

Price-Linked Subsidies and Health Insurance Markups*

Sonia Jaffe[†]

Mark Shepard[‡]

October 14, 2016

Abstract

Subsidies in the Affordable Care Act exchanges and other health insurance programs depend on prices set by insurers – as prices rise, so do subsidies. We study the economics of these “price-linked” subsidies compared to “fixed subsidies” set independently of market prices. We show that price-linked subsidies encourage higher prices, increasing the markups due to imperfect competition and raising subsidy costs for the government. Using a structural model estimated with administrative data from Massachusetts’ health insurance exchange, we find a non-trivial upward price distortion up to 5% of prices. We also simulate the market with only two insurers, and find distortions of 5-10%. Despite this cost, price-linked subsidies do have two advantages when the government is uncertain about health care cost. First, they insure low-income consumers against this cost risk, transferring it to the government. Second, they may indirectly link subsidies to the cost of uncompensated care, which is a key rationale for subsidizing insurance. We analyze these tradeoffs empirically under different cost shocks and find that the losses from higher prices outweigh the uncertainty benefits of price-linked subsidies for reasonable levels of cost uncertainty. We discuss an alternate policy that would eliminate the pricing distortion while maintaining the guaranteed affordability property of price-linking.

*The authors would like to thank Amitabh Chandra, David Cutler, Leemore Dafny, Keith Ericson, Jerry Green, Jon Gruber, Kate Ho, Scott Kominers, Robin Lee, Ariel Pakes, Amanda Starc, Pietro Tebaldi and seminar and conference participants at Boston University, Harvard economics department, Harvard Medical School, the University of Chicago economics department, and ASHEcon 2016 for their helpful comments. They thank the Lab for Economic Applications and Policy (LEAP) at Harvard University for funds for acquiring the data, and the Massachusetts Health Connector and its employees (especially Nicole Waickman and Marissa Woltmann) for access to and explanation of the data. The views expressed herein are our own and do not reflect those of the Connector or its affiliates. All mistakes are our own.

[†]Becker Friedman Institute, University of Chicago, spj@uchicago.edu. Jaffe gratefully acknowledges the support of a National Science Foundation Graduate Research Fellowship, as well as the hospitality of the Becker Friedman Institute for Research in Economics at the University of Chicago.

[‡]John F. Kennedy School of Government, Harvard University and NBER, mark_shepard@hks.harvard.edu Shepard gratefully acknowledges the support of National Institute on Aging Grant No. T32-AG000186 via the National Bureau of Economic Research, and a National Science Foundation Graduate Research Fellowship.

Public health insurance programs increasingly cover enrollees through regulated markets offering a choice among subsidized private plans. Long used in Medicare’s private insurance option (Medicare Advantage), this model was adopted for the Medicare drug program (Part D) in 2006 and more recently, in the insurance exchanges for low-income individuals created by the Affordable Care Act (ACA) in 2014. These programs aim to leverage the benefits of choice and competition, while using subsidies to make insurance more affordable and encourage enrollee participation. While a large body of work studies the role of subsidies in expanding coverage,¹ there is also growing recognition that the design of subsidies may affect insurers’ pricing incentives, potentially affecting the benefits of competition.²

We study the implications of a key subsidy design choice that arises in a variety of settings: whether to *link* subsidies to prices set by insurers. Many public programs take this “price-linked” subsidy approach. For instance, Medicare Part D links subsidies to market average prices, and the ACA exchanges link subsidies to the second-cheapest plan in the “silver” tier. Other programs, however, set subsidies at specific levels or based on external benchmarks not controlled by insurers. For instance, Medicare Advantage sets subsidies largely based on an area’s lagged costs in traditional Medicare. While linking subsidies to prices is convenient when there is uncertainty about market prices, it also raises competitive concerns. In the extreme case, with a monopoly insurer (or if insurers collude), price-linked subsidies could facilitate large price increases without any loss of demand, since subsidies would increase in tandem and consumer prices would remain unchanged.

We argue that economists should think of price-linked subsidies as involving a basic tradeoff. On the one hand, they create a competitive distortion, increasing consumers’ or government’s cost by increasing prices compared to “fixed subsidies” set independently of prices. On the other hand, price-linked subsidies create an indirect link between subsidies and cost shocks, which can be desirable in the face of uncertainty about health care costs. We use a simple theoretical model to show the intuition for these effects and to formalize the conditions under which they occur. We then analyze these tradeoffs empirically using a structural model estimated with administrative data from Massachusetts’s (ACA-like) insurance exchange for low income uninsured.

Our first contribution is to use a simple model of equilibrium pricing in an insurance exchange to analyze the competitive implications of price-linked subsidies. If a higher price yields a higher subsidy for the firm *relative* to other options in the market, the firm has an

¹See Gruber (2008) for a review of the literature on the rationales for and effects of subsidies and mandate penalties. Two prominent rationales for subsidies are adverse selection (Einav et al., 2010; Hackmann et al., 2015; Bundorf et al., 2012) and the cost of uncompensated care incurred by the uninsured (Mahoney, 2015).

²Recent work on this topic includes Decarolis et al. (2015) studying Medicare Part D, Curto et al. (2014) studying Medicare Advantage, and Tebaldi (2016) studying the California ACA exchange.

incentive to raise its price. Even in markets where subsidies are constant across plans,³ as in the the ACA, the subsidy usually does not apply to the “outside option” of not purchasing insurance. A higher subsidy decreases the cost of buying a market plan relative to not buying insurance. Each firm gains some of the consumers brought into the market by the higher subsidy, so each firm has an incentive to raise the price of any plan it thinks may affect the subsidy.

We show that price-linked subsidies can distort prices whenever markets are imperfectly competitive and the subsidy affects firms’ prices relative to a viable outside option. Since price-linked subsidies in effect remove the competition from the outside option, the distortion is larger when firms have more market power – a smaller own-price elasticity of demand – or compete more with the outside option – a larger elasticity of demand with respect to the price of the outside option. The average federally facilitated ACA exchange had just 3.9 insurers in 2014 (Dafny et al., 2015) and approximately 30% of population-weighted markets had just one or two insurers. Higher subsidies also bring additional federal dollars into a state (since subsidies are paid by the federal government), so the state governments that run many of the exchanges could have an incentive to coordinate pricing (collusion) by insurers, further decreasing competition.⁴ In both Massachusetts and the ACA context, the subsidies do not affect the mandate penalty for uninsurance, which is the main outside option for exchange consumers since people with employer coverage or Medicare are not subsidy-eligible.⁵ About half of the subsidy-eligible population in Massachusetts remained uninsured instead of purchasing insurance through the exchange, so it is not surprising that we find substantial substitution to the outside option.

Our analysis suggests a simple alternate policy that would remove the distortion: link the mandate penalty to prices in the *opposite* direction. Specifically, if the price of the pivotal plan exceeds an expected level, the difference would be applied to reduce the mandate penalty (and vice versa if its price were less than expected). Unlike fixed subsidies, this policy ensures the affordability of the pivotal (second-cheapest silver) plan. Unlike existing price-linked subsidies, this policy does not distort competition because the *total* incentive to buy insurance (the subsidy plus mandate penalty) is unaffected by prices. However, if there

³As noted in Enthoven (1988), there are additional incentive issues if subsidies vary across plans and a plan’s price directly affects its own subsidy (e.g. an employer subsidizes a 75% of the price of each plan, meaning that each \$1 price increase raises a plan’s subsidy by \$0.75).

⁴This effect is similar to the “fiscal shenanigans” concerns that have been documented for federal matching funds in Medicaid (Baicker and Staiger, 2005).

⁵Non-group insurance purchased outside of exchanges is likely to be dominated for low-income people relative to the heavily subsidized exchange plans. An interesting possibility that we do not model is that employer insurance could be an outside option through employers’ strategic decisions to drop coverage and let employees join exchanges. To the extent this were an important factor, our alternate policy (discussed below) would not fully correct the pricing distortion.

is substitution on other margins, such as changing jobs to an employer that offers insurance (Aizawa, 2016), this will not completely eliminate the distortion.

Our second main contribution is to analyze the welfare tradeoffs of price-linked subsidies relative to fixed subsidies (or our alternate policy). We base our welfare framework on a consumer surplus standard (commonly used in antitrust analysis), modified to include government costs and an externality of uninsurance – an important part of the rationale for subsidizing insurance.⁶ Higher prices generate a first-order loss in the higher costs to the government and consumers. The benefits of price-linking arise only when the regulator faces *uncertainty* about health care costs and market outcomes. Absent uncertainty, or if insurers face the same uncertainty as the government, the regulator can predict market prices and therefore the price-linked subsidy; it can replicate the same subsidy amount with a fixed subsidy but eliminate the pricing distortion, creating a pure gain for consumers.

Uncertainty creates two reasons price-linked subsidies may be desirable. First, price-linking lets the government insure low-income enrollees against the risk that prices grow faster than expected. This is closely related to an explicit rationale the ACA’s subsidy design: guaranteeing that post-subsidy premiums are “affordable” regardless of insurer prices. Second, a large component of the externality the uninsured impose on society is the the cost of “uncompensated care” – typically borne by hospitals and public clinics.⁷ The regulator wants the subsidy to track this associated externality; price-linking provides an indirect way of achieving this, since both uncompensated care costs and insurance prices generally rise with health care costs. Both of these potential benefits of price-linked subsidies require that the regulator face more uncertainty when setting the subsidy policy than insurers face when setting prices. For price-linking to be optimal, the regulator’s relative uncertainty must be large enough to outweigh the cost of the price distortion. The extent to which the insurers are better informed than the government is an open question,⁸ so in our empirical simulations, we consider a range of cost shocks to gauge the sensitivity of our conclusions.

Our paper’s third main contribution is to study the pricing distortion and welfare tradeoffs empirically using a structural model of insurance competition. To estimate the model, we use

⁶We choose not to use a social surplus standard because we think excluding insurer profits is a more realistic description of policymakers’ objective function in this setting. A social surplus standard would (to first-order) imply no social cost of the price distortion. The higher markups would simply be a transfer from the government to insurers. Also, with positive markups, a social surplus standard calls for subsidies because value is greater than cost, independent of any concerns specific to health insurance.

⁷There is growing evidence on the importance of uncompensated care for low-income people (Garthwaite et al., 2015; Finkelstein et al., 2015). Indeed, Mahoney (2015) proposes including these costs (which he connects with the threat of bankruptcy) as a rationale for the mandate penalty.

⁸States typically require insurers to submit cost reports, giving regulators substantial past cost information. Regulators can also draw on data from other markets (e.g., local Medicare costs) and cost predictions by professional actuaries. Nonetheless, insurers may still possess superior signals about future cost trends.

administrative plan enrollment and claims data from Massachusetts’s pre-ACA subsidized insurance exchange. We supplement this with data on the uninsured from the American Community Survey, in order to estimate the share of people who choose the outside option. An important model for the ACA, this market lets us observe insurance demand and costs for a similar low-income population.

We use standard demand estimation methods for micro-data, similar to Berry et al. (2004). Our model allows for adverse selection by letting both demand and cost parameters vary with demographics. An important strength of our method is that we use quasi-experimental variation to identify the key statistic of the responsiveness of demand to the price of the outside option.⁹ Two natural experiments in Massachusetts give us exogenous variation. The first is the introduction of the mandate penalty in December 2007.¹⁰ The second is a subsidy change in July 2007 that lowered all plans’ premiums for a single income group.¹¹ In both cases, we use income groups excluded from the change as control groups. The results are similar: each \$1 increase in the relative price of uninsurance raised insurance demand by 1%. We therefore match this statistic as a moment in our demand estimation.

The estimated demand and cost parameters allow us to calculate insurers’ profit functions, which we use to simulate Nash pricing equilibria under different subsidy policies and cost shocks.¹² Comparing the cheapest plan’s price between simulations with price-linked and fixed subsidies (set equal to the equilibrium price-linked subsidies) provides an estimate of the pricing distortion. Depending on the year and simulation specification, we estimate a distortion of \$3-\$22 in monthly prices, or 1-5% of plan prices. The estimate \$21 subsidy distortion in 2011 translates to \$42 million in annual subsidy costs for Massachusetts (\$5.2 billion if extrapolated nationally to the ACA). Instead of setting the fixed subsidy equal to the equilibrium price-linked subsidy, we can calculate the government costs and insurance rates for various subsidy levels. In 2011, with fixed subsidies, the government could have spent the same amount and gotten 2.72 percentage points more health insurance coverage or spent 5.33% less and gotten the same insurance coverage.

⁹We match this moment exactly by including a random coefficient on individuals’ utility of insurance (relative to uninsurance). The variance of this random coefficient determines substitution patterns between insurance and uninsurance, including the semi-elasticity moment we match.

¹⁰Chandra et al. (2011) study the same change and show that it led to an increase in exchange enrollment by relatively healthy individuals. Our identification strategy is similar to theirs, but we focus on estimating a demand semi-elasticity from the change.

¹¹Because a lower price of all plans affects relative prices the same as an increase in the mandate penalty, we can use these two changes to estimate the same statistic.

¹²We mostly follow Massachusetts’ subsidy rules in our simulations, which differed somewhat from the ACA. There was a single coverage tier (not bronze/silver/gold) and subsidies were linked to the cheapest plan. However, we exclude enrollees below 100% of poverty who are covered by Medicaid in the ACA and whose price sensitivity we cannot identify because all plans were free for them in Massachusetts.

Since about many ACA markets have only one or two insurers, we also simulate the market under the two subsidy regimes with only two insurers. (With one insurer, absent other controls, the subsidy is a money pump and there is no equilibrium.) Distortions are larger with only two insurers, particularly when there is a big gap in the costs of the two insurers. The price of the cheapest plan is \$49 higher under price-linked subsidies when its only competition is a high cost plan.

Returning to the full market with four to five insurers, we next simulate the performance of these policies under market-level cost uncertainty. We assume that regulators cannot observe the cost shock in time to adjust the fixed subsidy, but insurers observe the cost shock and price accordingly, leading the price-linked subsidy to vary with costs. We find that price-linked subsidies stabilize the insurance coverage rate, but the gain is small relative to the loss from the price distortion. The stability of coverage rates is an important point for predicting future coverage under the ACA: even if costs grow faster than expected (or selection is worse than expected), the adjustment of price-linked subsidies will stabilize coverage. However, even with the assumption about uncompensated care cost that is most favorable for price-linked subsidies, fixed subsidies do better unless costs differ from expectations by 12.5%. Such a shock seems improbably large given the 1.9 percentage point standard deviation in state-level annual cost growth from 1991-2009. If the government had substantial uncertainty about which customers would participate in a new market, and therefore could not estimate costs well, price-linked subsidies might do better for the first year or two, but in an established market, price-linking is only likely to maximize public surplus if the government observes cost with a many-year lag or if political economy or complexity constraints prevent adjusting fixed subsidies over time based on updated cost data.

As we discuss in Section 5, the pricing distortion we identify is relevant in programs besides the ACA that use price-linked subsidies and have a meaningful outside option. Our paper is related to a small but growing literature studying the (often unintended) competitive implications of subsidy policies – including Decarolis (2015) and Decarolis et al. (2015) studying Medicare Part D, and Liu and Jin (2015) studying the Federal Employees Health Benefits Program. We are (to our knowledge) the first to formalize and analyze the tradeoffs involved with linking subsidies to prices. The most closely related paper is recent work by Tebaldi (2016), who shows how to estimate a model of market equilibrium without cost data for the first year of California’s ACA exchange. Tebaldi compares the ACA’s price-linked subsidies to fixed subsidies (or “vouchers”) in one of his counterfactuals, finding a price distortion that is remarkably similar to our estimates from Massachusetts. Our paper differs both in its focus on the conceptual tradeoffs of price-linked subsidies (rather than Tebaldi’s focus on the methodological approach and the benefits of varying subsidies with age) and in

considering their performance under cost uncertainty.

The remainder of the paper is structured as follows. Section 1 uses a simple model to show our theoretical analysis of price-linked subsidies and explain our welfare framework. Section 2 describes our Massachusetts setting and data and presents evidence from two natural experiments in this market. Section 3 describes the structural model, the estimation strategy, and framework for the counterfactuals. Section 4 reports the parameter estimates and shows the counterfactual simulation results. Section 5 discusses the tradeoffs among subsidy structures and the implications of our findings for the ACA and other insurance markets. Section 6 concludes.

1 Theory

We adapt a standard discrete choice model of demand to allow for a mandate penalty and various subsidy policies. The conditions for firm profit maximization show the basic mechanism through which the subsidy structure affects prices and give a first-order approximation for the price distortion. We focus on the case relevant to our Massachusetts exchange setting, where each insurer offers a single plan. The multi-plan insurer case – which is relevant to the ACA setting – is conceptually quite similar. We discuss how the magnitude of the pricing distortion might differ in Section 5, and work out the formula in the Appendix A.

Insurers $j = 1, \dots, J$ offer differentiated products and compete by setting prices $P = \{P_j\}_{j=1\dots J}$. The exchange collects these price bids and uses a pre-specified formula to determine a subsidy $S(P)$ that applies equally to all plans.¹³ Subsidy-eligible consumers then choose which (if any) plan to purchase based on plan attributes and post-subsidy prices, $P_j^{cons} = P_j - S(P)$. Consumers can also choose an outside option of uninsurance. If they do so, they are subject to the legally applicable mandate penalty, $M(P)$, which could also depend on prices. Total demand for plan j , $Q_j(P^{cons}, M)$, is a function of all premiums and the mandate penalty.

We assume that insurers set prices simultaneously to maximize static profits, knowing the effects of these choices on demand and cost.¹⁴ As in most recent work on insurance (e.g. Einav et al., 2010; Handel et al., 2015), we assume fixed plan attributes and focus instead on pricing incentives conditional on plan design.¹⁵

¹³The subsidy and mandate penalty may differ across consumers based on their incomes or other characteristics. For ease of exposition, we do not show this case here but allow for income-specific subsidies in our empirical model.

¹⁴We therefore abstract from pricing dynamics or incomplete information. We discuss below how incorporating these into the model might affect the pricing distortion.

¹⁵Fixed attributes seem reasonable as a first approximation in regulated insurance markets like the Massachusetts or ACA exchanges, where plan attributes like cost sharing and covered services are fixed or heavily

Insurer costs are a function of both plan attributes and the enrollees who select its plan. In addition, the exchange uses risk adjustment to compensate plans based on the measured sickness of its enrollees. We allow for both selection and risk adjustment by letting a plan's average costs, $\bar{c}_j(P^{cons}, M)$, and risk adjustment transfer, $\bar{\phi}_j(P^{cons}, M)$, depend on prices (which affect consumer selection). The insurer's net (risk-adjusted) costs equal:

$$c_j^N(P) = \bar{c}_j(P^{cons}, M) - \bar{\phi}_j(P^{cons}, M).$$

If there are constant marginal costs and no adverse selection, then $c_j^N(P)$ is just a fixed c_j and $\frac{\partial c_j^N}{\partial P} = 0$.

The insurer profit function is:

$$\pi_j = (P_j - c_j^N) \cdot Q_j(P^{cons}, M).$$

A necessary condition for Nash equilibrium is that each firm's first-order condition is satisfied:¹⁶

$$\frac{d\pi_j}{dP_j} = \left(1 - \frac{\partial c_j^N}{\partial P_j}\right) Q_j(P^{cons}, M) + (P_j - c_j^N) \cdot \frac{dQ_j}{dP_j} = 0. \quad (1)$$

This differs from standard oligopoly pricing conditions in two respects. First, the $\frac{\partial c_j^N}{\partial P}$ term allows for selection to influence pricing, a standard consideration in insurance markets. Second, because of subsidies, the firm's price P_j enters consumer demand indirectly, through the subsidized premiums, $P^{cons} = P - S(P)$. As a result, the term dQ_j/dP_j (a total derivative) is not the slope of the demand curve, but a composite term that combines the slope of demand and any indirect effects on demand if P_j affects the subsidy or mandate penalty (via the regulatory formula). The total effect of raising P_j on demand is a combination of three effects:

$$\frac{dQ_j}{dP_j} = \underbrace{\frac{\partial Q_j}{\partial P_j^{cons}}}_{\text{Direct}} - \underbrace{\left(\sum_k \frac{\partial Q_j}{\partial P_k^{cons}}\right) \frac{\partial S}{\partial P_j}}_{\text{via Subsidy}} + \underbrace{\frac{\partial Q_j}{\partial M} \frac{\partial M}{\partial P_j}}_{\text{via Mandate}}. \quad (2)$$

The first term is the standard demand slope with respect to the consumer premium. The next two terms are the indirect effects via the subsidy (which lowers all plans' consumer constrained by regulation.

¹⁶These first-order conditions would be necessary conditions for Nash equilibrium even in a more complicated model in which insurers simultaneously chose a set of non-price characteristics like copays and provider network. Thus, the theoretical point we make about price-linked subsidies holds when quality is endogenous, though there may also be effects on quality and cost levels, which we do not capture.

premiums) and mandate penalty.

We can simplify this formula by imposing an assumption that is standard in most discrete choice models: that price enters the utility function linearly. This implies that (at least locally) only *price differences*, not levels, matter for demand.¹⁷ Thus, raising all prices (and the mandate penalty) by \$1 leaves demand unchanged: $\sum_k \frac{\partial Q_j}{\partial P_k^{cons}} + \frac{\partial Q_j}{\partial M} = 0 \forall j$. Using this condition to simplify Equation (2), we get:

$$\frac{dQ_j}{dP_j} = \underbrace{\frac{\partial Q_j}{\partial P_j^{cons}}}_{(-)} + \underbrace{\frac{\partial Q_j}{\partial M}}_{(+)} \cdot \underbrace{\left(\frac{\partial S}{\partial P_j} + \frac{\partial M}{\partial P_j} \right)}_{\text{Price-Linking}}. \quad (3)$$

The effective demand slope (for a firm’s pricing equation) equals the slope of the demand curve, plus an adjustment arising if policy creates a link between S or M and prices.

Intuitively, the adjustment effect depends on the magnitude of $\partial Q_j/\partial M$ because neither S nor M affect price differences among in-market plans but they both affect the price of all plans relative to the outside option. Since relative prices are what drive demand, the effect of S and M depends on how sensitive Q_j is to the relative price of the outside option. If there is no outside option (e.g., if the mandate penalty led to 100% coverage) or if few people buy insurance when M increases, the effect will be small. But if substitution is high, the effect will be large. Thus, a key goal of our empirical work is to measure $\partial Q_j/\partial M$.

Since the effect of the mandate penalty on demand for j is positive (understand standard assumptions), the adjustment effect will be positive if the subsidy and/or mandate penalty rise with prices (as occurs in the many subsidized insurance markets). This diminishes the (negative) slope of the demand curve, making *effective* demand less elastic, which increases equilibrium markups. Conversely, if the subsidy or penalty fell with prices, the policy would make effective demand *more* elastic, lowering markups. This is an interesting possibility but not one that we have seen implemented in any actual insurance programs.

We use this framework to analyze the impacts of different subsidy schemes. We start in Section 1.1 by showing the implications for price competition and markups. In Section 1.2, we introduce an objective function of the public regulator to analyze the tradeoffs involved between price-linked and fixed subsidies under uncertainty.

¹⁷This assumption is typically justified by the fact that prices are a small share of individuals’ incomes. Although we study a low-income population, subsidized premiums are also quite low. Average premiums in Massachusetts range from < 1% of income for lower income groups (below 150% of poverty) to about 5% of income for the highest income groups. In an insurance setting, linear-in-price utility can be seen as a transformed approximation to a CARA utility function, in which risk aversion is not affected by income effects.

1.1 Markups under Different Subsidy Policies

Fixed Subsidies

One policy option is for regulators to set the subsidy and mandate penalty based only on “exogenous” factors not controlled by market actors. We call this policy scheme “fixed subsidies” to emphasize that they are fixed relative to prices; however, subsidies may adjust over time and across markets based on exogenous factors (e.g., local costs in Medicare), as in the yardstick competition model of Shleifer (1985). Under fixed subsidies:

$$\frac{\partial S}{\partial P_j} = \frac{\partial M}{\partial P_j} = 0 \quad \forall j.$$

Since subsidies and the mandate penalty are unaffected by any plan’s price, dQ_j/dP_j in Equation (3) simplifies to the demand slope $\partial Q_j/\partial P_j^{cons}$. Even though there are subsidies, the equilibrium pricing conditions are not altered relative to the standard form for differentiated product competition. Of course, the subsidy and mandate may increase insurance demand – an outward shift in the demand curve – but our point is that they do not *rotate* the demand curve, as occurs with price-linking. Under this fixed subsidies benchmark, markups are:

$$Markup_j^F \equiv P_j - c_j = \frac{1}{\eta_j} \left(1 - \frac{\partial c^N}{\partial P_j} \right) \quad \forall j,$$

where $\eta_j \equiv -\frac{1}{Q_j} \frac{\partial Q_j}{\partial P_j^{cons}}$ is the own-price semi-elasticity of demand.

Price-Linked Subsidies

Alternatively, exchanges could *link subsidies to prices* (but again set a fixed M). This is the approach used in most public insurance markets, including Massachusetts, the ACA, and Medicare Part D. We focus on the type of subsidy used in Massachusetts and the ACA, where subsidies are linked to an order statistic on prices – the cheapest price (Massachusetts) or the second-cheapest silver plan price (ACA). This setup ensures that this plan’s post-subsidy premium equals an (income-specific) “affordable amount,” regardless of its pre-subsidy price. This can be good when prices are uncertain (as we discuss further below), but it also lessens price competition. For instance, under the Massachusetts design:

$$\begin{aligned} S(P) &= \min_j P_j - \text{AffAmt}, \quad M(P) = M \\ \implies \frac{\partial S(P)}{\partial P_j} &= 1, \quad \frac{\partial M}{\partial P_j} = 0 \quad \forall j, \end{aligned}$$

where \underline{j} is the index of the pivotal (cheapest) plan.

In equilibrium, this design makes a single plan completely pivotal for the subsidy ($\partial S/\partial P_{\underline{j}} = 1$), as long as its price does not increase to the level of the next-highest plan price. Of course, out-of-equilibrium, it may not be obvious which plan will be pivotal. We focus primarily on the equilibrium case, since it captures the main intuition, and discuss below how the result might differ with uncertainty about the identity of \underline{j} .

Effective demand for the pivotal plan is less elastic: $dQ_{\underline{j}}/dP_{\underline{j}} = \partial Q_{\underline{j}}/\partial P_{\underline{j}} + \partial Q_{\underline{j}}/\partial M$. Plugging this into Equation (1) and rearranging yields the following markup condition for the pivotal plan under price-linked subsidies:

$$Markup_{\underline{j}}^{\text{PLink}} \equiv P_{\underline{j}} - c_{\underline{j}} = \frac{1}{\eta_{\underline{j}} - \eta_{\underline{j},M}} \left(1 - \frac{\partial c^N}{\partial P_{\underline{j}}} \right),$$

where $\eta_{\underline{j},M} = \frac{1}{Q_{\underline{j}}} \frac{\partial Q_{\underline{j}}}{\partial M}$ is the semi-elasticity of demand for \underline{j} with respect to the mandate penalty.

Comparing Fixed and Price-Linked Subsidies

In Massachusetts price-linked subsidies lower the effective price sensitivity faced by the cheapest (pivotal) plan, leading to a higher equilibrium markup than under fixed subsidies. Though if the distortion is large enough that the cheapest plan would want to price above the second-cheapest plan, it instead sets a price equal to the second-cheapest plan.¹⁸

Like much of the related literature, we make the simplifying assumption that firms know what equilibrium they are in, so there is no uncertainty about which plan will be cheapest. In this case, the distortion only applies to the pivotal plan, though there may be strategic responses by other firms. In a model with uncertainty about others' prices (perhaps due to uncertainty about others' costs) then the distortionary term $\eta_{\underline{j},M}$ would be weighted by the probability of being the lowest price plan. The (ex-post) cheapest plan would have a smaller distortion, but there would also be effects on other plans' prices.

If the semi-elasticities of demand are constant across the relevant range of prices (equivalently, if own-cost pass-through equals one and cross pass-through is zero), we can derive an explicit expression for the increase in markups between fixed and price-linked subsidies:

$$Markup_{\underline{j}}^{\text{PLink}} - Markup_{\underline{j}}^{\text{F}} = \frac{\eta_{\underline{j},M}}{\eta_{\underline{j}} (\eta_{\underline{j}} - \eta_{\underline{j},M})} \left(1 - \frac{\partial c^N}{\partial P_{\underline{j}}} \right) > 0, \quad (4)$$

¹⁸This can create a range of possible equilibria with a tie among multiple cheapest plans, an issue we address in our simulations.

which is positive.¹⁹ as long as $\frac{\partial c^N}{\partial P_j} < 1$ – meaning the adverse selection not captured by risk adjustment is not really terrible.²⁰ Alternatively, if semi-elasticities are not constant, this expression can be thought of as an estimate of how much marginal costs would have had to decrease to offset the incentive distortion generated by price-linked subsidies.²¹

Alternate Policy: Price-Linked Subsidies and Mandate Penalty

Our model suggests a simple alternative to the standard price-linked subsidy design that would preserve the guaranteed affordability while eliminating the price distortion. Specifically, regulators could set a base mandate penalty M_0 and then *apply the subsidy to the mandate penalty* so that:

$$M(P) = M_0 - S(P).$$

In this case, the second term in Equation (3) equals zero, so the demand slope is unaffected by the price-linking: $dQ_j/dP_j = \partial Q_j/\partial P_j^{cons} \forall j$. The government could set M_0 so that the expected mandate penalty equaled the penalty under the current system, but the actual mandate penalty would depend on market prices.

Intuitively, this works because the *net* public incentive for consumers to buy insurance is $S + M$. Fixed subsidies hold both S and M fixed, regardless of insurer pricing. Standard price-linked subsidies link S to prices but leave M fixed, causing the pivotal plan to affect the net incentive $S + M$. Our alternate policy links S and M to prices in opposite directions, generating a constant $S + M$. Since plans do not care about the levels of S and M independent of their total, which is fixed, there is no incentive distortion.

1.2 Welfare Analysis

To assess the tradeoff between the price distortion and the ability of price-linked subsidies to ensure affordable prices for consumers and match the subsidy to the externality when costs are uncertain, we need to specify a regulatory objective function.

¹⁹Both $\eta_{j,M} = \frac{1}{Q_j} \frac{\partial Q_j}{\partial M}$ and $\eta_j = -\frac{1}{Q_j} \frac{\partial Q_j}{\partial P_j^{cons}}$ are positive under standard demand assumptions, and $\eta_j > \eta_{j,M}$ as long as there is at least one other in-market option besides j . (With no other plans, this subsidy rule creates a money pump, which absent other regulation, incentivizes an infinite markup.) Note that this expression only applies as long as the higher markup does not push the pivotal plan's price above the next-cheapest.

²⁰We think it is reasonable to assume that even if marginal consumers are more expensive, the average cost across all consumers, c^N , will not increase one-for-one with price.

²¹This is similar to the idea from Werden (1996) that, without assumptions about elasticities away from the equilibrium, one can calculate the marginal cost efficiencies needed to offset the price-increase incentives of a merger.

We use a welfare standard based on consumer surplus, choosing to exclude insurer profits, for two reasons. First, the public debate around health care reform centered around maximizing consumer benefit while minimizing government costs, with little weight put on insurer profits.²² Second, with imperfect competition, maximizing social welfare (including profits) generates higher subsidies for products with higher markups, simply because the social benefit exceeds the marginal cost. To avoid having this effect drive our results and to focus on issues specific to the exchange, we exclude profits from the regulator objective. In this case, it is socially optimal for a consumer to purchase insurance if its (social) benefit exceeds its price, just as in models of perfect competition, but here prices depend on the market and subsidy structure.

We also want the objective function to provide a rationale for subsidizing health insurance (as opposed to redistributing funds via the transfer system).²³ The two primary reasons for subsidizing health insurance are adverse selection – which pushes up insurance prices and leads some people to be inefficiently uninsured – and the externality of uninsurance – the social cost. Part of this externality is based on the uninsured’s use of “uncompensated care,” whose importance has been highlighted in recent work (Mahoney, 2015; Garthwaite et al., 2015; Finkelstein et al., 2015). These costs are borne by hospitals and public clinics, who are unable or unwilling to deny needed care. In order to match the level of subsidies in Massachusetts and under the ACA, we also allow for a pure paternalistic social disutility of people lacking insurance.

Let $U(P - S, M)$ be total consumer surplus, which is a function of consumer prices ($= P - S$) and the mandate penalty. Normalize the population size to 1, and let $D_j(P - S, M)$ be demand (market share) for plan j , with $D_0(P - S, M)$ being the share uninsured. Net government spending equals subsidies for the insured minus mandate penalty revenue:

$$G(P, S, M) = (1 - D_0(P - S, M)) \cdot S - D_0(P - S, M) \cdot M$$

Finally, define the average externality avoided per insured consumer as $\bar{E}(P - S, M, C)$, which can vary with consumer prices (which determine who is uninsured). The regulator seeks to maximize “*public surplus*”:

$$PS = \underbrace{U(P - S, M)}_{\text{Cons. Surplus}} - \underbrace{G(P, S, M)}_{\text{Govt. Cost}} + \underbrace{\bar{E}(P - S, M, C) \cdot (1 - D_0(P - S, M))}_{\text{Avoided Externality}}.$$

²²Indeed, policymakers seem eager to constrain insurer profits with policies like medical loss ratio limits.

²³There is a large literature about rationales for in-kind benefits, see Currie and Gahvari (2008) for a review.

In our structural model simulations, we use this public surplus measure to numerically analyze the tradeoffs involved between price-linked and fixed subsidies. Showing this analysis analytically is in general quite complicated, so in this section, we make some simplifying assumptions to allow us to illustrate the basic intuition.

Baseline: No Adverse Selection, Cost Uncertainty or Risk Aversion

Suppose the regulator has full information about costs, demand, and the externality of uninsurance and there is no adverse selection. Suppose also that utility is (locally) quasi-linear so that the subsidy and penalty net out as pure transfers between consumers and the government. In this case, there is little justification for price-linked subsidies. To see this, compare (1) the optimal price-linked subsidy policy with an optimal mandate penalty M^* and resulting equilibrium subsidy S^E and (2) a (generally sub-optimal) fixed subsidy policy that uses the same mandate penalty and fixes the subsidy at that equilibrium level, S^E . The different subsidy policies will result in different prices and consumer choices. To a first-order approximation, the change in public surplus equals:

$$PS^{\text{PLink}} - PS^{\text{F}} \approx \underbrace{-\sum_{j \neq 0} \Delta P_j D_j}_{\Delta \text{Cons. Surplus}} - \underbrace{\Delta D_0 \cdot (\bar{E} - (S^E + M^*))}_{\Delta \text{Govt. Cost and Externality}} \quad (5)$$

where $\Delta P_j = P_j^{\text{PLink}} - P_j^{\text{F}}$ and similarly for ΔD_0 .

The first term in Equation (5) is the change in consumer surplus. Since the policies have equal subsidies, the change in consumer surplus is approximately the price changes times baseline quantities (with changes in quantities being second-order). The pricing distortion means we expect ΔP_j to be positive. Assuming strategic complementarity (or strategic substitutability that is not too large), the average price will also fall. Thus, we expect the difference in consumer surplus will be negative.

The second term in Equation (5) is the change government costs and the externality. Lower prices and equal subsidies mean that fixed subsidies will have fewer uninsured ($\Delta D_0 > 0$). For intuition about $\bar{E} - (S^E + M^*)$, think about the regulator choosing the optimal mandate penalty (and knowing the resulting subsidy). If prices did not change with the mandate penalty, the regulator would set $M = \bar{E} - S^E(P)$ so the net subsidy, $M + S$ equals the externality. With price complementarity, prices increase with the mandate penalty, which is a reason to have a lower mandate penalty so $M^* + S^E < \bar{E}$. Therefore, $\Delta D_0 \cdot (\bar{E} - (S^E + M^*)) > 0$, fixed subsidies have lower government costs and externality than price-linked subsidies.

In Appendix B, we formalize this heuristic argument that absent uncertainty and adverse

selection, it is optimal for the subsidy to not increase with price for the case of logit demand (which gives price complementarity).

Adverse Selection

Adverse selection adds the difference in the average saved externality of the insured, $\Delta \bar{E} \cdot (1 - D_0)$, to the difference in public surplus. Under standard adverse selection the additional consumers drawn into the market by lower prices under fixed subsidies will be less sick than those already in the market, lowering the average saved externality of the insured. Therefore, $\Delta \bar{E} \cdot (1 - D_0) > 0$ so this effect increases public surplus under price-linked relative to fixed subsidies.

However, adverse selection also changes the second term in Equation (5). There are two effects on the optimal mandate penalty under price-linked subsidies. The regulator cares about the externality of the marginal consumer, so absent a price response would set $M^* + S^E = E^{\text{marg}} < \bar{E}$. On the other hand, since raising M pushes healthier people into the market which lowers firms' costs and leads to lower prices, the regulator will want a higher M . The net effect is ambiguous, but adverse selection could increase $\bar{E} - (S^E + M^*)$, increasing the difference in public surplus between price-linked and fixed subsidies, potentially outweighing the decrease due to $\Delta \bar{E}$. So it is not clear whether adverse selection makes price-linked subsidies better or worse relative to fixed subsidies.

Cost Uncertainty

When costs are uncertain, the regulator cannot predict the equilibrium subsidy and does not know the externality from uncompensated care and therefore faces uncertainty about what the *ex-post* optimal S and M will be.

Suppose that there are market-level cost shocks that insurers can observe but the government cannot (or can only observe with a substantial delay). Under price-linked subsidies, quantities do not vary much with costs because the increased subsidy cancels out the increased prices. Under a fixed subsidy, when costs are high prices will be high, so more consumers will be uninsured ($\Delta D_0 < 0$) and the externality will be high so $\bar{E} - (S^E + M^*) > 0$. This means $\Delta D_0 \cdot (S^E + M^* - \bar{E}) < 0$, hurting public surplus under fixed-subsidies. With costs lower than expected, the sign of both terms is reversed, so the combined affect is the same. Under cost uncertainty, the increase in markups from price-linked subsidies must be balanced against their ability to indirectly link subsidies to the externality of uninsurance.

Risk Aversion

When costs are uncertain, price-linked subsidies have the property of stabilizing post-subsidy consumer prices and transferring the cost risk to the government. Under our baseline case, which assumes locally quasi-linear utility, this risk transfer has no net welfare impact. However, for our low-income population, it is plausible that changes in insurance premiums materially affect the marginal utility of consumption. As a robustness check, we add the cost of this increased risk to our welfare measure.

Let Y be income and P_{j^*} be the price of a consumer's chosen plan. If consumers have concave utility over non-health insurance consumption, $u(Y - P_{j^*})$, then the change in utility from a price increase of ΔP is

$$\begin{aligned} u(Y - P_{j^*} - \Delta P) - u(Y - P_{j^*}) &\approx -u'(Y - P_{j^*})\Delta P + u''(Y - P_{j^*})\frac{(\Delta P)^2}{2} \\ &= -u'(Y - P_{j^*})\Delta P - \gamma\frac{u'(Y - P_{j^*})}{Y - P_{j^*}}\frac{(\Delta P)^2}{2} \end{aligned}$$

where γ is the (local) coefficient of relative risk aversion. In order to combine consumer utility and government expenditures, we implicitly assumed that prior to any cost shocks, the government set the transfer system “right” so the consumer's marginal utility of a dollar ($= u'(Y - P_{j^*})$) equaled the marginal value of a dollar to the government (which we normalize to one). In that case,

$$\Delta u \approx \Delta P - \frac{\gamma}{2}\frac{(\Delta P)^2}{Y - P_{j^*}}$$

When weighted by demand shares, the first term is the first-order effect on consumer surplus included our previous calculation. The second term is the additional cost of the risk falling on consumers. In Section 4.4, we use the cost of risk in our simulations to assess the empirical importance of this risk transfer.

2 Setting and Data

To understand the quantitative importance of the incentives created by price-linked subsidies, we estimate a model using data from Massachusetts' pre-ACA subsidized health insurance exchange (Commonwealth Care, or “CommCare”). Created in the state's 2006 health care reform, CommCare facilitates and subsidizes coverage for individuals earning less than 300% of the federal poverty line (FPL) and without access to insurance from an employer or another government program. This population is similar to those newly eligible for public

insurance under the ACA. There are 4-5 insurers offering plans during the period we study, making it an appropriate setting to study imperfect competition.²⁴

CommCare’s design is similar to the ACA exchanges but somewhat simpler. There are no gold/silver/bronze tiers – each participating insurer offers a single plan. That plan must follow specified rules for cost sharing and covered medical services. However, insurers can differentiate on covered provider networks and other aspects of quality like customer service. Importantly, these flexible quality attributes apply equally to enrollees in all income groups, a fact we use in estimating demand.

In CommCare, subsidies are linked to the price of the cheapest plan so that this plan costs an income-specific “affordable” amount.²⁵ A consumer’s premium for a plan is the plan’s price (set by the insurer) minus the subsidy for that consumer’s income group. In addition (and unlike the ACA), CommCare applied special subsidies for the below-100% of poverty group that made all plans free, regardless of their pre-subsidy price. We use this fact to aid demand estimation – since this group can purchase the same plans but face different (lower) relative prices.

Since CommCare’s eligibility criteria exclude people with access to other sources of health insurance (including other public programs and employer coverage), eligible individuals’ relevant outside option is uninsurance.²⁶ The price of uninsurance is the mandate penalty after its introduction in late 2007 (see discussion below). Like subsidies, the mandate varies across income groups and is set to equal half of each group’s affordable amount (i.e., half of the post-subsidy premium of the cheapest plan).

We link administrative data from CommCare with data on the uninsured from the American Community Survey (ACS), – an annual 1% sample of U.S. households administered by the Census²⁷ – to get a dataset of CommCare-eligible individuals, whether or not they chose to purchase insurance. We use data from January 2008, when the individual mandate is

²⁴By standard metrics, this is a highly concentrated market (e.g., its statewide HHI is between 2,500 and 3,500 in each year from 2008-2011), though it should be noted that it has more competitors than many of the ACA exchanges.

²⁵In 2009-2011, our main years of analysis, the affordable amount was \$0 for 0-100% and 100-150% of the federal poverty level (FPL), \$39/month for 150-200% FPL, \$77 for 200-250% FPL, and \$116 for 250-300% FPL. These amounts differed slightly in earlier years, and in all cases, we use the actual premiums for demand estimation. One notable change is that the amount for 100-150% FPL declined from \$18 in 2007 to \$0 in 2008 – a natural experiment that we study below.

²⁶In theory, individuals could buy unsubsidized coverage on a separate exchange (“CommChoice”), but these plans have less generous benefits and are more expensive because of the lack of subsidies. People whose employer insurance is deemed “unaffordable” (based on the employer covering less than 20%/33% of the cost of family/individual coverage) are also eligible for CommCare. Because this is likely to be a small group and we have no way of measuring them in the data, we do not attempt to adjust for these individuals.

²⁷We obtained ACS data from the IPUMS-USA website, Ruggles et al. (2015), which we gratefully acknowledge.

fully phased in, to June 2011, just prior to the start of CommCare year 2012,²⁸ when plan choice rules and market dynamics shifted considerably (see Shepard, 2016).

Administrative data from the CommCare program²⁹ let us observe (on a monthly basis) the set of participating members, their demographics, their available plans and premiums, their chosen plan, and their realized health care costs (via insurance claims). The availability of cost data is an advantage of the CommCare setting. (It is one of the only non-employer insurance markets with plan choice and cost data linked at the enrollee level.) For the ACS data, we restrict the sample to people who are uninsured and satisfy CommCare’s eligibility criteria based on age, income, and U.S. citizenship. We re-weight observations to correct for the the sampling in the ACS and the fact that it gives the stock of uninsured and not the flow of people into uninsurance. See Appendix C for more information on the data and sample construction.

We focus on active plan choices made by new enrollees.³⁰ This lets us abstract from inertia known to affect plan switching (Handel, 2013; Marzilli Ericson, 2014) and focus on the initial choices that are the primary driver of market shares. Although an approximation, this simplification has been used in structural work on insurance markets (e.g. Ericson and Starc, 2015) as a way of abstracting from the complex dynamics that inertia creates.

Table 1 shows summary statistics for two samples: the full sample, which we use for model estimation, and a sample restricted to 100-300% of poverty individuals, which we use for equilibrium simulations (to better match the eligible population in the ACA). The raw sample include 448,866 new CommCare enrollees and 4,562 uninsured from the ACS. The latter scales up to a population estimate (of “newly uninsured” people) of 355,749 people, implying that 44.2% of the eligible population is uninsured. While this estimate may seem high, recall that CommCare (like the ACA) is targeted at the relatively small subset of the population without other insurance options.³¹ For the simulation sample, about half of those who enroll in CommCare choose the cheapest plan.

The population is quite poor, with just over half having family income less than the

²⁸Because of the timing mismatch, where CommCare’s year runs from July to June while the ACS is a calendar year sample, we match CommCare years to averages from the two relevant ACS years.

²⁹This data was obtained under a data use agreement with the Massachusetts Health Connector, the agency that runs CommCare. All data are de-identified. Our study protocol was approved by the IRBs of Harvard and the NBER.

³⁰“New” enrollees also includes people who re-enroll after a break in coverage, since these people also must make active choices. We assume CommCare eligibility occurs exogenously due to factors like a job loss or income change. Though existing enrollees have an opportunity to switch plans annually during an open enrollment period, on average, only 5% of people switch plans. Therefore, the initial plan choice plays a primary role in determining demand.

³¹For the full ACS data for Massachusetts, we confirm that the uninsured rate is less than 5%, consistent with the perception of near-universal coverage.

Table 1: Summary Statistics

	All Income Groups (Estimation Sample)		100-300% Federal Poverty Line (Simulation Sample)	
	Newly Insured	ACS Uninsured	Newly Insured	ACS Uninsured
Counts				
Sample size	-	4,562	-	2,225
Newly eligible pop	448,866 (55.8%)	355,749 (44.2%)	218,565 (53.7%)	188,641 (46.3%)
Cheapest Plan's Share	44.6%	-	50.2%	-
Demographics				
Age	37.57 [13.73]	35.13 [12.75]	40.64 [13.35]	36.64 [12.28]
Male	0.49 [0.50]	0.67 [0.47]	0.43 [0.49]	0.62 [0.48]
Pre-subsidy Prices				
Average Plan	\$386	-	\$390	-
Cheapest Plan	\$367	-	\$370	-
Consumer Premiums				
Average Plan	\$23.31	-	\$47.87	-
Cheapest Plan	\$18.42	\$27.72	\$35.13	\$52.28
Mandate Penalty	\$8.77	\$13.16	\$18.01	\$24.82
Costs				
Observed	\$417 [\$1510]	-	\$372 [\$1485]	-
Predicted in Avg Plan	\$418 [\$150]	\$386 [\$151]	\$381 [\$123]	\$340 [\$114]

NOTE: This table shows counts and attributes for eligible consumers (CommCare enrollees and the ACS uninsured) from January 2008 to June 2011. "Pre-subsidy Prices" shows the enrollment-weighted average and the cheapest monthly price paid to firms for an enrollee (insured) or that would be paid (ACS). "Consumer Premiums" are the average (enrollment-weighted) and cheapest post-subsidy monthly prices consumers paid (insured) or would have paid (ACS) for plans. (The government pays the difference.) The mandate penalty – which is set by law as half of each income group's affordable amount – is what we calculate the uninsured paid and the insured would have paid under the Massachusetts policy. For each consumer the predicted cost is if they were to enroll in the average plan. Costs are above prices because they refer only to the new enrollees used in the estimation and new enrollees are more expensive than existing enrollees.

poverty line. 20.8% have incomes between 100-150% of poverty; 15.5% between 150-200%; 8.2% between 200-250%; and 4.2% between 250-300%. Consumers' ages range from 19-64; the uninsured are slightly younger than the insured. Pre-subsidy monthly prices average \$386, but subsidies are quite large. Below-poverty consumers pay nothing, and consumers with incomes 100-300% of poverty (the simulation sample) pay an average of \$48/month. The cheapest plan is somewhat less – \$35/month. We estimate the uninsured pay an average mandate penalty of \$24. (Since the insured are somewhat poorer, the average mandate penalty they would have paid is lower, \$18.) New enrollees are more expensive than average enrollees, so the average cost in our estimation sample, \$417, is actually above the average price. If we include current enrollees, costs are below price.³² Predicted costs for the uninsured are somewhat lower, consistent with that being a healthier population.

2.1 Natural Experiments

Conveniently, there are two sources of exogenous variation in the relative price of uninsurance. We use these natural experiments to allow us to identify a statistic that the theory says is key for our model: the responsiveness of demand for the cheapest plan with respect to the price of the outside option. In this section, we show the basic patterns graphically and the main estimates. Details are in Appendix D

Mandate Penalty Introduction Experiment

Though a requirement to obtain insurance took effect in July 2007, the requirement was not enforced by financial penalties until December 2007. Those earning more than 150% of poverty who were uninsured in December forfeited their 2007 personal exemption on state taxes – a penalty of \$219 (see Commonwealth Care, 2008). Starting in January 2008, the mandate penalty was assessed based on monthly uninsurance. The monthly penalties for potential CommCare customers ranged from \$17.50 to \$52.50, depending on income.

There was a spike in new enrollees into CommCare for people above 150% of poverty exactly concurrent to the introduction of the financial penalties in December 2007 and early 2008. Figure 1 shows this enrollment spike for the cheapest plan, which is proportional to the spike for all plans. To make magnitudes comparable for income groups of different size, the figure shows new enrollments as a share of the same plan's total enrollment in that income group in June 2008.³³

³²Higher costs in the sample is also part of why prices in our simulations are higher than observed.

³³We use June 2008 as a baseline because enrollment, which had been steadily growing since the start of CommCare, stabilizes around June 2008. Therefore, we treat June 2008 enrollment as an estimate of equilibrium market size.




Figure 1: New Enrollees in Cheapest Plan by Month

NOTE: These are monthly new enrollees (both first-time consumers and those re-enrolling after a break in coverage) into CommCare’s cheapest plan as a share of total June 2008 enrollment, so units can be interpreted as fractional changes in enrollment for each group. The vertical line is drawn just before the introduction of the mandate penalty, which applied only to the “150-300% Poverty” income group. The “150-300% Poverty (Other Years)” combines all years in our data *except* July 2007–June 2008. Premiums varied by region and income group, so the cheapest plan is defined at the individual level but held constant across the time frame.

We believe that this enrollment spike was caused by the financial penalties for a few reasons. There were no changes in plan prices or other obvious demand factors for this group at this time. As Figure 1 shows, there was no concurrent spike for people earning less than poverty (for whom penalties did not apply), and there was no enrollment spike for individuals above 150% of poverty in December-March of other years.³⁴ Additionally, Chandra et al. (2011) show evidence that the new enrollees after the penalties were differentially likely to be healthy, consistent with the expected effect of a mandate penalty in the presence of adverse selection.

We estimate the semi-elasticity associated with this response using a triple-differences

³⁴People earning 100-150% of poverty are omitted from this analysis because a large auto-enrollment took place for this group in December 2007, creating a huge spike in new enrollment. But the spike occurred only in December and was completely gone by January, unlike the pattern for the 150-300% poverty groups. This auto-enrollment did not apply to individuals above 150% of poverty (Commonwealth Care, 2008) so it cannot explain the patterns shown in Figure 1.

specification, analogous to the graph in Figure 1. With one observation per month, t , for both the treatment and control group, g , we estimate

$$\begin{aligned} NewEnroll_{g,t} = & \alpha_0 + \beta_0 \cdot DM_t + \gamma_0 \cdot Treat_t + \delta_0 \cdot X_t \\ & + (\alpha_1 + \beta_1 \cdot DM_t + \gamma_1 \cdot Treat_t + \delta_1 \cdot X_t) T_g + \varepsilon_{g,t}, \end{aligned}$$

where the dependent variable is new enrollment divided by that group’s enrollment in June 2008, $Treat_t$ is a dummy for the treatment period December 2007 - March 2008, DM_t is a dummy for the months December through March in all years, X_t is a vector of time polynomials and CommCare-year dummies, and T_g is a dummy for the treatment group. The coefficient of interest is γ_1 . We estimate that the mandate penalty caused a 22.5% increase in enrollment in the cheapest plan relative to its enrollment in June 2008. (See Appendix D for full regression results and robustness checks.) As we describe in Section 3, we match this statistic as a moment in the structural model. It helps identify unobserved heterogeneity in demand for insurance, ensuring that our model matches this key determinant of the size of the pricing distortion.

Affordable Amount Decrease Experiment

As a robustness check for the effects measured from the introduction of the mandate penalty, we use a change in the affordable amount that occurred in July 2007. (Plan prices were fixed from November 2006 to June 2008.) CommCare set subsidies so that the post-subsidy premium for the cheapest plan equaled a pre-set “affordable amount;” for fixed pre-subsidy prices, a \$1 decrease in the affordable amount has an equivalent effect on relative prices as a \$1 increase in the mandate penalty. This approach addresses the potential concern that the *introduction* of a mandate penalty may have a larger effect (per dollar) than a *marginal increase* in a penalty. It also allows us to obtain estimates for the 100-150% poverty group, who faced a \$0 mandate penalty. The estimated semi-elasticity for the 150-200% poverty group (the only group that faced changes both times) is very similar to the one estimated from the mandate penalty’s introduction.

3 Structural Model

To study the effects of subsidy structure, we estimate a demand system and cost function using the CommCare data and then simulate equilibrium under alternative subsidy regimes. In this section, we specify the model and simulation method and discuss the assumptions for identification.

3.1 Demand

Using the dataset of CommCare enrollees and eligible uninsured individuals described in Section 2, we estimate a random coefficient logit choice model for insurance demand. Consumers choose between CommCare plans and an outside option of uninsurance based on the relative price and quality of each option. Each consumer i is characterized by observable attributes $Z_i = \{r_i, t_i, y_i, d_i\}$: r is their region, t is the time period (year) they make their choice, y is income group, and d is their demographic group (gender crossed with 5-year age bins). We suppress the i subscript when the attribute is itself a subscript.

The utility for consumer i of plan j equals

$$u_{ij} = \underbrace{\alpha(Z_i) \cdot P_{yj}^{cons}}_{\text{Premium}} + \underbrace{\xi_j(Z_i)}_{\text{Plan Quality}} + \underbrace{\epsilon_{ij}}_{\text{Logit Error}} \quad j = 1, \dots, J$$

where P_{yj}^{cons} is the plan's premium for consumer i (which depends on income), $\xi_j(Z_i)$ is plan quality, and ϵ_{ij} is an i.i.d. type-I extreme value error giving demand its logit form. Price sensitivity can vary with income and demographics: $\alpha(Z_i) = \alpha_y + \alpha_d$. We capture plan quality with coefficients on plan dummies, which vary by region-year and region-income bins: $\xi_j(Z_i) = \xi_{j,r,t} + \xi_{j,r,y}$. We allow for this flexible form both to capture variation across areas and years (e.g., due to differing provider networks) and to aid in identification, as discussed below.

The utility of the outside option of uninsurance equals

$$u_{i0} = \underbrace{\alpha(Z_i) \cdot M_i}_{\text{Mandate Penalty}} + \underbrace{\beta(Z_i, \nu_i)}_{\text{Utility of Uninsurance}} + \underbrace{\epsilon_{i0}}_{\text{Logit Error}}$$

where M_i is the mandate penalty and $\beta(Z_i, \nu_i)$ is the utility of uninsurance. Rather than normalizing the utility of the outside option to zero (as is often done), we normalize the average plan quality ($\xi_j(Z_i)$) to zero, letting us estimate β , the utility of uninsurance, for different groups. We allow it to vary with observable factors and an unobservable component: $\beta(Z_i, \nu_i) = \beta_0 + \beta_y + \beta_r + \beta_t + \beta_d + \sigma\nu_i$, with $\nu_i \sim N(0, 1)$. The random coefficient captures the idea that the uninsured are likely to be people who, conditional on observables, have low cost of uninsurance. This allows us to better match substitution patterns – including the elasticity of insurance demand with respect to the mandate penalty.

Estimation and Identification

We estimate the model by simulated method of moments, incorporating micro moments with (group-specific) market shares and interactions of average chosen characteristics and consumer observables, an approach similar to Berry et al. (2004). These moments identify the coefficients in $\alpha(Z_i)$ and $\xi_j(Z_i)$, as well as the non-random coefficients in β . The assumptions underlying identification of the premium coefficients deserve special discussion. While most papers use instruments to address the concern of correlation between price and unobserved plan quality, we follow Shepard (2016) in using *within-plan* premium variation created by the exchange’s subsidy rules.³⁵ We use the fact that subsidies make all plans free for below-poverty enrollees, while higher-income enrollees pay higher premiums for the same plans. This structure also creates differential premium changes over time, which we use for identification. For instance, when a plan increases its price between years, its premium increases for higher income groups, but there is no premium change for below-poverty enrollees (since it remains \$0). Econometrically, the plan-region-year dummies ($\xi_{j,r,t}$) absorb premium differences arising from plan pricing and plan-region-income dummies ($\xi_{j,r,y}$) absorb persistent demand differences across incomes. The premium coefficients are identified from the remaining variation, which is entirely from within-plan differential changes across incomes.

In addition, we employ a novel approach to estimate the variance of the random coefficient on uninsurance (σ). We use the change in insurance coverage around the natural experiment of the introduction of a mandate penalty, as described in Section 2.1. Specifically, we match the estimated 22.5% coverage increase to our model’s predicted coverage increase for the same time period when M goes from zero to its actual level in early 2008. This identification works because of the classic intuition that σ affects substitution patterns. If there is more heterogeneity in the value of uninsurance, the uninsured will tend to be people with lower values of insurance who are unlikely to start buying insurance when the mandate penalty increases. Thus, higher values of σ generate less demand response to the mandate penalty, and vice versa. Appendix E shows details of the method and formulas of all moments.

3.2 Costs

To simulate pricing equilibrium, we need to model each insurer’s expected cost of covering a given consumer. We use the observed insurer costs in our claims data to estimate a simple cost function. We assume that costs are generated by a Poisson regression model (also known as a generalized linear model with a log link) with expected costs for consumer i in plan j

³⁵The discussion that follows closely follows that of Shepard (2016).

in year t of

$$E(c_{ijt}) = \exp(\mu X_{it} + \psi_{j,r,t}). \quad (6)$$

Costs vary with consumer observables (X_{it}) and a (region-year specific) plan effect, $\psi_{j,r,t}$. Although our claims data include a rich set of consumer observables, our inclusion of the uninsured population limits us to what we can also observe in the ACS: age-sex groups and income group. Our model nonetheless captures adverse selection through the correlation between insurance demand and demographics.³⁶ Costs vary across plans because of their different provider networks. We let the plan effects vary by region and year to capture network differences over time and across areas.

A concern with a basic MLE estimation of Equation (6) is that estimates of $\psi_{j,r,t}$ will be biased by selection on unobserved cost. This is particularly concerning because X_{it} includes a relatively coarse set of observables. To address this issue, we estimate the $\psi_{j,r,t}$ parameters in a separate version of Equation (6) with individual fixed effects. These estimates are identified based only on within-person cost variation when an individual switches plans or leaves the market and later re-enrolls in a different plan.³⁷ We then adjust observed costs by removing the estimated plan component and estimate the coefficients on individual characteristics from cross-person variation with all enrollees. We use the resulting predicted values of $E(c_{ijt})$ as our estimates of costs for each enrollee-plan possibility.³⁸

Risk Adjustment

Because firms can only set one price, their pricing incentives may be affected by adverse selection. Risk adjustment can mitigate or eliminate selection on observables across plans in the exchange. However, it does not mitigate the effects of adverse selection into any

³⁶The level of detail in our cost model is comparable to past structural work that includes uninsurance as an option (e.g. Ericson and Starc, 2015; Tebaldi, 2016). In general, adverse selection can also be driven by correlation between groups' cost estimates and their price-sensitivity (α) and utility of uninsurance (β) in the demand model. In our estimates, selection manifests primarily by younger individuals and males (lower costs consumers) having lower demand for insurance (via β).

³⁷This method eliminates selection if individuals' risk factors are stable over time or uncorrelated with plan changes. It will, however, miss selection on cost trends – e.g., if individuals tend to switch plans just before they get sick. We have experimented with using instrumental variables to deal with the selection problem but have not found another approach that works.

³⁸The above model captures insurers' medical costs. In addition, insurers incur administrative costs for functions like claims processing. Using plan financial reports, we estimate variable administrative costs by regressing administrative costs on plan enrollment, which results in an estimate of \$30 per member-month. Following Massachusetts rules, we assume that insurers are paid a per-member-month administrative fee of the same amount (\$30) that exactly nets out this cost. This fee is not subject to the risk adjustment described below. Thus, administrative costs drop out of the profit function, and our model is not affected by them.

insurance, since the exchange cannot transfer money away from people who never enroll. We model a simple risk adjustment system that fully addresses cross-plan selection but maintains the effect of selection into insurance.

Based on the policy in Massachusetts (similar to that of the ACA), we assume the exchange estimates a ϕ_{it} for each individual-year that indicates how costly they are expected to be relative to an average enrollee. It then adjusts payments accordingly so the expected profits that plan j earns from enrolling individual i at time t are

$$\pi_{ijt} = \phi_{it}P_{jt} - C_{ijt}.$$

We assume that ϕ_{it} perfectly captures the individual (non-plan) portion of the cost function: $\phi_{it} = \exp(\mu X_{it}) / \bar{c}_t$, where \bar{c}_t is the average of $\exp(\mu X_{it})$ across individuals who buy insurance. Thus, risk adjustment fully offsets individuals expected costs in a proportional way.³⁹ Expected profits for an individual simplify to $\pi_{ijt} = \phi_{it} (P_{jt} - \bar{c}_t \exp(\psi_{j,r(i),t}))$.⁴⁰

3.3 Equilibrium and Simulations

To draw closer parallel to the ACA markets, we base our simulations on consumers above poverty. We use the ACA affordable amounts and mandate penalties, but stick with the MA subsidy rule since we have no basis on which to predict demand for the different tiers available in the ACA. We use the demand and cost models to specify firm profits as a function of all plans' prices and the subsidy policy:

$$\pi_j = \sum_i (\phi_i P_j - C_{ij}) \cdot Q_{ij}(P^{Cons}(P))$$

where we have dropped t subscripts for convenience and $P^{Cons}(P)$ is the subsidy function mapping prices into consumer premiums. We assume that each insurer sets its price to maximize profits in static Nash equilibrium. This equilibrium is defined by the first-order conditions (FOCs) $\partial\pi_j/\partial P_j = 0$ for all j , given all other plans' prices. These FOCs can be expressed as:

$$P_j = \frac{1}{\phi_j^m} \left(C_j^m + \frac{\bar{\phi}_j}{-\frac{1}{Q_j} \frac{dQ_j}{dP_j}} \right) \quad (7)$$

³⁹In principle, we could also model imperfect risk adjustment by assuming that $\phi_{it}^{Imperfect} = (\phi_{it})^\gamma$ for $\gamma \in (0,1)$. Because the main predictions of our model do not depend obviously on cross-plan adverse selection, we have not explored this avenue.

⁴⁰We note the interesting property that with "perfect" risk adjustment under this system, markups are larger for sicker (higher ϕ_{it}) individuals. This appears to be a natural byproduct of the proportional risk adjustment system rather than something we could offset by using a different value for ϕ_{it} .

where $C_j^m \equiv \frac{dQ_j}{dP_j}^{-1} \sum_i C_{ij} \frac{dQ_{ij}}{dP_j}$ is the (average) cost of plan j 's marginal consumers, ϕ_j^m is the (analogously defined) marginal risk score, and $\bar{\phi}_j \equiv \frac{1}{Q_j} \sum_i \phi_i Q_{ij}$ is the average risk score for all consumers in plan j . For each year and subsidy policy, we simulate equilibrium numerically by searching for the price vector P that satisfies these equilibrium conditions for all insurers. We do this both for the set of insurers present in the market in 2009 and 2011 and for counter-factual markets with only two insurers.

Note that $\frac{dQ_j}{dP_j}$ in Equation (7) is a total derivative that incorporates any effect of increasing P_j on the subsidy, which occurs under price-linked (but not fixed) subsidies. This introduces a discontinuity in the FOC of the cheapest plan at the price of the second-cheapest plan: below it, they are subsidy-pivotal (so $\frac{dQ_j}{dP_j} = \frac{\partial Q_j}{\partial P_j} + \frac{\partial Q_j}{\partial M}$) while above it, they are not (so $\frac{dQ_j}{dP_j} = \frac{\partial Q_j}{\partial P_j}$). This means that, in some cases, multiple plans set the same cheapest price in equilibrium; this equilibrium price is generally not unique, instead there is a range of prices that the two cheapest plans may have in equilibrium. When this occurs in our simulations (for 2009 but not 2011), we show results for both the minimum and maximum possible prices at which the cheapest plans can tie.

Cost Uncertainty

We allow for cost uncertainty by first setting the fixed subsidies equal to the price-linked subsidies that emerge in equilibrium with baseline costs. We then simulate equilibrium under each subsidy policy when actual costs are scaled upward/downward by between -15% and +15%. This proportional adjustment applies to all plans' costs and to an individual's cost of uncompensated care. We assume that firms observe this cost shock and price based on it but that it is unobserved by the regulator. Thus, price-linked subsidies automatically adjust with the cost shock, while fixed subsidies do not. We also assume that the shock is unobserved to consumers so their demand for insurance is unaffected.⁴¹

3.4 Welfare

We calculate welfare based on the framework described in Section 1.2. For each subsidy policy and cost shock, we use the estimated demand parameters and the simulated equilibrium prices to calculate the fixed component of utility for each consumer for each plan, $\hat{u}_{ij} \equiv u_{ij} - \epsilon_{ij}$. These allow us to calculate each individual's expected consumer surplus $CS_i = \log \left(\sum_{j=0}^J \exp(\hat{u}_{ij}) \right)$ and choice probabilities ($\widehat{Pr}_{ij} = \exp(\hat{u}_{ij}) / \exp(CS_i)$). The

⁴¹This assumption on demand would hold, for instance, if consumers had to choose whether to buy insurance before observing health care prices, or if the price shock did not affect consumers' out-of-pocket costs of care when uninsured (e.g., if these were determined purely by ability to pay).

government’s expected net expenditures on each individual are $G_i = S - (S + M) \cdot \widehat{Pr}_{i0}$.

The externality avoided when an individual has insurance is composed of two parts. The first is the uncompensated care costs, which are, unfortunately, quite difficult to measure (see, e.g. Garthwaite et al., 2015). Instead of attempting to estimate them in our Massachusetts setting, we assume that uncompensated care costs are proportional to an individual’s expected costs in the average exchange plan, and therefore increases proportionally with the cost shock, Δ . We use,

$$C_{it}^{Uncomp} = \lambda \cdot \exp(\hat{\mu}X_{it})(1 + \Delta), \quad \lambda = \{.6, .8, 1\}$$

Using evidence from the Oregon health insurance experiment, Finkelstein et al. (2015) find that when individuals shift from Medicaid to uninsurance, third parties cover about 60% of what their costs would have been if insured (with individuals paying 20% and 20% being moral hazard). Mahoney (2012) finds that the uninsured only pay 20% of their health care bills. Since price-linked subsidies are more attractive when the regulator wants the subsidy to move more with price, we consider $\lambda = 1$ as the best-case scenario for price-linked subsidies.

To rationalize the observed level of subsidies, we include a second part of the externality, E_0 , that is driven by things other than uncompensated care costs – e.g. paternalism, moral responsibility, etc. For each simulation year and specification of λ , we calibrate E_0 so that absent a cost shock public surplus is maximized by a fixed subsidy of the baseline size. This calibration is conceptually important. Our goal is to understand how cost uncertainty affects the case for price-linked subsidies assuming that subsidies would be set optimally if costs were known. If we do not calibrate E_0 , our welfare results are partly driven by the deviation of subsidies from their optimal level at $\Delta = 0$. Empirically, we find that a positive E_0 is needed to rationalize the ACA’s subsidies, which we use as our baseline.

Putting it all together, the public surplus for individual i equals

$$PS_i = \log \left(\sum_{j=0}^J \exp(\hat{u}_{ij}) \right) - S + (S + M) \cdot \widehat{Pr}_{i0} + \left(1 - \widehat{Pr}_{i0} \right) \left(\lambda \cdot \exp(\hat{\mu}X_{it})(1 + \Delta) + \hat{E}_0 \right).$$

Note that we add the “saved externality” from when i has insurance instead of subtracting the externality when uninsured. The two are conceptually equivalent, and if we subtract the externality, the large direct surplus loss from higher costs makes it hard to see the differences between subsidy regimes. Total public surplus is the sum of PS_i across all individuals in our data.

We also consider the effect of consumer risk aversion, since, with cost uncertainty, price-linked subsidies have the advantage of shielding consumers from premium variation. As

discussed in Section 1.2, if γ is the coefficient of relative risk aversion, a price change of ΔP has an additional cost (or a decreased benefit) of $\frac{\gamma}{2} \frac{(\Delta P)^2}{c}$ relative to the effect on a risk-neutral agent. Chetty (2006) argues that $\gamma \leq 2$; to consider the case most favorable to price-linked subsidies, we use $\gamma = 2$. We define c based on an estimate of each consumers' monthly income minus their expected premium or mandate penalty.⁴²

4 Results

4.1 Estimated Parameters

Demand

The demand coefficients are summarized in Table 2. The average price coefficient ranges from -0.049 for those just above the poverty line to -0.024 for those making 250%-300% FPL; it does not vary systematically with age or gender. On average, consumers prefer Fallon to Network Health to BMC to NHP to CeltiCare, with the biggest difference being between Celticare and the rest of the plans. The cost of ininsurance is small for young males, but much larger for females and increasing in age. We allow the value of uninsurance to vary with observable and unobservables. Both generate large variance in the value of uninsurance, but unobservables seem to play a larger role, leading to a standard deviation of 2.59 across individuals, relative to the 0.99 standard deviation explained by observables.

Because the logit parameters can be hard to interpret, Table 3 shows the semi-elasticities with respect to own price and with respect to the mandate penalty.⁴³ The own-price semi-elasticity increases over time, largely because of the entry of a plan (CeltiCare) in 2010. Our estimates are quite a bit higher than the elasticities reported in Chan and Gruber (2010) (even after adjusting to allow for substitution to uninsurance, which they do not consider). Our results may differ because we allow for heterogeneity in price coefficients by income and demographics, and we also use below-poverty enrollees (for whom all plans are free) as a control group.

⁴²Specifically, for c we take the middle of a consumer's income bin, divide by 12 for monthly income and subtract off the average monthly premium (or mandate penalty) they pay. For fixed subsidies, we calculate (ΔP) for each plan relative to baseline, square it, weight by the individual's probability of purchasing each plan and divide by c . The results are not significantly different if we use the consumer price difference between price-linked and fixed subsidies because cost shocks have little effect on consumer prices under price-linked subsidies, so they are always close to baseline.

⁴³Semi-elasticity refers to the percentage change in demand resulting from a \$1 change in a price or penalty.

Table 2: Parameters in Demand Model

(a) Price and plan parameters				(b) Uninsurance parameters		
				Value of Uninsurance		
				Avg	S.E.	
Premium Coefficient						
100-150 % Pov	-0.049***	0.0067		Average	-0.123	0.353
150-200 % Pov	-0.029***	0.0079		5yr Age Bin	-0.091***	0.029
200-250 % Pov	-0.027***	0.0087		Female	-1.056***	0.270
250-300 % Pov	-0.024*	0.0131		Standard Deviation		
Plan Dummies (BMC=0)				Total		3.58
CeltiCare	-1.14***	0.147		From Observables		0.99
Fallon	0.103	0.116		From Unobservables (σ)		2.59
NHP	-0.342***	0.053				
Network Health	0.053	0.1064				

Note: This table reports the average and standard error of demand parameters across individuals and how they differ with demographic characteristics. The plan dummies are standardized by setting the value of BMC to zero.

Table 3: Average Semi-elasticities

Own Price Semi-Elasticity				Semi-Elasticity of Insurance w.r.t. Mandate Penalty	
By plan		By Year		By Year	
BMC	-2.06%	2008	-2.09%	2008	0.79%
CeltiCare	-3.09%	2009	-2.23%	2009	0.95%
Fallon	-2.70%	2010	-2.57%	2010	1.05%
NHP	-2.69%	2011	-2.59%	2011	1.09%
Network Health	-2.41%	All	-2.38%	All	0.98%

Note: The semi-elasticity is the percent change in demand induced by a \$1 change in price. The left panel reports the average across years of the own price semi-elasticity for each plan and the (share-weighted) average across plans for each year. The right panel reports the semi-elasticity of buying any insurance with respect to the mandate penalty. We allow the price-sensitivity parameters to vary by year and there are compositional changes in the demographics and plan shares which contribute to the variation.

Costs

The first part of Table 4 shows the differences in costs across plans, broken down by before and after 2010, when Celticare entered the market. Prior to Celticare’s entry Neighborhood Health had the lowest costs, 8% below the share-weighted average. Other plans’ costs for treating the same patient ranged from .5% to 23.7% higher. When Celticare entered its costs were 26% lower than Network Health’s.

Table 4: Cost Parameters

	Plan Effects				
	BMC	CeltiCare	Fallon	NHP	Network Health
2008-2010	-7.16%		+8.69%	14.16%	-7.68%
2010-2011	-1.32%	-31.39%	+9.69%	+14.89%	-7.32%
	Percent of Federal Poverty Line				
	100-150	150-200	200-250	250-300	
Relative to < Poverty	-27.09%	-25.55%	-27.67%	-34.75%	

Note: The plan effects give the percent difference between the expected cost for a consumer under that plan and the share-weighted average (for that consumer) in that time period. The plan parameters, ψ_j , are estimated in a Poisson regression with individual fixed effects; the reported percentages are $\exp \psi_j - 1$, normalized so that the share-weighted average is zero. The coefficients on income group are estimated in a Poisson regression from the full sample after removing the estimated firm component from the observed costs. The reported percentages are $\exp \mu_g - 1$, when the μ_g for the ‘Less than Poverty’ group is set to zero.

The second part of Table 4 summarizes the cost parameters for income groups. All income groups have substantially (25-35%) lower costs than the group below the federal poverty line, though costs are not strictly decreasing with income. We do not report the parameters for the demographic groups, but they are as expected – costs increase with age and are higher for females at young ages and higher for males at older ages.

4.2 Baseline Simulations

We start with the simplest case where we assume the government anticipated the market equilibrium and set the affordable amounts optimally. We compare the equilibrium under this price-linked policy to the equilibrium under a fixed subsidy policy where the government sets the subsidy equal to the equilibrium subsidy under the price-linked policy. We do this separately for 2009 and 2011 because the competitive dynamics in the market were very different before and after Celticare’s entry.

In 2009 Network Health and BMC are the cheapest plans and have fairly similar costs.

It is therefore not surprising that we find that under price-linked subsidies, the price of the second cheapest plan acts as a binding upper-bound for the cheapest plan. For a given level of the other plan’s price, each plan’s profit function has a kink when its price equals the other plan’s price, since below that level its price affects the subsidy and above it does not. This generates a range of equilibria where BMC and Network Health have the same price. We compare the extrema of this range to the equilibrium with a fixed subsidy set equal to the lowest price-linked equilibrium subsidy, thinking that was the equilibrium the government would be targeting when it set the affordable amount. In 2011, Celticare’s cost are enough lower than all the other plans, that the bound created by the second cheapest plan’s price is not binding and there is only one equilibrium.

Table 5 reports the prices, subsidy and insurance rate for each equilibrium. In 2009, depending on the equilibrium, prices for BMC and Network health are between \$3 and \$21 higher under price-linked subsidies than under fixed subsidies. For Celticare in 2011, the price is \$22 higher under price-linked subsidies. The cheapest plan gains market share, which pushes the change in average price to be larger than the change in the cheapest price, but this is dominated by the fact that other plans have smaller prices changes, and some have slightly higher prices under fixed subsidies, so the drop in average price is smaller than the drop in the lowest price. Because of the lower prices, the fixed subsidies lead to slightly higher rates of insurance purchase.

Table 5: Prices, Subsidies, and Insurance Rate under Price-Linked and Fixed Subsidies

	Subsidy Type	Min Price	Average Price	Average Subsidy	Share Insured
2009	Min Price-linked	\$458.41	\$461.17	\$393.80	57.3%
	Max Price-linked	\$477.26	\$480.76	\$412.32	57.8%
	Fixed	\$454.98	\$459.35	\$393.65	57.9%
	Min Difference	\$3.43	\$1.82	\$0.15	-0.6%
	Max Difference	\$22.29	\$21.40	\$18.66	-0.2%
		Difference			
2011	Price-Linked	\$424.61	\$448.19	\$362.79	37.5%
	Fixed	\$403.06	\$427.22	\$364.28	42.5%
	Difference	\$21.55	\$20.97	-\$1.49	-5.0%

Note: This is a comparison of the market equilibria under price-linked and fixed subsidies. In 2009, there are a range of equilibrium under price-linked subsidies; this reports statistics for the equilibrium with the minimum and maximum lowest price. The average price is weighted by plan share. The average subsidy is across income groups and regions.

Table 6 reports different components of welfare. The difference in consumer surplus is so

Table 6: Costs and Surplus under Price-Linked and Fixed Subsidies

Subsidy Type	Per Eligible Consumer per Month					
	Consumer Surplus	+ Saved Externality	- Gov Costs	= Public Surplus	Profits	
2009	Min Price-linked	\$9.99	\$262.76	\$197.08	\$75.68	\$34.04
	Max Price-linked	\$10.68	\$264.59	\$209.50	\$65.77	\$43.98
	Fixed	\$10.92	\$265.25	\$199.48	\$76.69	\$33.53
	Min Difference	-\$0.93	-\$2.48	-\$2.40	-\$1.01	\$0.51
	Max Difference	-\$0.24	-\$0.66	\$10.02	-\$10.92	\$10.46
2011	Price-Linked	-\$10.56	\$160.67	\$98.59	\$51.53	\$21.58
	Fixed	-\$6.10	\$179.79	\$117.75	\$55.94	\$21.98
	Difference	-\$4.46	-\$19.12	-\$19.16	-\$4.42	-\$0.40

Note: This is a comparison of the market equilibria under price-linked and fixed subsidies. In 2009, there are a range of equilibrium under price-linked subsidies; this reports statistics for the equilibrium with the minimum and maximum lowest price. ‘Per Eligible Consumer’ includes the uninsured. Consumer surplus is relative to the market not existing where consumers get the (dis)utility of uninsurance, but do not have to pay the mandate penalty. Government costs are subsidy expenditures minus mandate revenue. Saved externality is the sum across consumers of the probability that they do not buy insurance times their expected externality. Public surplus adds consumer surplus to the saved externality and subtracts government costs.

much smaller than the price change because it is per eligible consumer and only about 20% of eligible consumers purchase the cheapest plan. Celticare consumers were better off by the full \$21.55 in 2011. Consumer surplus is calculated relative an individual’s expected cost of uninsurance – the market not existing and there being no mandate penalty. In 2011, the subsidy was relatively low because Celticare was so cheap; if consumers did not like Celticare, they had to either pay the mandate penalty or had to pay a higher (than 2009) price for one of the other plans.⁴⁴ The negative consumer surplus indicates that on average they were worse off than not having insurance and not paying the mandate penalty. Government costs are the subsidies paid minus the mandate penalty revenue. The externality is calibrated to make the equilibrium subsidy optimal (see discussion in Section 3.4).

Total public surplus adds this saved externality to the average consumer surplus and subtracts government costs.⁴⁵ When the fixed subsidy is set equal to the price-linked subsidy – as in 2011 and the minimum price-linked equilibrium in 2009 – the differences in government costs and the saved externality are both due to different rates of insurance, so they push public surplus in opposite directions by about the same amount. Therefore, in 2011, the difference in public surplus is about equal to the difference in consumer surplus: \$4 per month per potential customer higher under fixed subsidies. In 2009, there is a price-linked

⁴⁴Table 5 shows that the difference between the average price and the minimum price was much larger in 2011 than in 2009.

⁴⁵See Section 1.2 for a discussion of why we focus on public surplus.

equilibrium where the difference in public surplus relative to the fixed subsidy is less than \$1, indicating the potential for the second cheapest plan to effectively limit the distortion – though under the ACA it would have to be the third cheapest plan.⁴⁶ Even with a second low-cost plan in the market, there are equilibria that are much worse for the government. If the plans coordinate on their preferred equilibrium, then relative to the low price-equilibrium, the subsidy is \$19(= \$22 – \$3) higher, leading to public surplus that is \$10 lower per *potential* consumer.

Profits, which are not included in public surplus, do not seem to vary systematically with the subsidy type, except in 2009 where price-linked subsidies are much more profitable if the insurers manage to coordinate on the higher-priced equilibrium. In the other cases, fixed subsidies increase price competition, but because of adverse selection, the lower prices draw lower-cost consumers into the market, so the net affect on producers can be slightly positive. (This does not imply that insurers would have been better off if they had colluded to lower prices because under the price-linked subsidies, that would have also lowered the subsidy, so they would not have gotten the resulting increase in demand and healthier consumers.)

Without calibrating the externality, which requires assuming that the price-linked equilibrium subsidy is optimal, we can still show how alternative levels of fixed subsidy could save the government money or increase insurance coverage relative to the price-linked subsidy. We find that with fixed subsidies, in 2011, the government could have spent the same amount and gotten 2.72 percentage points more health insurance coverage or have lowered net expenditures by 5.33% and gotten the same rate of insurance coverage. In 2009, it depends which price-linked equilibrium we compare to: .34-3.5 percentage points more insurance for the same cost or .57-6.11% lower spending for the same coverage.

Comparison to the reduced-form equation

We can plug the estimated semi-elasticities for the cheapest plan with respect to own price and with respect to the mandate penalty into the reduced form approximation of the distortion from Equation (4). With perfect risk adjustment, we get a distortion of \$33 for 2009 and \$28 for 2011. We think the difference from the structural estimates is due to strategic interactions with the other plans, which are not captured in the reduced-form equation. In order to convert the equivalent cost change into an estimated price increase, the reduced-form approach assumes that a firm’s semi-elasticity are unaffected by other firms’ prices.

⁴⁶Subsidies under the ACA are linked to the second cheapest silver plan, so it would be the third cheapest silver plan that could limited the distortion. Though as we see here, even then, there may be a higher-price equilibrium with a substantial subsidy distortion.

4.3 Fewer Insurers

Many of the ACA markets have only two insurers. To see how this effects the distortion we again simulate the market for price-linked and fixed subsidies with our demand system and costs estimated for 2011, but with only two insurance plans available. Table 7 shows the distortion for each pair of the four main insurers – Fallon was a smaller plan, not available in all regions.

The plans are listed in order of increasing costs: Celticare is the lowest cost, Network Health is substantially more expensive, BMC is slightly more expensive than Network Health, and NHP is by far the highest cost (see Table 4). All of the distortions are somewhat larger than what we estimate for the market with five insurers, but more notable is how much larger the distortion is when there is a large cost difference between the two plans. While the number of insurers definitely matters for the price levels, the type of insurers – how close their costs are – really matters for the price distortion.

Table 7: Price Distortion with Two Insurers in 2011

		(Lower cost)	→	(Highest cost)
		Network Health	BMC	NHP
(Lowest cost)	Celticare	\$24.46	\$35.20	\$49.66
↓	Network Health		\$29.03	\$34.74
(Higher cost)	BMC			\$32.17

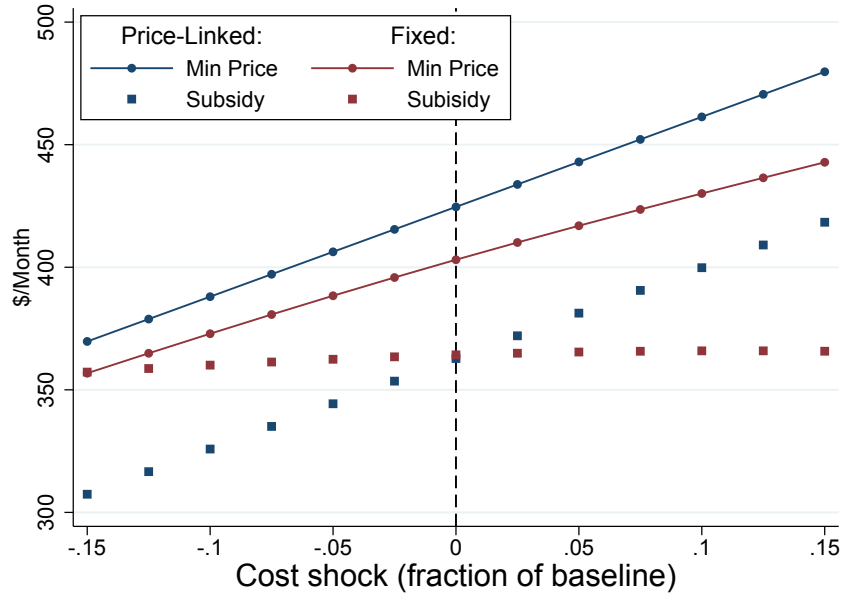
(All 5 plans in the market: \$21.55)

Note: This table shows the difference in the cheapest price under price-linked and fixed subsidies for markets with only two insurers. The plans are listed in order of increasing costs: Celticare, Network Health, BMC, NHP.

4.4 Cost Uncertainty

The baseline analysis assumed that the government knows health care costs when setting the subsidy level. We now look at how a fixed subsidy system fairs when costs differ from baseline. For each subsidy regime, we recalculate the firm-pricing equilibrium when costs differ from baseline by up to 15% (costs ranging from 85% to 115% of estimated costs). Figure 2 shows the price of the cheapest plan and the subsidy. Unsurprisingly, prices increase with costs; they increase slightly faster under price-linked subsidies than under fixed subsidies. (A 15% cost shock is an average change of about \$56, so the slope of the fixed subsidy a little less than 1.0.) Figure 3 shows the share of people who purchase insurance. Under

Figure 2: Price of the cheapest plan under cost shocks




Note: This figure shows the price of the cheapest plan under price-linked and fixed subsidies for cost shocks of -15% to 15% of baseline. The results are for simulation year 2011.

fixed subsidies, consumer’s premiums increase with prices, so substantially fewer people buy insurance when costs are higher. Depending on how the externality of uninsurance changes with costs, it may be optimal for fewer people to buy insurance when it is more expensive.

Based on the literature discussed in Section 3.4, we consider the cases where a \$1 increase in costs increases the externality of uninsurance by 60¢, 80¢, or \$1. Figure 4 shows public surplus in 2011 under each of these assumptions, again for costs ranging from 85% to 115% of baseline costs. When the externality tracks the cost, the optimal fixed subsidy changes a lot with costs and so if there is a positive cost shock of more than about 15% or a negative cost shock larger than about 12%, the disadvantage of the mis-calibrated subsidy level outweighs the advantage of the lower price and public surplus is lower under a fixed subsidy than price-linked subsidies. If the externality moves less than 1-for-1 with costs, the fixed subsidies do much better. The top panel shows that if the change in externality is 60% of the change in cost then the fixed-subsidy is better than the price-linked subsidy even for cost shocks of > 20%.

The Medicare national health expenditures data indicate that the standard deviation in state-level annual cost growth from 1991-2009 is 1.9 percentage points. Even if costs could only be observed with a three-year lag (implying a standard deviation of 5.7 percentage point if costs are serially uncorrelated), a 12.5% market-level cost shock is more than two standard deviations away from expectations. The uncertainty might be somewhat greater

Figure 3: Share buying insurance under cost shocks




This figure shows the share of consumers who purchase insurance under price-linked and fixed subsidies for cost shocks of -15% to 15% of baseline. The results are for simulation year 2011.

for the narrower population on the exchanges, but we still think 12.5 percentage point shock to the *growth rate* of health care costs is quite high. Political limits on the government's ability to update fixed subsidies over time based on new cost data may justify price-linked subsidies, but our simulations suggest that they are very hard to justify based on purely informational limits.

In 2009 the comparison between price-linked fixed subsidies under cost shocks depends a lot on the equilibrium. In the maximum price equilibrium, fixed subsidies do better for cost shocks up to 20%. In the minimum price equilibrium, price-linked subsidies do better with cost shocks of about 5%. In this best case scenario, when the second cheapest plan is an effective cap on the price of the cheapest plan and the plans do not find the profit maximizing equilibrium, the distortion of price-linked subsidies is smaller. Therefore, it does not require as large a cost shock for the disadvantage of having the subsidy not match the externality to outweigh the price distortion. Somewhat surprisingly, public surplus under the minimum price-linked subsidy equilibrium increases very slightly with cost under the assumption that externalities move one-for-one with costs. It appears that the cap from the second cheapest plan causes prices to increase less than costs; when the externality increases more than the price, the benefit of the market goes up.

Figure 4: Public surplus under cost shocks



Note: This figure shows public surplus under price-linked and fixed subsidies for cost shocks of -15% to 15% of baseline. Each graph corresponds to a different assumption about how much the externality changes with costs – in percentage terms. Public surplus includes the *saved externality*, effectively adding back the direct loss of higher cost to all scenarios so that difference does not swamp the difference between subsidy regimes.

Risk Aversion

Under cost uncertainty, fixed subsidies shift risk from the government to consumers, which is potentially bad for welfare if consumers are risk adverse. However, in this context, allowing for risk aversion has a negligible effect. As described in Section 3.4, we calculate the risk cost of fixed subsidies for each cost shock; it is never more than 15¢. Table 8 shows the public surplus under price-linked and fixed subsidies when the externality increases 80¢ for every \$1 increase in costs. The first two-lines are the same as the middle graph in Figure 4; the last line adjusts the public surplus under fixed subsidies for the risk cost to consumers.⁴⁷ The differences are small. A cost shock of > 15% makes price-linked subsidies barely better than fixed subsidies.

Table 8: Welfare with Risk Aversion

	Cost Shock						
	-0.15	-0.1	-0.05	0	0.05	0.1	0.15
Price-Linked	52.02	50.94	49.91	48.9	47.91	46.95	46
Fixed	52.05	53.28	53.82	53.5	52.25	50.15	47.44
Fixed, Risk-adjusted	51.96	53.23	53.8	53.5	52.24	50.12	47.37

Note: This table shows simulated public surplus under different sized cost shocks for price-linked and fixed subsidies in 2011, assuming that the externality increases 80¢ for every \$1 increase in costs. The last line includes the cost of having risk-adverse consumers bear the price risk.

5 Implications for Subsidized Insurance Markets

Economists tend to see the externalities from people being uninsured as the primary reason to subsidize insurance. Political rhetoric, on the other hand, tends to focus on wanting to make health insurance “affordable” for all. We have not tried to incorporate “affordability” into the welfare metric, but implications for what conditions favor different structures are similar.

We have discussed three possible ways of implementing subsidies:

1. price-linked subsidies with a fixed mandate penalty (the ACA policy),
2. price-linked subsidies also applied to the mandate penalty,

⁴⁷Cost shocks under price-linked subsidies also generate some risk for consumers, because though the affordable amount is fixed, the premiums they pay for plans other than the cheapest ones move with the price differences. However these are negligible – always less than .04¢ (\$.0004) – so we do not include them.

3. fixed subsidies.

Option 1 allows the government to simultaneously take on price risk, guarantee affordability for consumers, and have the subsidy for insurance track the externality of uninsurance; however, it distorts firms' pricing incentives. Options 2 and 3 avoid the pricing distortion, but under uncertainty they transfer price risk (premiums or mandate penalty) to the consumer and do not allow the net incentive for insurance to track the externality. If costs are much lower than expected, fixed subsidies can also push post-subsidy prices to their zero lower-bound.⁴⁸ Under uncertainty, the government has to trade off the cost of the pricing distortion caused by option 1 against the benefit of providing affordability and better calibrating the subsidy to the externality. The government could also do a mix of fixed total subsidy (with either fixed or varying equilibrium mandate penalty) and price-linked subsidies – e.g. set a baseline subsidy amount and have it increase 50¢ on the dollar with the pivotal plan's price.

The more competitive the market or the more uncertainty the government faces, the more likely that the benefits of price-linked subsidies will outweigh the costs. In more competitive markets, the government will want to use the information from market prices because doing so does not distort prices as much.⁴⁹ The more uncertainty the government faces, the more valuable the information from firms' prices is. If the government is primarily concerned about the externality of uninsurance, then it is uncertainty about *health care costs* that matter; if the government is concerned about affordability, then it is uncertainty about the *pricing equilibrium* that matters. Either way, the more information the government has, the more it will want to rely on a fixed total subsidy which does not distort pricing. If variance in the mandate penalty is less politically concerning than variance in prices then our suggestion of applying the subsidy to a higher initial penalty may be optimal.

5.1 ACA, Medicare, and Employer Insurance

Given the available data, we believe that \$21 is a reasonable estimate of the potential distortion of the price-linked subsidies in the Massachusetts CommCare market. However, a couple caveats are in order. CommCare had additional price regulation rules in place, which we do not model, which may have mitigated the effects of this distortion. We look at the early years of the market; if the trend we see towards more elastic demand over time is robust, it may become less of an issue. We also cannot be confident that pricing and entry in the market had reached equilibrium.

⁴⁸Price-linked subsidies also applied to the mandate penalty can push the mandate penalty to its zero lower-bound when costs are higher than expected.

⁴⁹In competitive markets, firms have less pricing power and the outside option tends to be a less important source of competition, so the pricing distortion will generally be lower.

The Massachusetts market has now been converted into an ACA exchange. These exchanges differ in several respects: plans of three generosity tiers, multi-plan insurers, additive risk adjustment, subsidies that depend on the second cheapest silver plan, expansion to higher income groups, inclusion of unsubsidized consumers. The unsubsidized consumers, estimated to be 20% of the market, will mitigate the distortion if they do not all buy bronze level plans. In addition medical loss ratios may be effective in limiting firms' profit margins.

Other factors may exacerbate the distortion. A *New York Times* analysis found that 58% of counties served by the federal exchange, have two or fewer insurers.⁵⁰ While normally the distortion would be capped by the price of the third-cheapest silver plan, counties with one or two insurers may not have a third-cheapest plan (or it may be controlled by the same insurer as the second-cheapest). Multi-plan insurers may also face an additional incentive to distort the price of a plan to raise the subsidy, because their other plans benefit from the increased subsidy. Though if they do not offer a bronze plans and most consumers on the margin of uninsurance chose bronze plans, the price distortion would be mitigated. While we cannot estimate what the distortion in the ACA exchanges, our estimates suggest that it has the potential to be substantial and thus should be an important policy consideration.

The magnitudes will differ, but this potential for distortion exists in any market where (1) firms have some market power⁵¹ and (2) there is the possibility of substitution to an unsubsidized outside option. Price-linked subsidies exacerbate existing market power by effectively removing the competition of the outside option; they cause a larger distortion when firms have more market power and the margin of substitution to the outside option is more important.

Insurers in Medicare Advantage, Medicare Part D, and employer-sponsored insurance programs all have market power and these markets have some substitution to an outside option. In the Medicare Advantage market, which uses fixed baseline subsidies,⁵² our results suggest that switching to price-linked subsidies would create a pricing distortion unless the subsidy was also applied to traditional Medicare (as in "premium support" proposals). If the government does not care whether consumers use traditional Medicare or private plans, price volatility of the outside option is not an issue in this case. In order to avoid the potential for negative prices for traditional Medicare (if the Medicare Advantage bids are high), a

⁵⁰Moreover, about 20% of markets have just one insurer (Abelson et al., 2013). In theory, in single-insurer markets the insurer could raise price arbitrarily without losing subsidized consumers. They will only be limited by medical loss ratio requirements. These areas are disproportionately small and rural, but the distortions there are potentially sizable.

⁵¹The presence of market power is not a restrictive condition; it merely requires that a \$1 price increase does not cause a firm's demand to fall to zero (as would be the case in perfect competition).

⁵²However, after applying this baseline subsidy, Medicare reduces the subsidy when a plan reduces its price below the benchmark. This creates a different kind of pricing distortion, which may be significant.

premium support system would need to count the cost of traditional Medicare as another bid in the market.

Medicare Part D uses flat, price-linked subsidies based on a national enrollment-weighted average of plan price bids. Because all plans' prices affect the subsidy through this average, our theoretical distortion applies to all plans – not just a subset of potentially pivotal silver plans as in the ACA – but the distortion for each plan is smaller. It is approximately proportional to the national market share of the plan's parent insurer, the largest of which is United Health Group with 28% in 2011 (see Decarolis, 2015).⁵³

Employers typically pick a small menu of options for their employees and set subsidies based on prices (either implicitly or explicitly). To the extent that an employer's chosen insurer(s) have market power, this can lead to the same pricing distortion. Since tax rules limit employers ability to subsidize employees' outside options, employers who want to keep prices down should strive to make their subsidies not depend, even implicitly, on insurers' prices.

6 Conclusion

This paper considers the distortion of pricing incentives generated by price-linked subsidies in health insurance exchanges, an important topic for economists analyzing these markets and policy makers designing and regulating them. We highlight this distortion in a simple theoretical model and derive a first-order approximation of its size. We then two natural experiments in the Massachusetts exchange to get structural demand estimates and simulate the market under alternative subsidy policies. In 2011, we find an upward distortion of the subsidy-pivotal cheapest plan's price of \$21 or 5% of the average price of insurance. The potential budgetary effect is substantial: Massachusetts had about 1.9 million member-months of subsidized coverage in fiscal-year 2011, so a \$21 increase in monthly subsidies would translate to \$42 million per year in government costs. For the ACA, using the CBO projection of 20 million subsidized enrollees annually, a \$14 per month subsidy increase would translate to \$5.2billion per year in federal spending.

We do not view these numbers as a precise estimate of either the historical distortion in Massachusetts or the expected distortion in the ACA exchanges. Rather, we think that they indicate that the pricing distortion we identify in theory should be of practical concern. In addition to analyzing the ACA markets as data becomes available, we hope future research will measure the relevant elasticities in Medicare Advantage, Medicare Part D, and employer-sponsored insurance programs, to assess the importance of this pricing distortion in those

⁵³Decarolis (2015) discusses why the distortion from the low-income subsidy is likely larger.

markets.

Price-linked subsidies have advantages. The right tradeoff between the firms' pricing incentives on the one hand and affordability and consumer incentive concerns on the other depends on the level of competition in the market and the precision of the regulator's cost estimates; it will not be the same for every market. We hope our analysis allows for a better understanding of the tradeoffs involved.

References

- Abelson, R., Thomas, K., and McGinty, J. C. (2013). Health care law fails to lower prices for rural areas. *New York Times*, October 23, 2013.
- Aizawa, N. (2016). Health insurance exchange design in an empirical equilibrium labor market model.
- Baicker, K. and Staiger, D. (2005). Fiscal shenanigans, targeted federal health care funds, and patient mortality. *Quarterly Journal of Economics*, 120(1):345–386.
- Berry, S., Levinsohn, J., and Pakes, A. (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68.
- Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and welfare in health plan choice. *The American Economic Review*, 102(7):3214–3248.
- Chan, D. and Gruber, J. (2010). How sensitive are low income families to health plan prices? *The American Economic Review*, 100(2):292–296.
- Chandra, A., Gruber, J., and McKnight, R. (2011). The importance of the individual mandate – evidence from Massachusetts. *New England Journal of Medicine*, 364(4):293–295.
- Chetty, R. (2006). A new method of estimating risk aversion. *The American Economic Review*, 96(5):1821–1834.
- Commonwealth Care (2008). Report to the Massachusetts Legislature: Implementation of the health care reform law, chapter 58, 2006–2008.
- Currie, J. and Gahvari, F. (2008). Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, pages 333–383.
- Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2014). Can health insurance competition work? evidence from Medicare Advantage. Technical report, National Bureau of Economic Research.
- Dafny, L., Gruber, J., and Ody, C. (2015). More insurers lower premiums: Evidence from initial pricing in the health insurance marketplaces. *American Journal of Health Economics*, 1(1):53–81.
- Decarolis, F. (2015). Medicare Part D: Are insurers gaming the low income subsidy design? *The American Economic Review*, 105(4):1547–1580.

- Decarolis, F., Polyakova, M., and Ryan, S. P. (2015). The welfare effects of supply-side regulations in Medicare Part D. *NBER Working Paper 21298*.
- Einav, L., Finkelstein, A., Cullen, M. R., et al. (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877–921.
- Enthoven, A. C. (1988). *Theory and Practice of Managed Competition in Health Care Finance*. North-Holland, Amsterdam.
- Ericson, K. M. and Starc, A. (2015). Pricing regulation and imperfect competition on the massachusetts health insurance exchange. *Review of Economics and Statistics*, 97(3):667–682.
- Finkelstein, A., Hendren, N., and Luttmer, E. F. (2015). The value of medicaid: Interpreting results from the oregon health insurance experiment. *NBER Working Paper 21308*.
- Garthwaite, C., Gross, T., and Notowidigdo, M. J. (2015). Hospitals as insurers of last resort. *NBER Working Paper 21290*.
- Gruber, J. (2008). Covering the uninsured in the United States. *Journal of Economic Literature*, 46(3):571–606.
- Hackmann, M. B., Kolstad, J. T., and Kowalski, A. E. (2015). Adverse selection and an individual mandate: When theory meets practice. *The American Economic Review*, 105(3):1030–1066.
- Handel, B., Hendel, I., and Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83(4):1261–1313.
- Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *The American Economic Review*, 103(7):2643–2682.
- Liu, Y. and Jin, G. Z. (2015). Employer contribution and premium growth in health insurance. *Journal of Health Economics*, 39:228–247.
- Mahoney, N. (2012). Bankruptcy as implicit health insurance. *NBER Working Paper 18105*.
- Mahoney, N. (2015). Bankruptcy as implicit health insurance. *The American Economic Review*, 105(2):710–746.
- Marzilli Ericson, K. M. (2014). Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64.
- Ruggles, S., Genadek, K., Goeken, R., Grover, J., and Sobek, M. (2015). Integrated public use microdata series: Version 6.0 machine-readable database.
- Shepard, M. (2016). Hospital network competition and adverse selection: Evidence from the massachusetts health insurance exchange. *NBER Working Paper 22600*.
- Shleifer, A. (1985). A theory of yardstick competition. *The RAND Journal of Economics*, pages 319–327.
- Tebaldi, P. (2016). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA.
- Werden, G. J. (1996). A robust test for consumer welfare enhancing mergers among sellers of differentiated products. *Journal of Industrial Economics*, 44(4):409–413.

Appendix

A Theory with Multi-Plan Insurers (ACA Case)

In the ACA insurers must offer a plan in each of multiple tiers – bronze, silver, gold, and platinum.⁵⁴ Subsidies are set equal to the price of the second-cheapest silver plan minus a pre-specified “affordable amount.”⁵⁵ If an insurer does not offer a bronze plan *and* consumers on margin of uninsurance mostly pick bronze plans, that will mitigate the distortion. More generally, though, the fact that insurers are providing additional plans provides an even greater incentive for an insurer to increase the price of its silver plan because the higher subsidy increases demand for the insurer’s non-silver plans as well – again by inducing more customers to enter the market.

Suppose each firm j offers plans in tiers $l = \{(B)ronze, (S)ilver, (G)old, (P)latinum\}$. Assuming no adverse selection or perfect risk adjustment, the insurer maximizes profits:

$$\max_{P_{j1} \dots P_{jL}} \sum_l (P_{jl} - c_{jl}) Q_{jl}(P^{cons}, M),$$

where $P_{jl}^{cons} = P_{jl} - S(P)$ with $S(P) = P_{2nd,S} - \text{AffAmt}$. Following the same steps as in the text, the first-order condition for the silver plan is:

$$\frac{\partial \pi_j}{\partial P_{jS}} = Q_{jS}(\cdot) + \sum_l (P_{jl} - c_{jl}) \frac{dQ_{jl}}{dP_{jS}} = 0.$$

The markup with fixed subsidies is:

$$Markup_{jS}^F = \frac{1}{\eta_{jS}} + \frac{1}{-\frac{\partial Q_{jS}}{\partial P_{jS}^{cons}}} \sum_{l \neq \text{silver}} (P_{jl} - c_{jl}) \frac{\partial Q_{jl}}{\partial P_{jS}^{cons}}$$

The second term reflects the fact that the insurer captures revenue from consumers who switch to its other plans when it raises the price of its silver plan. This is a standard effect in settings with multi-product firms. However, with the subsidy set based on the second

⁵⁴Platinum plans cover 90% of medical costs (comparable to a generous employer plan today); gold covers 80% of costs; silver covers 70% of costs; and bronze covers 60% of costs. Consumers with incomes below 250% of poverty also receive so called “cost-sharing subsidies” that raise the generosity of silver plans.

⁵⁵This subsidy is applied equally to all plans (with a cap ensuring that no premium is pushed below \$0), ensuring that at least two silver plans (and likely some bronze plans) cost less than the affordable amount for low-income consumers.

cheapest silver plan, the markup for that subsidy-pivotal plan is:

$$Markup_{jS}^{PLink} = \frac{1}{\eta_{jS} - \eta_{jS,M}} + \frac{\sum_{l \neq S} (P_{jl} - c_{jl}) \left(\frac{\partial Q_{jl}}{\partial P_{jS}^{cons}} + \underbrace{\frac{\partial Q_{jl}}{\partial M}}_{\text{Additional Distortion}} \right)}{\left(\frac{\partial Q_{jS}}{\partial P_{jS}^{cons}} + \underbrace{\frac{\partial Q_{jS}}{\partial M}}_{\text{Additional Distortion}} \right)}. \quad (8)$$

The fact that other plans offered by the firms also gain some of the consumers driven into the market by the additional subsidy generates an additional distortion.

How much larger the distortion is in the multi-product ACA case is not certain, and we do not have data to credibly estimate its size. We discuss some of the issues in translating our estimates for Massachusetts to the ACA case in Section 5.

B Optimality of fixed subsidies under known costs and logit demand

Suppose the government sets both a fixed component of the subsidy (which could be zero) and decides to what extent it will depend on the price of cheapest plan. So the total subsidy is $S = S_0 + \alpha P_{\underline{j}}$. We assume the fixed part is set optimally and analyze how α affects welfare. Using

$$\frac{dW}{dS_0} = \frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial S} \frac{\partial W}{\partial P_j} = 0$$

we get

$$\begin{aligned} \frac{dW}{d\alpha} &= \frac{\partial W}{\partial S} \cdot P_{\underline{j}} + \sum_j \frac{\partial P_j}{\partial \alpha} \frac{\partial W}{\partial P_j} \\ &= \sum_j \frac{\partial W}{\partial P_j} \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_{\underline{j}} \right) \\ &= - \sum_j D_j \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_{\underline{j}} \right) \end{aligned} \quad (9)$$

Calculating how each firm changes its price in response to a change in S or α requires inverting a $J \times J$ matrix, but the weighted sum can be calculated in closed form.

The firms' first order conditions are

$$\begin{aligned} (p_j - c_j)(1 - D_j - \alpha D_0) &= 1 & j = \underline{j} \\ (p_j - c_j)(1 - d_j) &= 1 & j \neq \underline{j} \end{aligned}$$

Differentiating with respect to S gives

$$\frac{\partial p_{\underline{j}}}{\partial S}(1 - D_{\underline{j}} - \alpha D_0) = (p_{\underline{j}} - c_{\underline{j}}) \cdot \left((D_{\underline{j}} + \alpha D_0) \sum_k \frac{\partial p_k}{\partial S} D_k + \left(\frac{\partial D_{\underline{j}}}{\partial p_{\underline{j}}} - D_{\underline{j}}^2 \right) \frac{\partial p_{\underline{j}}}{\partial S} + D_{\underline{j}} D_0 - \alpha D_0 (1 - D_0) \right)$$

$$\begin{aligned} \frac{\partial p_{\underline{j}}}{\partial S} \left((1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0) \right) &= (D_{\underline{j}} + \alpha D_0) \sum_k \frac{\partial p_k}{\partial S} D_k - D_0(\alpha - \alpha D_0 - D_{\underline{j}}) \\ \frac{\partial p_{\underline{j}}}{\partial S} \left((1 - D_{\underline{j}})^2 + D_{\underline{j}} \right) &= D_{\underline{j}} \sum_k \frac{\partial p_k}{\partial S} D_k + D_{\underline{j}} D_0 \quad j \neq \underline{j}. \end{aligned}$$

Taking the wighted sum gives

$$\begin{aligned} \sum_j \frac{\partial p_j}{\partial S} D_j &= \sum_{j \neq \underline{j}} \frac{D_j}{(1 - D_j)^2 + D_j} \left(D_j \sum_k \frac{\partial p_k}{\partial S} D_k + D_j D_0 \right) \\ &\quad + D_{\underline{j}} \frac{D_{\underline{j}} + \alpha D_0}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \sum_k \frac{\partial p_k}{\partial S} D_k - D_{\underline{j}} \frac{D_0(\alpha - \alpha D_0 - D_{\underline{j}})}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)}. \end{aligned}$$

Rearranging

$$\begin{aligned} \sum_j \frac{\partial p_j}{\partial S} D_j &\left(1 - \sum_{j \neq 1} \frac{D_j}{(1 - D_j)^2 + j} - \frac{D_{\underline{j}}(D_{\underline{j}} + \alpha D_0)}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \right) \\ &= \sum_{j \neq 1} \frac{D_j^2 D_0}{(1 - D_j)^2 + D_j} - \frac{D_{\underline{j}} D_0 (\alpha - \alpha D_0 - D_{\underline{j}})}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)}. \quad (10) \end{aligned}$$

Differentiating the first-order conditions with respect to α gives

$$\begin{aligned} \frac{\partial p_{\underline{j}}}{\partial \alpha} (1 - D_{\underline{j}} - \alpha D_0)^2 &= \frac{\partial D_{\underline{j}}}{\partial \alpha} + \alpha \frac{\partial D_0}{\partial \alpha} + D_0 + \sum_k \frac{\partial p_k}{\partial \alpha} \left(\frac{\partial D_{\underline{j}}}{\partial p_k} + \alpha \frac{\partial D_0}{\partial p_k} \right) \\ \frac{\partial p_{\underline{j}}}{\partial \alpha} \left((1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0) \right) &= p_{\underline{j}}(D_{\underline{j}} D_0 - \alpha D_0(1 - D_0)) + D_0 + (D_{\underline{j}} + \alpha D_0) \sum_k \frac{\partial p_k}{\partial \alpha} D_k \\ \frac{\partial p_{\underline{j}}}{\partial \alpha} \left((1 - D_{\underline{j}})^2 + D_{\underline{j}} \right) &= D_{\underline{j}} \sum_k \frac{\partial p_k}{\partial \alpha} D_k + P_{\underline{j}} D_{\underline{j}} D_0 \end{aligned}$$

Taking the wighted sum gives

$$\begin{aligned} \sum_j \frac{\partial p_j}{\partial \alpha} D_j &= \sum_{j \neq 1} \frac{D_j}{(1 - D_j)^2 + D_j} \left(D_j \sum_k \frac{\partial p_k}{\partial \alpha} D_k + P_{\underline{j}} D_j D_0 \right) \\ &\quad + D_{\underline{j}} \frac{D_{\underline{j}} + \alpha D_0}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \sum_k \frac{\partial p_k}{\partial \alpha} D_k + D_{\underline{j}} \frac{D_0(1 + p_{\underline{j}}(D_{\underline{j}} + D_0 - 1))}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \end{aligned}$$

Rearranging

$$\begin{aligned} \sum_j \frac{\partial p_j}{\partial \alpha} D_j &\left(1 - \sum_{j \neq 1} \frac{D_j^2}{(1 - D_j)^2 + D_j} - \frac{D_{\underline{j}}(D_{\underline{j}} + \alpha D_0)}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \right) \\ &= \sum_{j \neq 1} \frac{D_j^2}{(1 - D_j)^2 + D_j} p_{\underline{j}} D_0 + \frac{D_{\underline{j}} D_0 (1 + p_{\underline{j}}(D_{\underline{j}} + D_0 - 1))}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)}. \end{aligned} \quad (11)$$

Plugging (10) and (11) into (9) gives

$$-\sum_j D_j \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_{\underline{j}} \right) = -\frac{\frac{D_{\underline{j}} D_0}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)}}{1 - \sum_{j \neq 1} \frac{D_j^2}{(1 - D_j)^2 + D_j} - \frac{D_{\underline{j}}(D_{\underline{j}} + \alpha D_0)}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)}} < 0,$$

where the inequality follows from

$$\begin{aligned} 1 - \sum_{j \neq 1} \frac{D_j^2}{(1 - D_j)^2 + D_j} - \frac{D_{\underline{j}}(D_{\underline{j}} + \alpha D_0)}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \\ > 1 - \sum_{j \neq 1} \frac{D_j^2}{D_j} - \frac{D_{\underline{j}}(D_{\underline{j}} + \alpha D_0)}{(1 - D_{\underline{j}} - \alpha D_0)^2 + D_{\underline{j}}(1 - \alpha D_0)} \\ &> 1 - \sum_{j \neq 1} D_j - (D_{\underline{j}} + D_0) = 0. \end{aligned}$$

Since welfare is decreasing in α , it is better to have fixed subsidies $\alpha = 0$ than price-linked subsidies $\alpha = 1$.

C Data details

For age, we restrict to ages 19-64; seniors are in Medicare and low-income children are eligible for Medicaid. For income, we restrict to those with household income $\leq 300\%$ of poverty. We define households as ‘‘health insurance units,’’ a variable included on the IPUMS ACS that is intended to approximate the household definition used by public insurance programs. We exclude non-citizens because most are ineligible, and the rare exception (long-term green card holders) cannot be measured. We exclude the uninsured who are likely eligible for Medicaid (rather than CommCare) – parents up to 133% poverty and the disabled. We correct for

the fact that the ACS is a sample by weighting ACS observations by their “person weight” – the Census-defined factor for scaling up to a population estimate

Although this lets us simplify the demand model, it creates an additional complication in combining the ACS and CommCare data. While we can differentiate new and existing consumers in the CommCare data, the ACS only lets us observe the total *stock* of uninsured in each year. To convert this into a comparable flow of “new uninsured,” we assume that the ratio of flow-to-stock among the eligible uninsured equals the ratio for CommCare enrollees, which we can estimate in our data. This assumption would hold exactly, for instance, if the market were in steady state and choice of uninsurance were uncorrelated with the likelihood of transitioning out of eligibility.⁵⁶ More generally, it can be seen as an approximation to the true, more complicated process. Once CommCare reaches a steady-state size (by the start of 2009), we estimate a fairly constant annual flow-to-stock ratio of 78%. To apply this to our data, we rescale all weights for ACS observations by a multiple of 0.78.

D Natural Experiments

This section provides more background on the natural experiments we use and details of the estimation as well as some robustness checks.

D.1 Mandate Penalty Introduction

We estimate excess new enrollments in December 2007-March 2008 relative to the trend in nearby months, using enrollment trends for people earning less than poverty as a control group. We estimate the effect through March 2008 for two reasons. First, the application process for the market takes some time, so people who decided to sign up in January may not have enrolled until March. Second, the mandate rules exempted from penalties individuals with three or fewer months of uninsurance during the year, meaning that individuals who enrolled in March avoided any penalties for 2008. However most of the effect is in December and January, so focusing on those months does not substantially affect our estimates.

We collapse the data to the income group-month level and calculate the new enrollees in the cheapest plan for each group and month, normalized by the same plan’s total enrollment for the income group in June 2008. Before running the specification described in the text, we separate out the effect for each month December 2007 through March 2008:

$$NewEnroll_{g,t} = \alpha_0 + \sum_{m \in DM} (\beta_{0m} + \gamma_{0m} \cdot Treat_t) \cdot 1_m + \delta_0 \cdot X_t + \left(\alpha_1 + \sum_{m \in T} (\beta_{1m} + \gamma_{1m} Treat_t) \cdot 1_m + \delta_1 \cdot X_t \right) T_g + \varepsilon_{g,t}$$

⁵⁶Formally, suppose n individuals become newly eligible each year, while a fraction ϕ of eligibles from the previous year lose eligibility. Then letting N_t be the stock of eligibles in year t , we know that $N_t = N_{t-1}(1 - \phi) + n$. In steady state, $N_t = N_{t-1} = n/\phi$, implying a flow-to-stock ratio of $n/N_t = \phi$. If the likelihood of choosing uninsurance is uncorrelated with entry/exit rates from eligibility, then this ratio will hold for both CommCare enrollees and the CommCare-eligible uninsured.

where DM is the months December through March, 1_m is an indicator for month m , $Treat_t$ is a dummy for the treatment period, X_t is a vector of time polynomials and CommCare-year dummies,⁵⁷ and T_g is a dummy for the treatment group. The difference-in-difference coefficients of interest are γ_{1m} 's.

Table 9: Introduction of the Mandate Penalty.

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment			
	(1)	(2)	(3)
150-300% Poverty x Dec2007	0.112*** (0.003)	0.110*** (0.006)	0.103*** (0.007)
x Jan2008	0.073*** (0.004)	0.067*** (0.006)	0.069*** (0.006)
x Feb2008	0.043*** (0.005)	0.033*** (0.006)	0.033*** (0.006)
x Mar2008	0.025** (0.006)	0.027*** (0.006)	0.020*** (0.007)
Total	0.253*** (0.017)	0.237*** (0.023)	0.225*** (0.024)
Control Group (< 100% Poverty) Dummies for Dec-March		X X	X
Observations	51	102	102
R-Squared	0.969	0.923	0.925

Robust s.e. in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

NOTE: This table performs the difference-in-difference regressions analogous to the graphs in Figure 1. The dependent variable is the number of new CommCare enrollees who choose the cheapest plan in each month in an income group, scaled by total group enrollment in that plan in June 2008. There is one observation per income group (the 150-300% poverty treatment group, plus the <100% poverty control group in columns (2) and (3)) and month (from April 2007 to June 2011). All specifications include CommCare-year dummy variables⁵⁸ and fifth-order time polynomials, separately for the treatment and control group, to control for underlying enrollment trends. (The CommCare-year starts in July, so these dummies will not conflict with the treatment months of December to March.) Specification (3) also includes dummy variables for all calendar months of December-March for the treatment group, to perform the triple-difference. See the note to Figure 1 for the definition of new enrollees and the cheapest plan.

Table 9 presents the regression results. Column (1) starts with a baseline single-difference specification that estimates based only on enrollment for the 150-300% poverty group in December 2007-March 2008 relative to the surrounding months. Column (2) then adds the <100% poverty group as a control group, to form the difference-in-difference estimates. Finally, Column (3) adds dummies for December-March in all years, forming the triple difference that nets out general trends for those months in other years. Despite the relatively

⁵⁷The CommCare-year starts in July, so these dummies will not conflict with the treatment months of December to March.

small number of group-month observations, all the relevant coefficients are highly significant. Summing across treatment months, the total excess enrollment after the mandate penalty introduction was between 22-25% of equilibrium market size.

Table 10 takes the final, triple-difference specification from Table 9 and breaks the analysis down into narrower income groups (by 50% of poverty interval, the narrowest we have). The coefficients are a little larger for the higher income groups – about 25% instead of 21% – who faced higher mandate penalties. Given the mandate penalties of \$17.50-\$52.50, the coefficients imply that each \$1 increase in the mandate penalty raised demand by between 1.2% for the 150-200% of poverty group and 0.48% for the 250-300% of poverty group, with a weighted average of 0.97%.

Table 10: Introduction of the Mandate Penalty

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment			
	Income Group (% of Poverty Line)		
	150-200%	200-250%	250-300%
Income group x Dec2007	0.101*** (0.007)	0.113*** (0.007)	0.091*** (0.008)
x Jan2008	0.061*** (0.007)	0.085*** (0.007)	0.086*** (0.007)
x Feb2008	0.026*** (0.007)	0.045*** (0.006)	0.051*** (0.006)
x Mar2008	0.020** (0.008)	0.020*** (0.006)	0.022*** (0.007)
Total	0.208*** (0.026)	0.263*** (0.022)	0.250*** (0.024)
Observations	102	102	102
R-Squared	0.927	0.922	0.920

Robust s.e. in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

NOTE: This table runs the same analysis as Column (3) in Table 9, broken down by income group. See the footnote in that table for a description of the control variables.

We would like to interpret the increases in enrollment as being the result of the \$17.50–\$52.50 monthly mandate penalty that went into effect in January 2008. However, the 2007 uninsurance penalty of \$219 was assessed based on coverage status in December 2007,⁵⁹ making that month’s effective penalty much larger. If consumers were only buying insurance because of the larger December penalty, we would expect many of them to leave the market soon after the monthly penalty dropped to the lower level. To assess this story, Figures 5a and 5b plot the probability of market exit within 1 and 6 months, respectively, for new enrollees. Each point represents a distinct group of new enrollees in the corresponding month

⁵⁹In addition, an exemption was given for individuals who applied for CommCare in 2007 and enrolled on January 1, 2008. However, December 31, 2007, was the main advertised date for assessing coverage status.

shown on the x-axis, and the y-axis value is the share who exited within 1 or 6 months. While there is a general upward trend over time, there is no jump in either series for people who joined in December 2007. The large spike for people enrolling in March 2008 is due to an unrelated income verification program.⁶⁰ This analysis suggests that consumers were not enrolling for just December to avoid the \$219 penalty and leaving soon after because of the lower penalty.

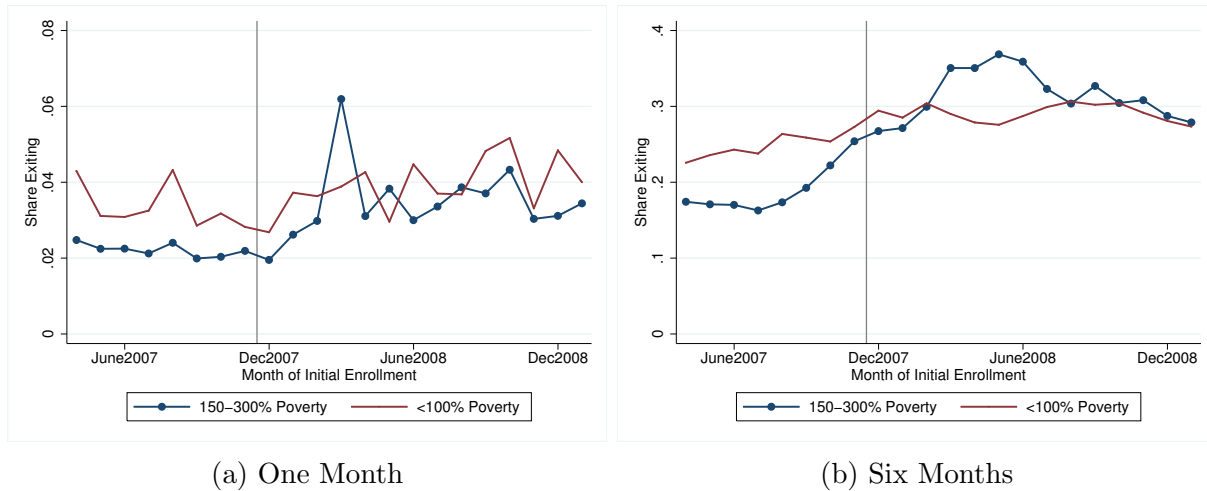


Figure 5: Share of New Enrollees Exiting Within the Specified Number of Months.

NOTE: These graphs show the rate of exiting CommCare coverage within one month (left figure) and six months (right) of initial enrollment among people newly enrolling CommCare in a given month. The spike among new enrollees in March 2008 reflects the start of an income-verification program for the 150-300% poverty group in April 2008. See the note to Figure 1 for the definition of new enrollees and the cheapest plan.

D.2 Affordable Amount Decrease Experiment

This natural experiment addresses a concern with our first method: that the *introduction* of a mandate penalty may have a larger effect (per dollar of penalty) than a *marginal increase* in penalties. Some individuals may obtain coverage to avoid the stigma of paying a penalty, but this stigma might not change when mandate penalties increase. An argument against the stigma explanation is that the legal mandate to obtain insurance had been in place since July 2007, though without financial enforcement. In addition, individuals below 100% of poverty were also required to obtain insurance (again without financial enforcement). However to the extent there is a stigma specifically from paying a fine for non-coverage, this concern is valid.

The most significant changes in the affordable amount occurred in July 2007 for consumers between 100-150% of poverty. Prices in CommCare usually change in July, but in

⁶⁰The income verification program took effect in April 2008 for individuals above 150% of poverty (but earlier on for people <100% poverty). Prior to April, income group at enrollment was based partly on self-reporting, and changes in income over time were also supposed to be self-reported. The verification program uncovered a large number of ineligible people, who were dis-enrolled in April 2008 and subsequent months. This event can also explain the upward trend in exits within 6 months leading up to April 2008.

the first year of the program, prices were held fixed from November 2006 to June 2008, so firms' price-bids did not change at the same time as this change in the affordable amount. For the first half of 2007, their affordable amount was \$18 and premiums ranged from \$18 for the cheapest plan to \$74.22 for the most expensive. In July 2007, CommCare eliminated premiums for this group, so all plans became free. We can think of this as the combination of two effects: (1) The affordable amount was lowered from \$18 to zero, and (2) the premium of all plans besides the cheapest one were differentially lowered to equal the cheapest premium (now \$0). The second change should unambiguously lower enrollment in what was the cheapest plan, since the relative price of all other plans falls. Therefore, if we use the actual policy change to estimate the effect on new enrollment in the (formerly) cheapest plan, this will be a lower bound on the effect of just lowering the affordable amount (the effect we want to estimate).

As a control group, we use the 200-300% poverty group, whose affordable amounts were essentially unchanged in July 2007. The affordable amount for consumers 200-250% poverty was unchanged at \$70, while the amount for 250-300% poverty was lowered by just \$1 from \$106 to \$105. To the extent this slightly increased enrollment for the control group, it would only bias our estimates downward. We exclude other incomes from our controls for several reasons. First, we do not use the below 100% poverty group because of its somewhat different enrollment history and trends. Whereas the groups above poverty only started joining CommCare in February 2007, the below 100% poverty group became eligible in November 2006 and had a large influx in early 2007 due to an auto-enrollment. Second, we exclude the 150-200% of poverty group from the controls because their affordable amount also fell non-trivially (from \$40 to \$35) in July 2007. While this smaller change does not produce as dramatic of an enrollment spike, we show in Appendix ?? that their enrollment increase was consistent with semi-elasticity calculated from the the mandate penalty introduction results presented in Table 10.

Table 11 presents the regression results. Again, in Column (1) we just look at the enrollment difference for the 100-150% poverty treatment group relative to trend, captured by CommCare-year dummies and time polynomials. Column (2) then adds the 200-300% poverty control group to form the difference-in-difference estimates. The last column does a triple-difference, further netting out changes in July-October of other years. The coefficients change a bit more between specifications, but they all imply significant enrollment increases of at least 15 percentage points. Our preferred specification of the triple-difference implies a semi-elasticity of .98%.

We use the smaller changes in the affordable amount that occurred for the 150-200% poverty group in July 2007 as a check on the semi-elasticity estimated for this group. Table 12 shows the results. Our preferred triple-difference estimate from column (3), indicates that the \$5 affordable amount decrease translated into a 6.3% increase in demand for the cheapest plan. This implies a semi-elasticity of 0.0126, quite similar to the semi-elasticity of 0.0119 that we found for this group using the mandate penalty introduction.

Table 11: Decrease in the Affordable Amount.

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment			
	(1)	(2)	(3)
100-150% Poverty x July2007	0.064*** (0.005)	0.060*** (0.008)	0.063*** (0.008)
x Aug2007	0.052*** (0.005)	0.031*** (0.008)	0.035*** (0.009)
x Sep2007	0.022*** (0.005)	0.022 (0.008)	0.016* (0.008)
x Oct2007	0.062** (0.005)	0.054*** (0.008)	0.055*** (0.008)
Total	0.200*** (0.021)	0.156*** (0.031)	0.169*** (0.032)
Control Group (200 – 300% Poverty) Dummies for July-October		X X	X
Observations	52	104	104
R-Squared	0.988	0.981	0.981

Robust s.e. in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

NOTE: This table performs regression analysis of the decrease in the affordable amount with specifications analogous to those in Table 9. The dependent variable is the number of new CommCare enrollees who chose the cheapest plan in each month in an income group, scaled by total group enrollment in that plan in June 2008. There is one observation per income group (the 100-150% poverty treatment group, and the 200-300% poverty control group in columns (2) and (3)) and month (from March 2007 to June 2011). All specifications include CommCare-year dummy variables and fifth-order time polynomials, separately for the treatment and control group, to control for underlying enrollment trends. (The first CommCare year ended in June 2008, so there is no conflict between the CommCare-year dummies and the treatment months of July to October 2007.) Specification (3) also includes dummy variables for all calendar months of July-October for the treatment group, to perform the triple-difference. Where applicable, specifications also include dummy variables to control for two unrelated enrollment changes: (a) for 100-150% poverty in December 2007, when there was a large auto-enrollment spike, and (b) for 200-300% poverty in each month from December 2007 to March 2008, when there was a spike due to the introduction of the mandate penalty. See the note to Figure 1 for the definition of new enrollees and the cheapest plan.

Table 12: Decrease in the Affordable Amount: 150-200% Poverty

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment			
	(1)	(2)	(3)
150-200% Poverty x July2007	0.013*** (0.005)	0.017** (0.007)	0.018** (0.008)
x Aug2007	0.029*** (0.005)	0.016** (0.007)	0.022*** (0.008)
x Sep2007	0.013** (0.005)	0.011* (0.006)	0.015** (0.007)
x Oct2007	0.008 (0.005)	0.009 (0.006)	0.009 (0.006)
Total	0.063*** (0.019)	0.052** (0.025)	0.063** (0.027)
Control Group (200 – 300% Poverty) Dummies for July-October		X X	X
Observations	52	104	104
R-Squared	0.960	0.957	0.958

Robust s.e. in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

NOTE: This table performs the difference-in-difference regressions analogous to those in Table 11 (see the note to that table for additional information), but with a treatment group of enrollees 150-200% poverty, whose affordable amount dropped from \$40 to \$35 in June 2008. The control group is enrollees 200-300% poverty, whose affordable amounts were essentially unchanged at that time.

E Structural Model Details

Let θ refer to all the parameters to be estimated. Given logit errors, the plan choice probabilities are

$$P(j_i^* = j | Z_{it}, \nu_i, \theta) = \frac{\exp(\tilde{u}_{ij})}{\sum_{k=0}^J \exp(\tilde{u}_{ik})}$$

We estimate the model by simulated method of moments, incorporating micro moments with an approach similar to Berry et al. (2004). For each individual i (with their associate Z_i) we draw a ν_i .

For each ξ we have a moment for the corresponding plan and group, g (either region-year or region-income group) that matches the observed share of consumers in that group who chose that plan, s_j^{Obs} , to the expected share given θ . If n_g is the number of individuals in group g , we have

$$F_{j,g}^1(\theta) = s_{j,g}^{Obs} - \frac{1}{n_g} \sum_{i \in g} Pr(j_i^* = j | \theta, Z_i, \nu_i)$$

For each β there is also a corresponding group, h – income, demographic, region, or

year. We use moments analogous to those above for the share of uninsured, separately by age-gender groups, income groups, regions, and year; however, because the uninsured data are based on relatively small samples in the ACS, we do not interact these categories. The corresponding moments are

$$G_{0,h}^1(\theta) = s_{0,h}^{Obs} - \frac{1}{n_h} \sum_{i \in h} Pr(j_i^* = j | \theta, Z_i, \nu_i).$$

We also match the covariance of plan premium and individual attributes. Following Berry et al. (2004), we use

$$G^2(\theta) = \frac{1}{n} \sum_j \sum_i P_{ij}^{Cons} Z_i (1\{j_i^{Obs} = j\} - Pr(j_i^* = j | \theta, Z_i, \nu_i))$$

where $P_{i,0}^{Cons} = M_i$. This helps us identify the different price-elasticity parameters.

The final set of moments helps identify the variance of the random coefficients by matching the estimated insurance demand response from the natural experiments discussed in Section 2.1. If there is substantial heterogeneity in the value of insurance, the uninsured will tend to be people with very low idiosyncratic values of insurance; since they are not close to the margin of buying coverage, an increase in the mandate penalty will not increase their demand for insurance very much. Thus, higher values of σ are likely to generate less demand response to the mandate penalty, and vice versa. We match the simulated change for the cheapest plan to observed 22.5% change in demand:

$$G^3(\theta) = \sum_i ((1 + 22.5\%) Pr(j_i^* = j_{min} | Z_i, \nu_i, \theta, M_i^{Pre}) - Pr(j_i^* = j_{min} | Z_i, \nu_i, \theta, M_i^{Post})).$$