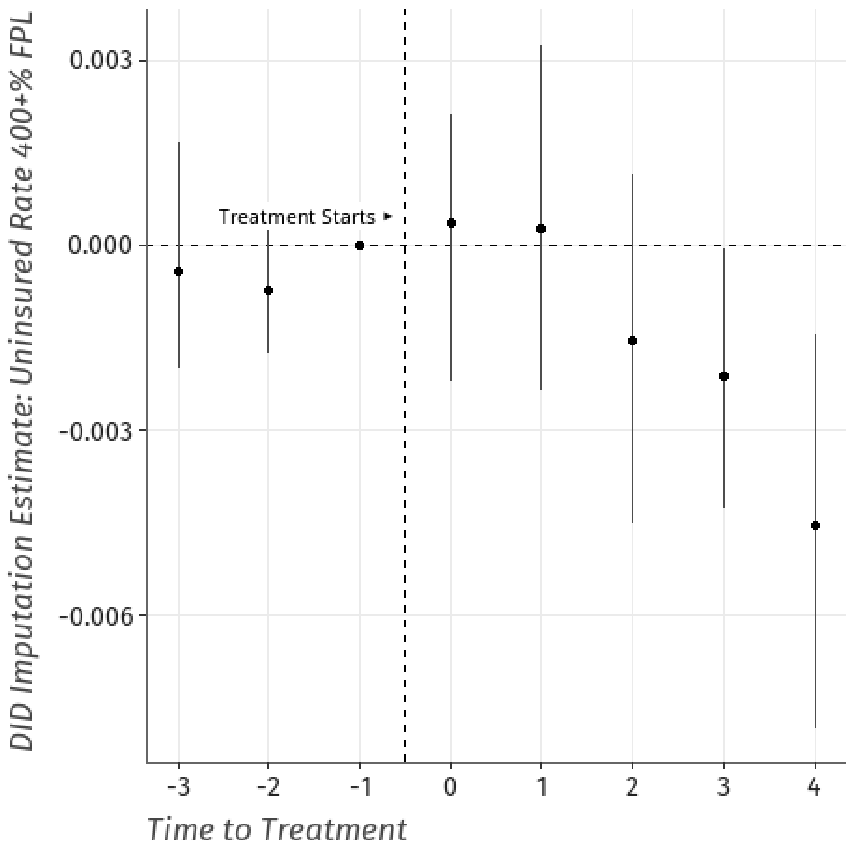

Note: The figure plots the coefficients and 95 percent confidence interval for the imputed difference-in-difference estimates for whether or not someone is uninsured. The regression includes state and year fixed effects, as well as individual level controls. Year "0" refers to the year Medicaid expanded and Year "-1" is the omitted category. Standard errors are block bootstrapped at the state level. Data is from the American Community Survey.

B.4 Falsification and Robustness:

One may be concerned that the enrollment effects shown in Table 2 in Section 4.1.3 are not unique to the unsubsidized population. That is, individuals 150-400% FPL could also be gaining insurance through the HIXs at increased rates. We examine this possibility by using the imputed difference-in-differences method described in Section 4.1.2 for the uninsured rate and HIX enrollment by individuals that receive subsidies on the exchanges. Columns (1) and (2) of Table B1 show that we fail to find enrollment effects for the subsidized population. One could also be concerned that the decrease in the uninsured rate for the unsubsidized population shown in Column (4) of Table 2 is due to changes in the employer-sponsored market. Column (3) of Table B1 shows that we find a very small and statistically insignificant effect for the employer-sponsored insurance rate for the unsubsidized population.

Table B1: Placebo Checks: Enrollment Effects

	150-400% FPL:		400+% FPL:
	(1)	(2)	(3)
	Uninsured Rate (% point):	P(HIX Purchase):	Employer-Sponsored Insurance Rate:
Estimated ATT	−0.414	−0.001	0.001
90% Confidence Interval	[-0.597, 1.193]	[-0.014, 0.013]	[-0.001, 0.004]
Pre-Expansion Sample Mean	14	0.435	0.962
Observations	3,631,926	499,980	3,595,818

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table shows the results of the difference-in-differences estimation as well as corresponding pre-expansion sample means for (1) the uninsured rate for subsidized individuals, (2) HIX enrollment by subsidized individuals, and (3) the employer-sponsored insurance rate by individuals 400+% FPL. (i.e., unsubsidized population). All models include person controls, year fixed effects, and state fixed effects. Standard errors for all models are block bootstrapped at the state level. Column (2) considers the 150-400% FPL population (i.e., subsidized population) who are either uninsured or purchase on the exchanges (labeled as direct purchase in the ACS). Column (3) is the entire 400+% population. Columns (1) and (3) includes states that expanded Medicaid in 2014 and the years of ACS data utilized are 2012-2018. Column (2) excludes states that expanded Medicaid in 2014 (as there is no pre-data) and the years of ACS data utilized are 2014-2018.

In Table B2 we check to ensure that the choice to include the 2018 ACS data is not driving our results for Column (3) (and therefore column (5)) in Table 2. Columns (1) and (3) of table B2 show the imputed difference-in-differences estimate for HIX enrollment for individuals 400+% FPL using 2014-2018 data and 2014-2017 data, respectively. Columns (2) and (4) present the Demand IV estimates utilizing the premium results from Table 2 and Columns (1) and (3), respectively. Columns (3) and (4) are the estimates presented in Table 2. We see that the exclusion of 2018 data has no statistically significant difference on the enrollment estimate or the demand estimate.

One may additionally be concerned about how sensitive our estimates our to modelling choices. To test the robustness of our results to different combinations of modelling choices, we estimate specifications curves which allow us to try many different combinations of modelling choices and view the results graphically.

Table B2: IV Estimate: Robustness Check

	2014-2018:		2014-2017:	
	(1)	(2)	(3)	(4)
	P(HIX Purchase):	Demand IV	P(HIX Purchase):	Demand IV
Estimated ATT	0.027***	−0.071***	0.029***	−0.078***
90% Confidence Interval	[0.011, 0.04]	[-0.11, -0.043]	[0.012, 0.043]	[-0.125, -0.05]
Implied Intercept	-	0.96***	-	0.981***
90% Confidence Interval		[0.864, 1.099]		[0.887, 1.14]
Pre-Expansion Sample Mean	0.74	-	0.74	-
Implied Post-Expansion Mean	0.767	-	0.769	-
Implied No-MHIU Mean	0.811	-	0.818	-
Observations	269,537	-	213,208	-

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table shows the results of the difference-in-differences estimation as well as corresponding pre-expansion sample means for HIX enrollment by subsidized individuals and the corresponding demand estimates. Columns (1) and (3) use the 150-400% FPL population who are either uninsured or purchase insurance. Columns (1) and (3) present the difference-in-difference estimates for the impact of Medicaid expansion on HIX enrollment using the ACS data from 2014 to 2018 and 2014 to 2017, respectively. Columns (3) and (4) match Columns (3) and (5) from Table 2. Columns (1) and (3) include person controls, year fixed effects, and state fixed effects. Standard errors for all models are block bootstrapped and clustered at the state level. The results are robust to the inclusion of 2018 data.

As Figure B.2 shows, our results are robust to the inclusion/use of various controls and the inclusion/exclusion of Arkansas. Of the 112 regressions, all are statistically significant at the 10% level. Our preferred specification, shown in the red line, is on the conservative side.

Figure B.3 shows the specification curve for the imputed difference-in-difference estimate for the direct purchase rate of those 100-150% FPL on the HIXs. Our results are robust to the inclusion/use of various controls, the inclusion/exclusion of Arkansas, and the years of data utilized with all 104 regressions being statistically significant at the 10% level. Our preferred specification, shown with the red line, represents a conservative estimate.

The specification curve for the imputed difference-in-difference estimate for the log share of those 100-150% FPL of the HIXs is shown in Figure B.4. As shown in the bottom portion of the figure, there is no discernible pattern indicating that our results are robust to various specifications. Of the 88 regressions, all are statistically significant at the 10% level. While our preferred specification (shown in the red line) is on the larger side of a decrease, the figure shows that the point estimates are quite tight around -0.45 and -0.47.

Figure B.2: Specification Curve: Premium Results

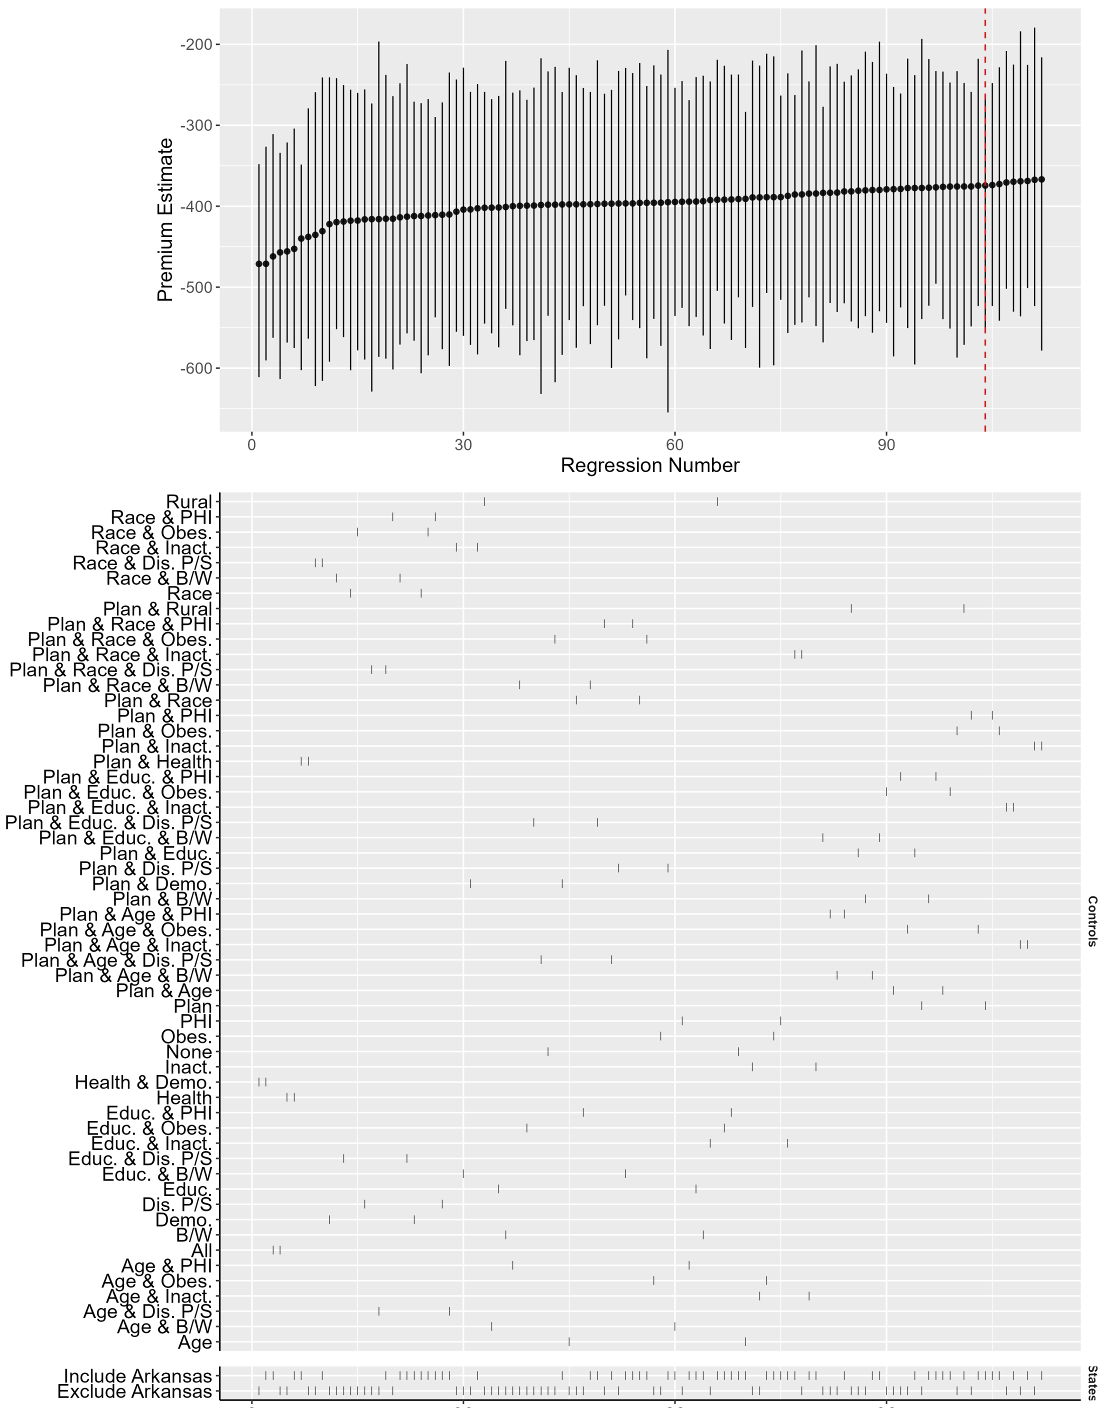

This figure plots estimated premium changes, with 90% confidence intervals, from 112 different regressions. The numbering on the x-axis indicates the regression number. The panels below the primary curve indicate the modeling options that can be varied for the specification curve. The options include various controls variables. Options shaded black are statistically significant at the 10% level; options shaded grey are not. You can tell how each choice impacts the point estimate by patterns in the second half of figure. Options that appear more often on the left of the graph tend to drive the estimate away from zero and options that are more on the right tend to drive the estimate toward zero. The red line indicates our preferred specification. Standard errors for all models are block bootstrapped and clustered at the state level. The various controls are combinations of Plan, Health, Demographic (Demo.), Education (Educ.), Race, Age, Rural, Poor Health Indicators (PHI), Disease Prevalence/Screening (Dis. P/S), Obesity (Obes.), Birthweight (B/W), and Inactivity (Inact.).

Figure B.3: Specification Curve: Direct Purchase

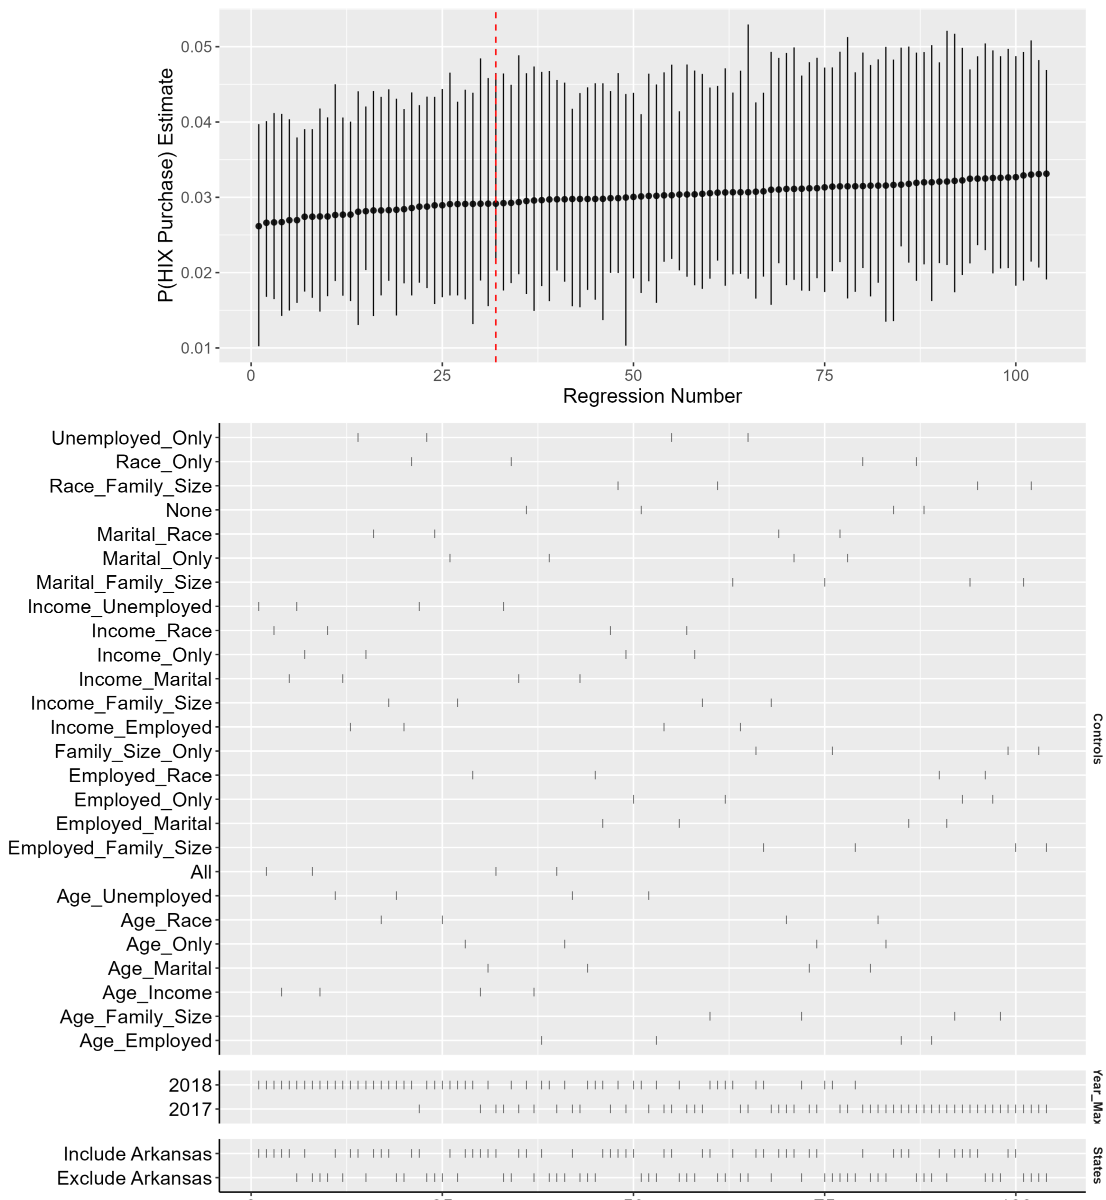

This figure plots the imputed difference-in-difference estimates for the direct purchase rate of those 100-150% FPL on the HIXs, with 90% confidence intervals, from 104 different regressions. The numbering on the x-axis indicates the regression number. The panels below the primary curve indicate the modeling options that can be varied for the specification curve. The options include various controls variables. Options shaded black are statistically significant at the 10% level; options shaded grey are not. You can tell how each choice impacts the point estimate by patterns in the second half of figure. Options that appear more often on the left of the graph tend to drive the estimate away from zero and options that are more on the right tend to drive the point toward zero. The red line indicates our preferred specification. Standard errors for all models are block bootstrapped and clustered at the state level.

Figure B.4: Specification Curve: Log Share 100-150% FPL

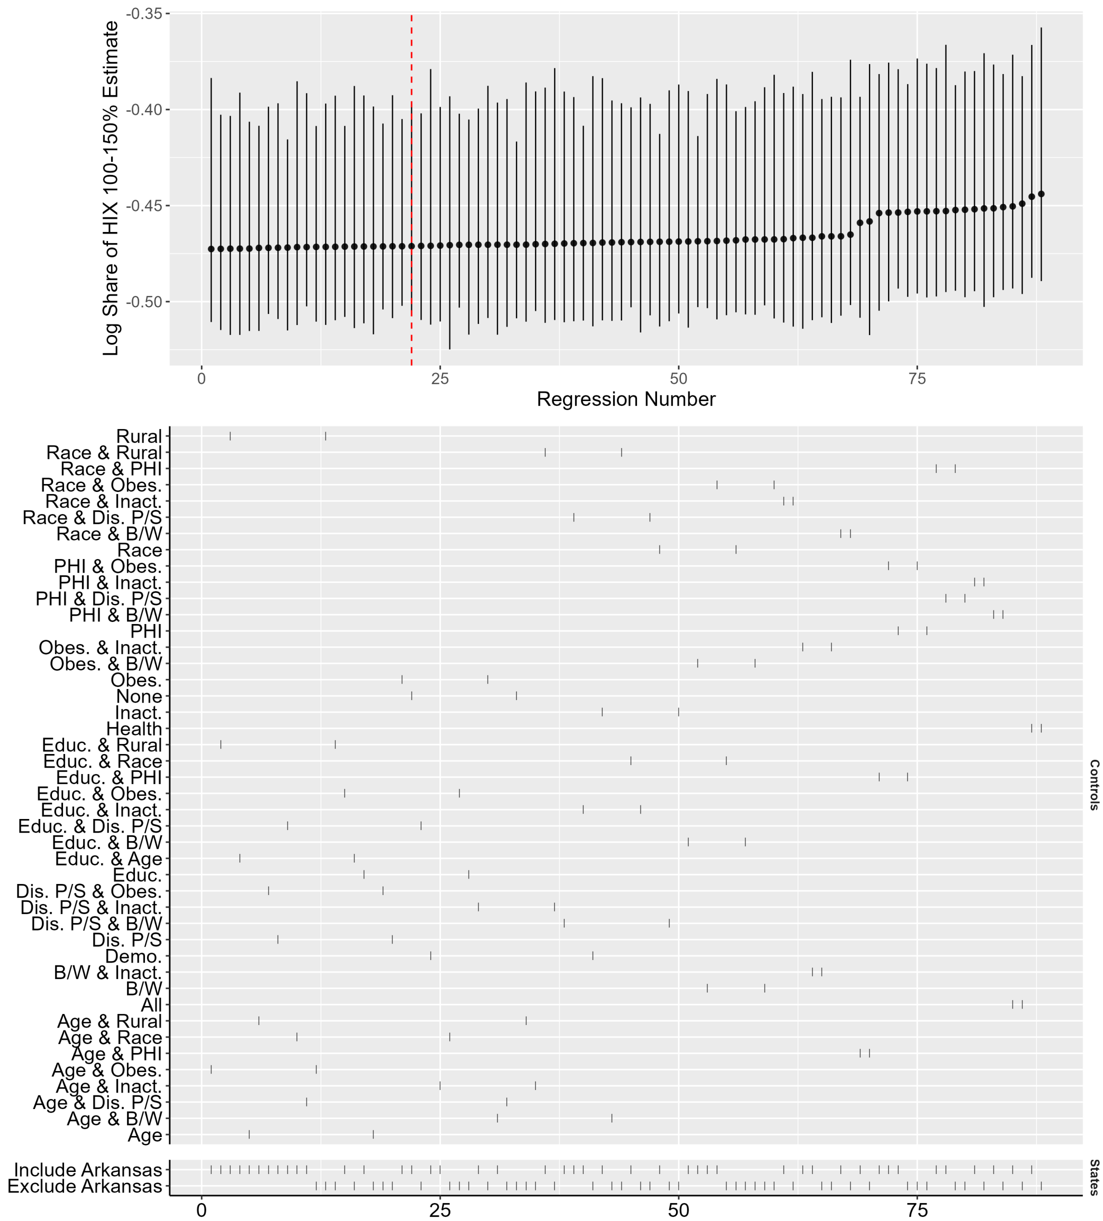

This figure plots the imputed difference-in-difference estimates for the log share of those 100-150% FPL on the HIXs, with 90% confidence intervals, from 88 different regressions. The numbering on the x-axis indicates the regression number. The panels below the primary curve indicate the modeling options that can be varied for the specification curve. The options include various controls variables. Options shaded black are statistically significant at the 10% level; options shaded grey are not. You can tell how each choice impacts the point estimate by patterns in the second half of figure. Options that appear more often on the left of the graph tend to drive the estimate away from zero and options that are more on the right tend to drive the estimate toward zero. The red line indicates our preferred specification. Standard errors for all models are block bootstrapped and clustered at the state level. The various controls are combinations of Health, Demographic (Demo.), Education (Educ.), Race, Age, Rural, Poor Health Indicators (PHI), Disease Prevalence/Screening (Dis. P/S), Obesity (Obes.), Birthweight (B/W), and Inactivity (Inact.).

Online Appendix C Moral Hazard Estimation Robustness

C.1 Summary Statistics:

Table C1 shows summary statistics for 2014-2018. We separate healthcare expenditures and visits into total, discretionary, and inpatient where total spending is the sum of discretionary, inpatient, dental and vision.[45,46] Inpatient visits and expenditures are used as partial placebo analyses because these visits include emergency procedures and should be less discretionary.[47]

One may be concerned about the accuracy of the MEPS data given that much of the information is self-reported. In the event that the household survey respondent provided inaccurate or incomplete information, MEPS will update or supplement the household survey response. For more information please refer to https://meps.ahrq.gov/data_stats/download_data/pufs/h201/h201doc.shtml#Household10 and https://meps.ahrq.gov/survey_comp/mpc_data_collection.jsp. MEPS does not verify the income of respondents. However, respondents are contacted in advance of the MEPS interviews with information about the records (including income) that will be asked about during the interview and the respondents are provided with a small financial incentive for their time and in maintaining these records. Please refer to https://www.meps.ahrq.gov/data_files/publications/mr1/mr1.shtml#Household for more information. Misreported incomes may lead to measurement error in our data. As long as the misreporting of income is not correlated with health status this misreporting should have an attenuating bias.

One may be also be concerned about the lack of metal level data for the MEPS. Metal level data is not available in 2014 and close to 37% of those on the exchanges do not provide information on the metal level of their plan. Therefore, we do not consider the metal level in our analysis. Sprung and Anderson (2018) find that 80-90% of those in the 100-200% FPL range select a silver plan.

Because we exploit the sharp discontinuity in the CSRs at 150% FPL, the first four columns of Table C1 are separated by the 9% block above and below the cutoff. On average, individuals on the ACA marketplaces with incomes below 150% FPL cutpoint spend more on healthcare and visit their healthcare provider more frequently than those slightly above the 150% FPL cutpoint.

The final column of Table C1 examines the balance across the CSR eligibility threshold for the predetermined covariates: age, educational attainment, sex, race and marital status. If treatment cannot be manipulated, there should be no systematic differences across the treatment eligibility threshold for these predetermined covariates. While there are some differences (particularly in sex and race), none of the predetermined covariates are statistically different at the 5% level.[48]

C.2 Specification Checks:

One may be worried about the robustness of our results based on various specifications. To address this, we estimate what is known as a specification curve (Simonsohn *et al.*, 2020). The purpose of a specification curve is to graphically summarize how the estimates from the model change based on various potential modeling choices made by the researcher. Figure C.1 plots our

[45] Discretionary expenditures/visits excludes services that are not covered by health insurers on the exchanges (i.e., dental services and vision) as well as inpatient services. Inpatient services are excluded as these are used as a partial placebo check.

[46] The inpatient numbers are unconditional on using inpatient care. The average inpatient visit expenditure is $18,892 for those with at least one inpatient visit.

[47] Lavetti *et al.* (2023) also find that inpatient expenditures do not seem impacted by the CSR Subsidies on the HIXs in Utah.

[48] For a more formal examination of the balance of predetermined covariates please see Appendix Section C.3.

Table C1: Summary Statistics: Medical Expenditure Panel Survey

Variable	140% to 149% FPL:		151% to 160% FPL:		Diff. Means:
	Mean	St. Dev.	Mean	St. Dev.	t
Total Healthcare Expenditures	3,810.43	8,653.18	2,061.65	4,838.57	1.63
Total Healthcare Visits	5.98	8.38	4.11	7.02	1.58
Total Discretionary Expenditures	2,434.50	6,795.21	872.43	1,488.60	2.06
Total Discretionary Visits	5.13	8.06	3.27	6.44	1.67
Inpatient Healthcare Expenditures	914.23	4,794.55	973.57	4,242.68	−0.086
Inpatient Healthcare Visits	0.06	0.28	0.068	0.30	−0.20
Age	43.39	15.51	40.39	16.00	1.25
Education – High School	0.13	0.34	0.12	0.33	0.12
Sex – Male	0.32	0.47	0.42	0.50	−1.34
Race – White	0.76	0.43	0.65	0.48	1.65
Race – Black	0.12	0.33	0.16	0.37	−0.76
Marital Status – Single	0.29	0.45	0.25	0.44	0.53
N	84		88		-

Note: This table shows the summary statistics for the 2014 to 2018 Medical Expenditure Panel Survey for individuals below the age of 65 from the ACA marketplaces. Columns (1) and (2) provide the summary statistics for those individuals living from 140% to 149% of the FPL. Columns (3) and (4) provide the summary statistics for those living 151% to 160% of the FPL. The final column displays the t-value for a difference in means test for those with income 140% to 149% compared to 151% to 160% of the FPL.

estimated elasticity, with 90% confidence intervals, from 2,016 different regression discontinuities, which are the result of different combinations of our modeling choices. We had modeling control over (1) the use and size of the donut; (2) the size of the RD bandwidth; (3) the type of RD kernel; (4) the transformation of the expenditure variable; (5) the type of control variables; and (6) the type of fixed effects.[49] Model options shaded black are statistically significant at the 10% level; whereas model options shaded grey are not. The visual patterns in the bottom half of Figure C.1 show how each choice impacts the elasticity estimate. Options that appear more often on the left of the figure tend to drive the elasticity estimate away from zero and options that are more on the right tend to drive the elasticity estimate toward zero.

Every point estimate is below zero with 80.2% of these estimates being statistically below zero. The first takeaway from the specification curve is that the inclusion of controls moves our elasticity estimate towards zero. While the majority (66%) of results with no donut are statistically significant, Figure C.1 indicates that the use of a donut moves our elasticity estimate away from zero with larger donuts moving the estimate further from zero. This is consistent with an attenuating bias due to measurement error in income. Additionally, our choice of log + 1 vs. inverse hyperbolic sine does have a small effect. The latter tends to estimate a slightly more elastic healthcare demand, and thus more moral hazard, than the more traditional log transformation. Finally, the kernel bandwidth size, kernel choice, or choice of fixed effects do not seem to have a noticeable impact on the elasticity estimates.

[49]Because health expenditures are right-skewed and zero-inflated, we also use the inverse hyperbolic sine, which allows us to retain zero-valued observations. Elasticities are calculated by Bellemare and Wichman (2020).

Figure C.1: Specification Curve

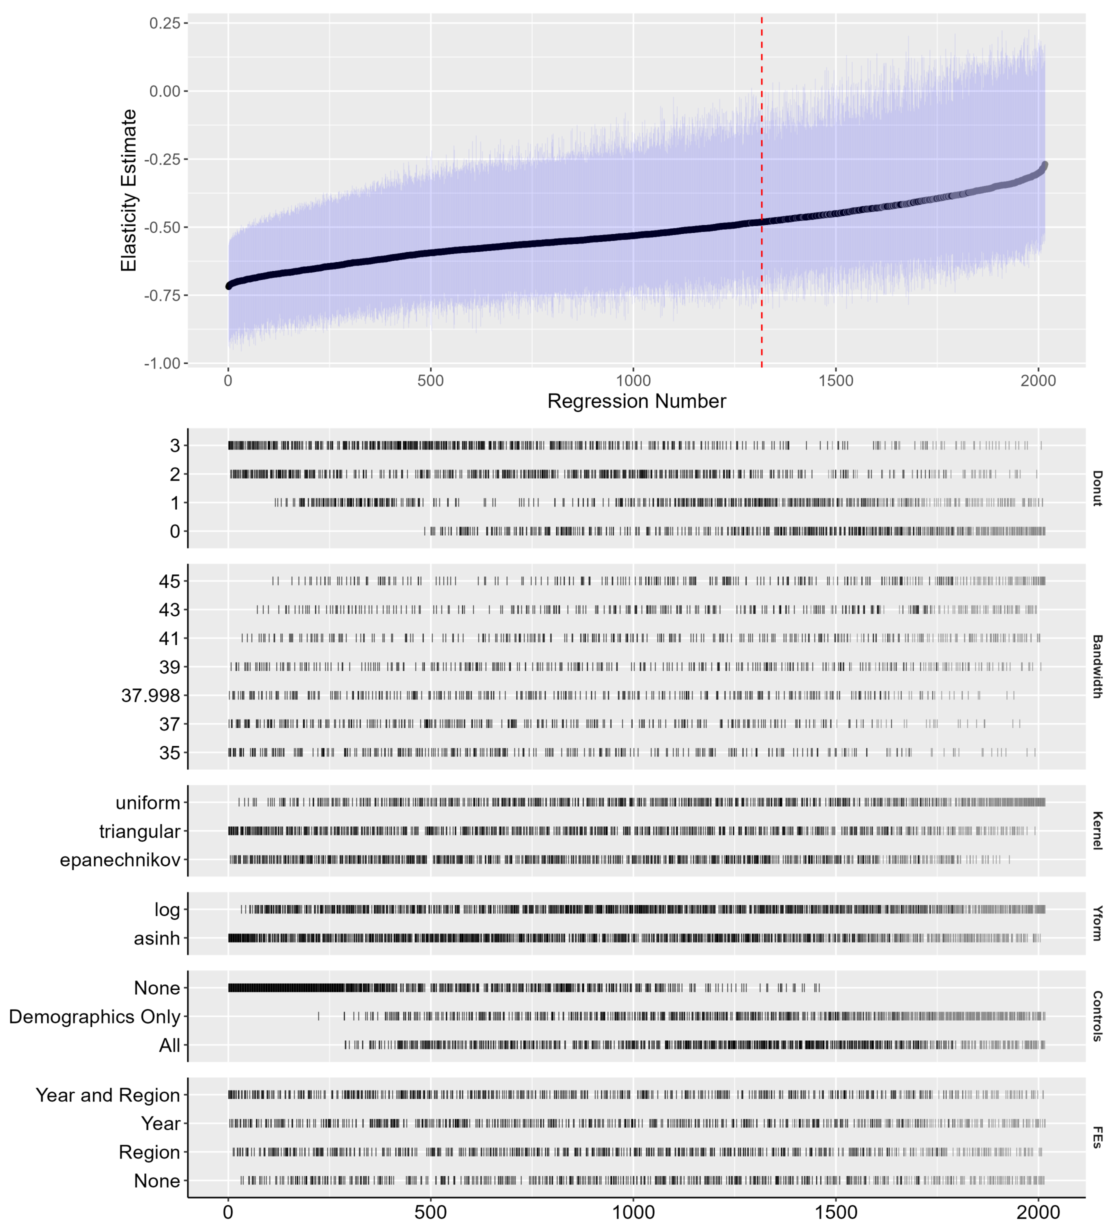

This figure plots estimated elasticities, with 90% confidence intervals, from 2,016 different regression discontinuities. The numbering on the x-axis indicates the regression number. The panels below the primary curve indicate the modeling options that can be varied for the specification curve. The options include (1) the use and size of a donut, (2) the size of the RD bandwidth, (3) the RD kernel, (4) the transformation of the expenditure variable, (5) the use and type of control variables, and (6) the choice of fixed effects. Options shaded black are statistically significant at the 10% level; options shaded grey are not. You can tell how each choice impacts the elasticity estimate by patterns in the second half of figure. Options that appear more often on the left of the graph tend to drive the elasticity away from zero and options that are more on the right tend to drive the elasticity toward zero. The red line indicates our preferred specification. The health control is the total predicted healthcare costs. Additional control variables are marital status, sex, age, education, family size and race. Note the estimates and confidence intervals are not bias corrected as the bias correction is to account for bandwidths derived from MSE optimization, which the specification curve does not use.

C.3 Balance Tests:

For a more formal examination of the balance of predetermined covariates, we examine whether a regression discontinuity exists at the income threshold of 150% FPL for our predetermined covariates. If there are no systematic differences across the treatment eligibility threshold (i.e., CSR eligibility cannot be manipulated) then no discontinuity should exist for predetermined covariates. Table C2 provides the results examining the presence of a discontinuity at 150% FPL for our predetermined covariates. For all of the covariates, the point estimates are statistically insignificant and in most cases quite small indicating that there is no evidence of a discontinuity for these covariates.

Table C2: Formal Continuity Based Balance Test: Predetermined Covariates

Variable:	CER-Optimal Bandwidth:	RD Estimator:	P-Value:	Number Observations:
Age	42.197	−0.922	0.691	796
Sex - Male	35.964	0.04	0.57	693
Marital Status - Single	35.35	0.034	0.664	681
Marital Status - Married	30.966	0.064	0.397	603
Race - White	43.043	−0.062	0.427	813
Race - Black	33.023	0.01	0.868	640
Eduation - High School	38.645	0.002	0.978	742
Predicted Health Risk	33.27	−0.205	0.308	648

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table shows the results of the regression discontinuities controlling for census region and year for the predetermined covariates as a balance test. The sample is for individuals younger than 65 with insurance from the ACA marketplaces. The coverage error rate bandwidths were calculated using the Calonico *et al.* (2014) method and the kernel is uniform (as recommended in Lee and Lemieux (2010)). The results using the mean squared error bandwidths are qualitatively similar but the CER bandwidths are used on the recommendation of Cattaneo *et al.* (2019). Results are clustered at the family level.

C.4 Non-linear Effects:

An additional concern is that the chosen cutpoint is exhibiting a discontinuity due to non-linearities in the data rather than as a result of the treatment. To check that our cutpoint at 150% FPL is indeed capturing the treatment effect of the CSR subsidies, we re-estimate equation (12) using alternative cutpoints. If there are non-linearities in the data, we would expect estimates far away from 150% FPL to potentially be significant. If our results are being driven solely by the treatment, we would expect estimates away from the true cutpoint to be insignificant and close to zero.

Figures C.2 and C.3 graphically show the point estimates for healthcare expenditures and visits, respectively, for each 1% increment in the 140-160% FPL range. Figures C.2 and C.3 show a distinct U-shape with the nearly all of the statistically significant results (plotted as triangles) clustered around the true cutpoint. This indicates that the minimum is very near our treatment threshold of 150% FPL, which is what one would expect if the treatment is driving the results.

Discontinuity Plot: – Comparing Cutpoints

Figure C.2: Expenditures

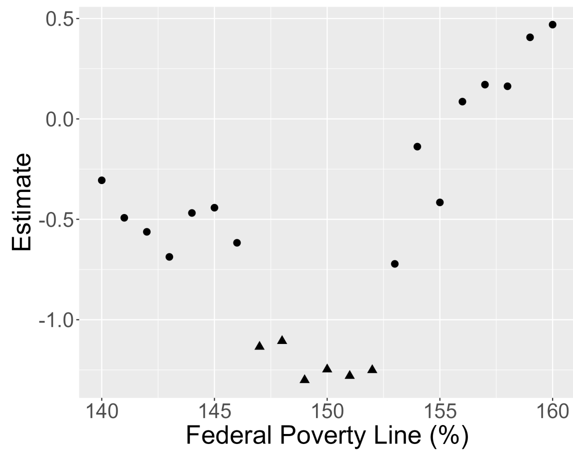

Figure C.3: Visits

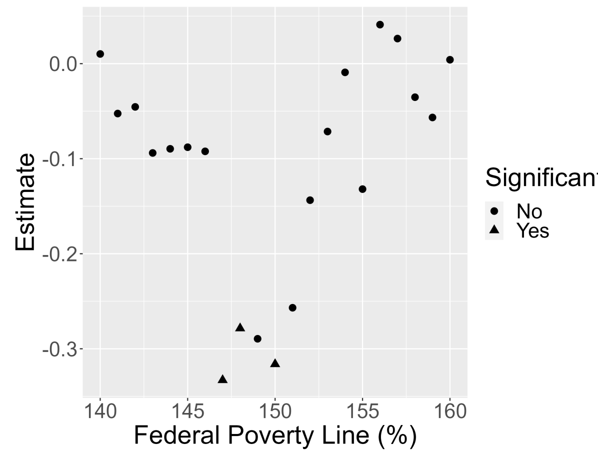

Note: This figure plots the estimate for regression discontinuities using cutpoints from 140-160% FPL of total expenditures (a) and visits (b) for consumers in the health insurance exchanges. Triangle estimates are statistically significant at the 10% level.

C.5 Manipulation of Running Variable:

Another threat for the validity of a regression discontinuity design is if individuals are able to manipulate the running variable. For our analysis, this would imply that individuals' are manipulating their income by working fewer hours, fewer jobs, etc. If this is occurring, it will violate the continuity assumption of the running variable, which is required for identification in a regression discontinuity design. To determine if this is the case, we run the McCrary (2008) sorting test, which checks for a discontinuity in the density of the data along the running variable. The density plot, shown in Figure C.4, does not display any clustering around our treatment threshold of 150% FPL. Additionally, the McCrary (2008) sorting test yields a p-value of 0.72 meaning that we fail to reject the null hypothesis that there is no sorting into treatment assignment.

Figure C.4: McCrary Sorting Test:

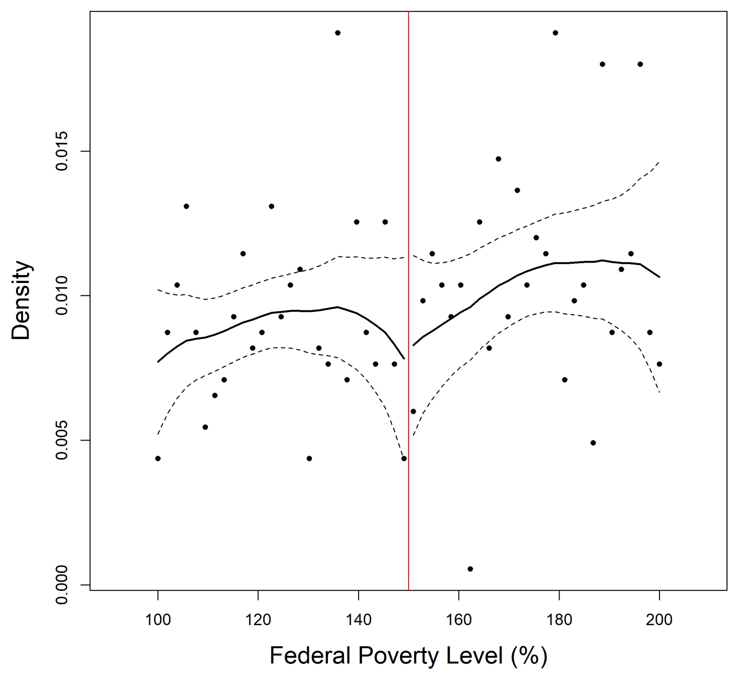

Note: This figure plots the density of the those in the non-Medicare eligible individuals who purchase health insurance through the health insurance exchanges along the running variable (income relative to the federal poverty line (FPL)) and checks for a discontinuity at 150% of the FPL.

C.6 Discussion of Adverse Selection:

To explore the third potential channel for adverse selection, we first provide intuition for why, in our research design, one should expect adverse selection and moral hazard to work in opposite directions. This result runs counter to much of the literature where moral hazard and adverse selection impact costs in the same direction (e.g., Chiappori and Salanié, 2000; Finkelstein and Poterba, 2014; Einav and Finkelstein, 2018). In table 4, we show empirically that adverse selection appears to mute the effect of moral hazard, consistent with how our empirical strategy would reflect adverse selection. However, this effect is very small and not statistically or economically significant. Our setting differs from prior literature in the groups that are compared. In Chiappori

Figure C.5: Health Status Distribution:

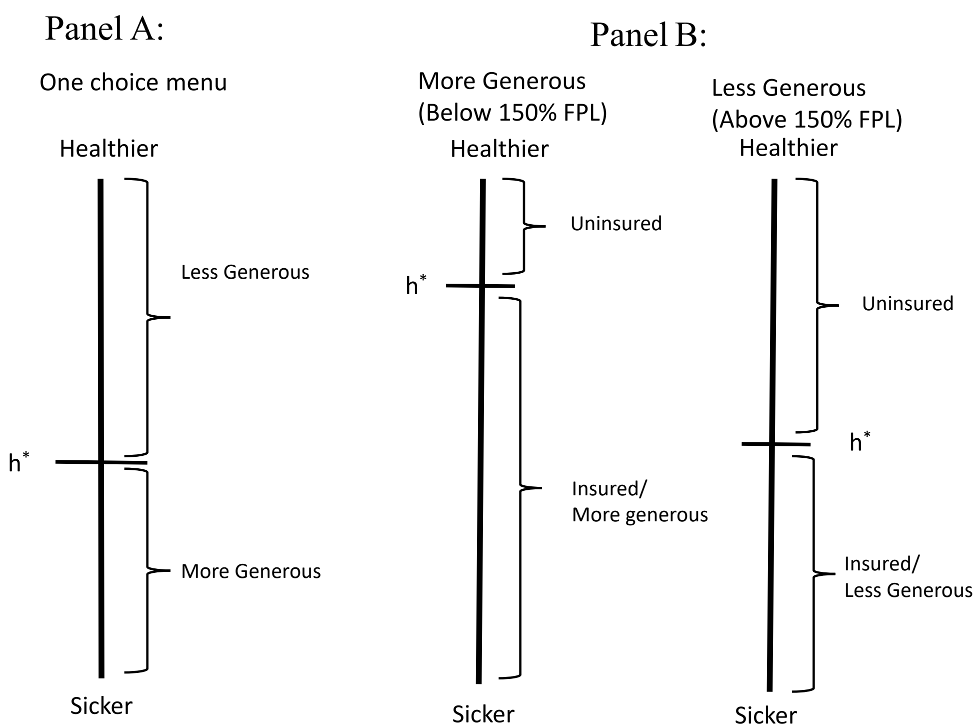

Note: This figure shows graphically how selection works in our setting. In Panel A, people sort into more or less generous insurance based on their health status. Panel B then shows how this sorting changes with plan generosity. The left figure in Panel B shows how the sorting will change when the plan becomes more generous and the right figure shows how the sorting will change when the plan becomes less generous (i.e., consumers are eligible for smaller CSR subsidies).

and Salanié (2000), and Finkelstein and Poterba (2014), the empirical strategy involves determining whether observationally similar individuals, who face the same menu of contracts, choose more or less generous insurance in a way that is correlated with their insurance usage or health status. Adverse selection implies that sicker consumers enroll in more generous insurance. Moral hazard implies that consumers will use more care with more generous insurance. Both moral hazard and adverse selection are causing higher utilization in the more generous plan.

In our setting, individuals face only a binary choice, insured vs. uninsured, where the generosity of the insurance option is assigned based on their income and premiums remain constant.[50] We then compare utilization for those who chose to be insured, across the generosity levels. More generous plans induce more of the population to purchase insurance. Because of adverse selection, the sickest people were already insured and the marginally insured are healthier. This makes the *average* enrollee in the more generous plan less costly *ex ante*, and mutes the impact of moral hazard.

First, some notation. Assume a continuum of consumers indexed by their health status h, with higher h indicating healthier consumers, and let $g(h, \alpha)$ represent consumer h's willingness to pay for a plan with generosity α. Finally, let $c(h, \alpha)$ represent a consumer h's expected cost given a level of generosity α. We first impose the setting has adverse selection, that everyone prefers more generous insurance, all else equal, and finally, moral hazard.

Assumption A5: Selection Assumptions

Assumption A5a $\frac{\partial g(h,\alpha)}{\partial h} < 0$ for all h

Assumption A5b $\frac{\partial g(h,\alpha)}{\partial \alpha} > 0$ for all h

Assumption A5c $\frac{\partial c(h,\alpha)}{\partial \alpha} > 0$ for all h

Assumption A5a states that healthier people have a lower willingness to pay for insurance (i.e., adverse selection). Assumption A5b states that all consumers have a higher willingness to pay for more generous benefits (i.e., vertical differentiation in plans). Assumption A5c is moral hazard, that more generosity increases costs holding health status fixed.

First, to show how our setting is distinct, we consider the empirical strategy which is more common in the literature, where people who choose a more or less generous plan within a menu are compared. Similar to the model in Cutler and Reber (1998), define $h^*(p, \alpha)$ as the marginal consumer who is the healthiest consumer purchasing the more generous insurance plan. From assumption A5a, all consumers who are sicker than the marginal consumer ($h < h^*$) will purchase more generous insurance and all consumers who are healthier ($h > h^*$) will purchase less generous insurance. This sorting is represented graphically in Panel A of Figure C.5.

Result 3a: The more generous plan has higher costs than the less generous plan because it has a sicker population due to adverse selection. This is immediate due to assumption A5a and the existence of a marginal consumer.

Result 3b: The more generous plan has higher costs than the less generous plan due to moral hazard. This is immediate due to assumption A5c.

Result 3c: The key result is that adverse selection and moral hazard work in the same direction.

Result 3c is a simple consequence of Result 3a and 3b, the more generous plan has higher costs due to both moral hazard and adverse selection.

Panel B of Figure C.5 highlights the difference between the more common empirical strategy and our empirical strategy. Now, consider two choice settings with different levels of generosity,

[50] There is some ability for consumers to choose different plans, however 80-90% of consumers at this income level choose a silver plan due to the generosity of the CSR subsidies (Sprung and Anderson, 2018).

but a similar distribution of h and similar premiums.[51] On the left, enrollees choose between more generous insurance and being uninsured. On the right, they choose between less generous insurance and being uninsured.

When the insurance option is more generous, the marginal person is healthier as all consumers have a higher willingness to pay for that plan, Assumption A5b. Therefore, on the left, a higher percentage of the population enroll in insurance. In this case, adverse selection implies that the more generous plan has the healthier risk pool, as the average enrollee is healthier on the left panel than the right panel. This is our Result 4a.

Result 4a: The more generous plan has lower costs than the less generous plan because it has a healthier population due to adverse selection.

Result 4b: The more generous plan has higher costs than the less generous plan due to moral hazard. This is immediate due to assumption A5c.

Result 4c: The key result is that adverse selection and moral hazard work in the opposite direction. This is a simple consequence of Result 4a and 4b.

To summarize, in our empirical strategy the presence of adverse selection will mute the effect of moral hazard.

C.7 Employer Sponsored Insurance Market Falsification Test:

Our empirical strategy of a regression discontinuity, which compares people with similar incomes, is meant to address the concern that lower-income people are generally sicker than higher-income people. In the limit, the average person at 149% FPL should have a similar health status to the average person at 151% FPL. To test this assumption empirically, Table C3 presents a falsification test using the employer-sponsored insurance (ESI) market. As with Table 4, we examine total expenditures/visits (panel A), inpatient expenditures/visits (panel B), and CSR eligible expenditures/visits (panel C). Panel (B) presents the results for the partial placebo check of inpatient expenditures and visits. The bandwidth used in Table C3 matches the bandwidth used in Table 4. If there is a discontinuity in population health status across the 150% FPL threshold, we would expect to see a discontinuity in the ESI population.

None of the estimates in Table C3 are statistically significant. Additionally, the point estimates are all are much smaller in magnitude, suggesting that the increase in cost in Table 4 is not driven by lower-income individuals being less healthy at the threshold.

[51] The premium subsidies are continuous through the 150% FPL threshold which is the focus of our empirical strategy, so holding premiums constant is reasonable in our setting.

Table C3: Falsification Test: Regression Discontinuity Models for Healthcare Expenditures and Utilization: Employer-Sponsored Insurance

	Expenditures:		Visits:	
	(1)	(2)	(3)	(4)
Total Healthcare (A):				
Over 150% FPL	−0.254	−0.203	−0.07	−0.051
	(0.168)	(0.149)	(0.054)	(0.047)
Implied \$ Change (μ)	751.924	616.059	0.357	0.265
Implied Elasticity	−0.192	−0.157	−0.058	−0.043
Placebo - Inpatient Healthcare (B):				
Over 150% FPL	0.08	0.106	−0.003	−0.003
	(0.103)	(0.096)	(0.003)	(0.003)
Implied \$ Change (μ)	−72.206	−96.901	0.000	0.000
Implied Elasticity	0.071	0.095	−0.003	−0.002
Discretionary Healthcare (C):				
Over 150% FPL	−0.261	−0.208	−0.057	−0.039
	(0.168)	(0.149)	(0.05)	(0.042)
Implied \$ Change (μ)	520.184	425.418	0.217	0.150
Implied Elasticity	−0.197	−0.161	−0.047	−0.032
Bandwidth	37.998	37.998	37.998	37.998
Number of Observations (LATE)	4,558	4,558	4,558	4,558
Year FEs?	Yes	Yes	Yes	Yes
Census Region FEs?	Yes	Yes	Yes	Yes
Demographic Controls?	Yes	Yes	Yes	Yes
Health Controls?	No	Yes	No	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. This table presents a falsification test using individuals younger than 65 with insurance from their employer. This table shows the results of the regression discontinuities for the log of total spending and visits on healthcare. Bandwidths were calculated using the Calonico *et al.* (2014) method and the kernel is uniform (as recommended in Lee and Lemieux (2010)). The bandwidths used are the optimal bandwidths from Table 4, which were calculated using the HIX market population. The additional covariates are predicted health status as well as sex, marital status, race, age, census region, year.

C.8 Oster Test for Potential Selection:

Table C4: Oster Test for Potential Selection

Outcome: *Discretionary Expenditures (Log)*						
Estimate:	Est.	S.E.	t-value	$R^2_{Y\sim D\mid\mathbf{X}}$	$RV_{q=1}$	$RV_{q=1,\alpha=0.05}$
LATE With No Selection:	-0.984	0.479	-2.054	0.6%	7.4%	0.3%
Lower Bound with Selection at Predictive Strength:						
1x ACG Score:	-0.832	0.430	-1.934	–	–	–
2x ACG Score:	-0.680	0.375	-1.814	–	–	–
4x ACG Score:	-0.375	0.226	-1.663	–	–	–

Note: This table presents the results of our Oster (2019) test for total CSR eligible expenditures excluding inpatient expenditures. The first row depicts our estimate, standard error, t-value, and R^2 values used for the test. The latter rows depict the bounds on our estimate due to various levels of potential selection. These values were calculated through Cinelli and Hazlett (2020). Note, these standard errors do not cluster by family and therefore do not match the standard errors shown in Table 4. However, the estimates here utilize the same bandwidth as Table 4.