

Price-Linked Subsidies and Health Insurance Markups*

Sonia Jaffe[†]

Mark Shepard[‡]

May 3, 2017

Abstract

Subsidies in many health insurance programs depend on prices set by competing 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 weaken price competition, leading to higher markups and subsidy costs for the government. The size of the price distortion depends on how demand responds to a change in the “price” of uninsurance, which we estimate in Massachusetts using natural experiments in the subsidized health insurance exchange. Reduced-form calculations suggest prices were 4-11% higher with price-linked subsidies than they would have been with fixed subsidies.

We next analyze the welfare effects of these price changes to see why governments may chose to use price-linked subsidies. We show that that price-linked subsidies make sense only if (1) there is uncertainty about costs/prices, and (2) *optimal* subsidies increase as prices rise. We argue that it may be optimal for health insurance subsidies to rise with prices because doing so both indirectly links subsidies to market-wide shocks affecting the cost of “charity care” used by the uninsured and insures consumers against cost risk. Structural estimation allows us to simulate market equilibria under different policies and cost shocks. Without a cost shock, price-linking increases prices by up to 6% relative to fixed subsidies. For levels of cost uncertainty that are reasonable in a mature market, we find that the losses from higher prices outweigh the benefits 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, Neale Mahoney, Luca Maini, Tom McGuire, Ariel Pakes, Daria Pelech, Maria Polyakova, Jonathan Skinner, Amanda Starc, Pietro Tebaldi and seminar and conference participants at Boston University, Harvard economics department, Harvard Medical School, University of Chicago economics department, the NBER Health Care Program, 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.

[‡]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: Medicare Part D links subsidies to market average prices; the ACA exchanges link subsidies to the second-cheapest plan in the “silver” tier. Other programs, however, set “fixed subsidies” independently of prices – at specific levels or based on external benchmarks not controlled by insurers: Medicare Advantage sets benchmarks that determine subsidies based on an area’s lagged costs in traditional Medicare. Shifting towards fixed, inflation-indexed subsidies for insurance is also a key feature of proposed Republican alternatives to the ACA. This choice of subsidy type is also relevant to other programs such as school vouchers and housing subsidies, where the government subsidizes individuals’ purchases from private providers.

While linking subsidies to prices is convenient when there is uncertainty about market prices, it also raises concerns about competition. In the extreme case, with a monopoly insurer (or if insurers collude), price-linked subsidies allow large price increases without any loss of demand, since subsidies would increase in tandem and consumer prices would remain unchanged. We provide a simple formula to approximate the price increase from price-linked subsidies and estimate the relevant parameters using data from Massachusetts’ pre-ACA subsidized exchange. We then use a structural model to analyze the welfare effects of the price differences and allow for the benefits of price-linked subsidies under uncertainty.

First, 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. The intuition for price-linked subsidies incentivizing higher prices is simple: if a higher price yields a higher subsidy for the firm *relative* to other options in the market, the firm has

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 charity 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.

an incentive to raise its price. Even in markets where subsidies are constant across plans,³ as in 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 that may affect the subsidy. While fixed subsidies can affect prices by shifting the demand curve, they do not *rotate* the demand curve in this way that clearly distorts prices upwards.⁴

Our model of equilibrium insurance exchange pricing gives a simple first-order approximation of this distortion that depends only on the semi-elasticity of demand with respect to own price and the semi-elasticity of demand with respect to the “price” of uninsurance (and potentially on adverse selection). Since price-linked subsidies effectively remove 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). These conditions seem particularly applicable to the exchanges set up by the ACA. In federally facilitated ACA exchanges in 2014, there were an average of just 3.9 insurers, and approximately 30% of consumers had only one or two insurers to choose from (Dafny et al., 2015).⁵ Furthermore, the outside option of uninsurance is likely to be an important factor, since a large share of the subsidy-eligible population in both Massachusetts (as we show) and the ACA (Avalere Health, 2015) remain uninsured.

We draw on two natural policy experiments in Massachusetts’ pre-ACA subsidized insurance exchange to estimate the semi-elasticity of demand with respect to the mandate penalty. An important model for the ACA, the Massachusetts market lets us observe insurance demand and costs for a similar setting and low-income population. Using unaffected income groups as a control, we study the effect on demand of the state’s introduction of a mandate

³As noted in Enthoven (1988) and Liu and Jin (2015), 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 75% of the price of each plan, meaning that each \$1 price increase raises a plan’s subsidy by \$0.75).

⁴Our analysis suggests a simple alternate policy that could preserve price-linked subsidies while mitigating the distortion: link the mandate penalty to prices in the *opposite* direction, thus making the *total* incentive to buy insurance (subsidy plus penalty) invariant to prices. Specifically, if the price of the pivotal plan exceeds an expected level, the mandate penalty would be reduced by the difference (and vice versa if prices were lower than expected). Under certain conditions, this would eliminate the distortion from price-linking. However, if there is substitution on other margins – such as changing jobs to an employer that offers insurance (Aizawa, 2016) – this will not completely eliminate the distortion.

⁵The joint federal-state nature of the program (subsidies are federally funded but insurers are state-regulated) also raises concerns that states have an incentive to coordinate pricing (collusion) to bring additional federal subsidy dollars into their states – though we hasten to add that there is no evidence that this has yet occurred. This effect is similar to the “fiscal shenanigans” concerns that have been documented for federal matching funds in Medicaid (Baicker and Staiger, 2005).

penalty and of a subsidy increase that lowered premiums of all plans in the exchange (but left the penalty for uninsurance unchanged). The results are similar: each \$1 increase in the relative price of uninsurance raised insurance demand by about 1%.⁶ Using own-price elasticities from either Chan and Gruber (2010) or the structural model we estimate subsequently, our first-order approximation estimates that price-linked subsidies caused the monthly price of the cheapest plan to be \$17-\$45 (4-11%) higher.

Second, to allow for strategic interactions and adverse selection, and to better address the welfare questions discussed below, we estimate a structural model of demand and costs. To estimate the model, we supplement the administrative plan enrollment and claims data from Massachusetts' pre-ACA exchange with data on the uninsured from the American Community Survey; this allows us to estimate the share of eligible enrollees who purchase insurance. 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 match the model's predicted effect of the mandate penalty on insurance demand to the observed effects of the policy changes described above. This quasi-experimental variation identifies the variance in the random component of individuals' utility of insurance, which is key for the effect of the subsidy policy on prices.

The estimated demand and cost parameters allow us to calculate insurers' profit functions; we use these to simulate Nash pricing equilibria under price-linked subsidies and under fixed subsidies – with the fixed subsidy amounts set equal to the realized values under price-linking. Across several simulation years and assumptions, we find a non-trivial upward distortion in the price of the cheapest plan (to which Massachusetts' subsidies are linked) of \$4-26 per month, or 1-6% of baseline prices. By switching to different levels of the fixed subsidy, the government could achieve the same coverage at 6.1% lower taxpayer cost, or achieve 1.3% greater coverage at the same cost, relative to price-linked subsidies. Although modest, these effects imply meaningful increases in government costs. For instance, the \$24/month subsidy distortion (in our simulations for 2011) would translate into \$46 million in annual subsidy costs for Massachusetts, and over \$3 billion if extrapolated nationally to the ACA.

We also find that the distortion could be much larger in markets with few competitors, as in many ACA exchanges. In simulations where we exogenously remove all but two insurers, we find price increases that range from 6-12% of baseline prices, depending on the identity of the two competitors. The largest distortions occur when the gap between the costs of the

⁶Our results are consistent with past work that studies the effect of Massachusetts' mandate penalty on coverage, including Chandra et al. (2011) and Hackmann et al. (2015). We compare our estimates to theirs in Section 3.

two insurers is large – i.e., if a low-cost plan competes against a high-cost plan.

Lastly, to convert these price changes to welfare effects and to see why governments may choose price-linked subsidies, we define a public objective function. The social planner cares about consumer surplus, government costs, a negative “externality” of uninsurance – an important part of the rationale for subsidizing insurance (as we discuss below) – and (down-weighted) insurer profits. 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 leading to higher prices compared to fixed subsidies. Even if insurer profits are given full weight, the higher prices for insurers are accompanied by lower enrollments under price-linked subsidies, making any gains to insurer profits smaller than the loss to consumers and the government. 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.

The potential benefits of price-linking arise only if two conditions hold: (1) the government is uncertain about market costs or prices and (2) the *optimal* subsidy is larger in states of the world where prices are unexpectedly high. If the government has enough information to predict the pricing equilibrium or the optimal subsidy is independent of prices, the government can just set a fixed subsidy at the optimal level, thereby avoiding the pricing distortion from price-linked subsidies.

Acknowledging that government may be uncertain about insurance prices, we propose two reasons why the optimal subsidy may increase with prices. First, an explicit rationale for the ACA’s subsidy design was guaranteeing that post-subsidy premiums be “affordable” regardless of insurance prices; price-linking lets the government insure low-income enrollees against higher than expected prices. Second, a large component of the externality that the uninsured impose on society is the cost of “charity care” – typically borne by hospitals and public clinics.⁷ The regulator wants the subsidies to track this externality. Price-linking provides an indirect way to achieve this if both charity care costs and insurance prices rise with market-level health care costs. Note, this requires that insurers have more information on costs than regulators do when setting subsidies; otherwise, prices have no additional signal value.⁸

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

⁸The extent to which insurers are better informed than the regulator is an open question. States 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. In our empirical simulations, we consider a range of cost shocks to gauge the sensitivity of our conclusions.

We simulate the market equilibrium under each subsidy policy for a range of cost shocks – observed by insurers but not by the regulator – that shift both insurer costs and the uninsurance externality. As expected, when there is zero cost shock, fixed subsidies result in higher welfare. But as costs diverge from the regulator’s expectations, the gap narrows, and with a large enough shock, price-linked subsidies may do better. Intuitively, price-linked subsidies automatically rise/fall in response to the cost shock, whereas fixed subsidies cannot adjust and are set at an increasingly sub-optimal level. We also find that price-linked subsidies stabilize the insurance coverage rate, so that even if costs are much higher than expected, coverage falls very little.⁹

However, we find that cost shocks need to be quite large before price-linking improves welfare. Even under the assumptions most favorable to price-linking, fixed subsidies do better for all cost shocks between -12.5% and +15%. Although shocks larger than this are reasonable in a new market (e.g., the ACA in its first years), they seem less relevant in a mature market where (as is true in practice) regulators have lagged cost information from insurer rate filing reports. In this case, the relevant uncertainty is about the cost *growth* rate over time. As a rough benchmark, the standard deviation of state-level cost growth from 1991-2009 (in the National Health Accounts data) was 1.9% points for annual growth, and 4.8% points even over three years. This makes cost shocks larger than 12.5% an extreme tail outcome. Of course, if political economy constraints prevent regulators from updating fixed subsidies based on recent cost data – e.g., if subsidies are indexed by statute to an inflation rate that is too low – the automatic adjustment inherent in price-linking could be more attractive.

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, particularly under uncertainty. Most closely related is concurrent work by Tebaldi (2016), who studies California’s ACA exchange and considers fixed subsidies (or “vouchers”) as a counterfactual. While Tebaldi focuses on the specifics of the ACA context and the benefits of age-specific subsidies, we analyze the conceptual and welfare tradeoffs of price-linked subsidies and their performance under cost uncertainty.

Our paper also relates to a broader literature estimating equilibrium under imperfect

⁹This is an important property for analyzing future coverage under the ACA’s price-linked setup: even if insurers exit and prices rise, coverage among the subsidized population is unlikely to fall because *post-subsidy* premiums will be stabilized.

competition in health insurance markets. This includes previous work on the Massachusetts insurance exchanges (Ericson and Starc, 2015, 2016; Shepard, 2016), and Medicare insurance markets (e.g., Town and Liu, 2003; Starc, 2014; Curto et al., 2014; Polyakova, 2016; Ho et al., 2015), and other settings (e.g., Handel, 2013; Handel et al., 2015; Ho and Lee, 2013). We discuss the broader implications of our analysis of price-linked subsidies for other health insurance settings in Section 6.

The remainder of the paper is structured as follows. Section 1 uses a simple model to show our theoretical analysis of price-linked subsidies. Section 2 describes our Massachusetts setting and data. Section 3 presents evidence from natural experiments in the market and calculates the first-order approximation of the price difference. Section 4 describes the structural model, presents the estimated demand and cost parameters, and simulates the pricing equilibria under each subsidy policy. Section 5 presents the welfare framework and the externality of uninsurance; it then presents the estimated welfare differences with and without uncertainty. Section 6 discusses the tradeoffs among subsidy structures and the broader implications of our findings. Section 7 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 for our data, where each insurer offers a single plan. We discuss how the pricing distortion might differ with multi-plan insurers (as in the ACA) in Section 6; the formulas are in 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.¹⁰ Subsidy-eligible consumers then choose which (if any) plan to purchase based on plan attributes and post-subsidy prices, $P_j^{cons} = P_j - S(P)$. If consumers choose the outside option of uninsurance, 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. 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.

¹⁰The 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 we do allow for income-specific subsidies in our empirical model.

Costs vary across insurers and enrollees. We allow for risk selection by letting a plan's average costs, $\bar{c}_j(P^{cons}, M)$, depend on prices, which affect the set of consumers who select a plan. The exchange uses a risk adjustment transfer, $\bar{\phi}_j(P^{cons}, M)$, to compensate plans based on the measured sickness of its enrollees (which also varies with prices). The insurer's net (risk-adjusted) costs equal:

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

The insurer profit function is:

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

A necessary condition for Nash equilibrium is that each firm's first-order condition is satisfied:¹¹

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

This differs from standard oligopoly pricing conditions in two respects. First, the $\frac{\partial c_j^{Net}}{\partial P}$ term allows for selection to influence pricing, a standard consideration in insurance markets. Second, 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) 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

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

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 premiums) and the mandate penalty.

We can simplify this formula by imposing an assumption that is standard in most discrete choice models: that (at least locally) price enters the utility function linearly. This implies that only *price differences*, not levels, matter for demand.¹² Thus, raising all prices (and

¹¹These 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.

¹²This assumption is typically justified by the fact that prices are a small share of income. Although we study a low-income population, post-subsidy premiums in our setting are just 0-5% of income, and price differences are even smaller. In an insurance setting, linear-in-price utility can be seen as a transformed

the mandate penalty) by \$1 is simply a lump-sum transfer that 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 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 or if few additional 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 estimate $\partial Q_j / \partial M$.

Under price-linked subsidies, the “Price-Linking” term in (3) is positive, as we formalize below. Since under standard assumptions $\partial Q_j / \partial M$ is also positive, the adjustment term in (3) is positive. This diminishes the (negative) slope of the demand curve, making *effective* demand less elastic, which increases equilibrium markups.¹³

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 shift the insurance demand

approximation to a CARA utility function, in which risk aversion is constant with income.

¹³If subsidies were negatively linked to prices (a negative “Price-Linking” term), then policy could make effective demand more elastic, lowering markups. This is an interesting possibility, but not one that we have seen implemented in practice.

curve – which can affect equilibrium markups, as shown by Decarolis et al. (2015), but they do not *rotate* the demand curve. Markups are:

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

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). If, as was the policy in Massachusetts, the regulator wants to ensure that the cheapest plan’s post-subsidy premium equals an (income-specific) “affordable amount,” regardless of its pre-subsidy price, then

$$S(P) = \min_j P_j - \text{AffAmt}$$

so $\frac{\partial S(P)}{\partial P_{\underline{j}}} = 1$, where \underline{j} is the index of the pivotal (cheapest) plan.

Demand for the pivotal plan is effectively 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:

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

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

1.2 Comparing Fixed and Price-Linked Subsidies

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.¹⁴

Like much of the related literature, we assume that the market reaches equilibrium where firms effectively know each other’s prices, 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 (e.g., 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 distortionary effects on other plans’ prices.

¹⁴This can create a range of possible equilibria with a tie among multiple cheapest plans, an issue we address in our simulations.

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:

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

which is positive under standard assumptions.¹⁵ 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.¹⁶ We can also express the difference as a percent change in markup; if adverse selection is constant over the relevant price range then

$$\frac{Mkup_{\underline{j}}^{\text{PLink}} - Mkup_{\underline{j}}^{\text{F}}}{Mkup_{\underline{j}}^{\text{F}}} = \frac{\eta_{j,M}}{(\eta_{\underline{j}} - \eta_{j,M})}. \quad (5)$$

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* (in addition to the insurance plans) so that:

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

The key feature of this policy is that $\frac{\partial M}{\partial P_j} = -\frac{\partial S}{\partial P_j}$, so the “Price-Linking” term in Equation (3) equals zero. As a result, subsidy policy does not diminish the slope of the demand curve. 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 it holds fixed $S+M$, the *net* public incentive for consumers to buy insurance. Since plans’ pricing cannot impact this net incentive for insurance, the distortion of price-linking goes away. (Fixed subsidies do the same by holding *both* S and M fixed with prices.) Of course, this result relies on the assumption that only the net incentive, $S + M$, matters for insurance demand, not the level of S and M individually. If

¹⁵Both $\eta_{j,M} = \frac{1}{Q_j} \frac{\partial Q_j}{\partial M}$ and $\eta_{\underline{j}} = -\frac{1}{Q_j} \frac{\partial Q_j}{\partial P_j^{\text{cons}}}$ are positive under standard demand assumptions, and $\eta_{\underline{j}} > \eta_{j,M}$ as long as there is at least one other in-market option besides \underline{j} . In any equilibrium where adverse selection is not sufficiently bad to unravel the market, $\frac{\partial c^{\text{Net}}}{\partial P_j} < 1$. This expression only applies as long as the higher markup does not push the pivotal plan’s price above the next-cheapest.

¹⁶This 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.

some consumers were unaware of the penalty or could avoid paying it (e.g., by applying for a religious or hardship exemption), this assumption would not hold perfectly. In our empirical work, we focus on comparing fixed and price-linked subsidies, rather than this alternate policy.

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 facilitated and subsidized coverage for individuals earning less than 300% of the federal poverty line (FPL) who do not have 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.¹⁷

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.

Administrative data from the CommCare program¹⁸ 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 few non-employer insurance markets with plan choice and cost data linked at the individual level.

In CommCare, subsidies are linked to the price of the cheapest plan so that this plan costs an income-specific “affordable amount.”¹⁹ 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

¹⁷By 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.

¹⁸This 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.

¹⁹In 2009-2011, our main years of analysis, the affordable amount was \$0 for <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.

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.²⁰ 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 Census²¹ – to get a dataset of CommCare-eligible individuals, whether or not they chose to purchase insurance. 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. We use data from January 2008, when the individual mandate is fully phased in, to June 2011, just prior to the start of CommCare year 2012,²² when plan choice rules and market dynamics shifted considerably (see Shepard, 2016). See Appendix B for more information on the data and sample construction.

We focus on active plan choices made by new enrollees.²³ This lets us abstract from inertia known to affect plan switching (Handel, 2013; 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

²⁰In 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. Some eligible consumers may have had access to employer insurance that was deemed “unaffordable” (based on the employer covering less than 20%/33% of the cost of family/individual coverage). 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.

²¹We obtained ACS data from the IPUMS-USA website, Ruggles et al. (2015), which we gratefully acknowledge.

²²Because 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.

²³“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.

way of abstracting from the complex dynamics that inertia creates. However, for additional precision and to better match average costs in the market, we estimate our cost model using all enrollees (both new and current).

Table 1 shows summary statistics for three samples: the full sample, used for cost estimation; the new enrollees, used for demand estimation, and the subsample of individuals in above 100% of poverty, used for equilibrium simulations (to match the ACA exchanges' eligible population). The raw sample includes 455,556 CommCare enrollees and 4,562 uninsured from the ACS. The population is quite poor, with about half having family income less than the poverty line. Consumers' ages range from 19-64; the uninsured are slightly younger than the insured.

Our simulations focus on the population above poverty, where 45.5% of eligible individuals are uninsured. While this estimate may seem high, recall that CommCare (like the ACA) is targeted at the subset of the population without other insurance options.²⁴ Of those who enroll, about 44% choose the cheapest plan; pre-subsidy monthly prices average \$399, but subsidies are quite large. Enrollees pay an average of \$46/month; the cheapest plan is only \$35/month for above poverty consumers. We estimate the above-poverty uninsured pay an average mandate penalty of \$25. (Since the insured are somewhat poorer, the average mandate penalty they would have paid is lower, \$16.) Predicted costs for the uninsured are somewhat lower than observed (and predicted) costs for the insured, consistent with the uninsured being a slightly healthier population.

3 Reduced-form Estimation

The first-order approximation for the change in price depends on the semi-elasticity of demand for the cheapest plan with respect to the mandate penalty and with respect to its own price. We use natural experiments in the Massachusetts market to estimate these two numbers.

3.1 Response to the Mandate Penalty

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 to the price of the outside option. In this section, we show the basic patterns graphically and report the main estimates. Details are in Appendix B.

²⁴For 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

	CommCare			ACS (Uninsured)	
	All Cost Est	New Enrollees Demand Est	Above Poverty Simulation	New Enrollees Demand Est	Above Poverty Simulation
Counts					
Unique Individuals	455,556	326,033	253,200	.	.
Sample size	.	.	.	4,562	2,225
Avg per month	161,871	10,679	82,906	8,531	69,084
Demographics					
Age	39.7	37.6	42.6	35.1	36.6
Male	47.2%	48.5%	42.2%	66.9%	62.3%
Income < Poverty	48.8%	51.3%	0.0%	47.0%	0.0%
100% – 200%	38.2%	35.2%	74.6%	29.4%	55.4%
200% – 300%	13.0%	13.4%	25.4%	23.6%	44.6%

	CommCare	ACS (Uninsured)
<i>Simulation Sample (Above Poverty)</i>		
Share	54.5%	45.5%
Cheapest Plan Market Share	43.9%	
Pre-subsidy Prices		
Average Plan	\$399	
Cheapest Plan	\$378	
Consumer Premiums		
Average Plan	\$46.44	
Cheapest Plan	\$35.09	\$52.27
Mandate Penalty	\$16.49	\$24.83
<i>Costs Sample (All Enrollees)</i>		
Observed Cost	\$372.81	
Predicted in Avg Plan	\$373.24	\$366.77

NOTE: The top panel shows counts and attributes for different subsets of eligible consumers (CommCare enrollees and the ACS uninsured) from January 2008 to June 2011. The second panel summarizes the prices and costs. “Pre-subsidy Prices” shows the enrollment-weighted average and the cheapest monthly price paid to firms for an enrollee. “Consumer Premiums” are the (enrollment-weighted) average 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.

Mandate Penalty Introduction Experiment

Our first strategy uses the mandate penalty’s introduction. Under the Massachusetts health reform, a requirement to obtain insurance took effect in July 2007. However, this 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 depended on income and ranged from \$17.50 (for 150-200% poverty) to \$52.50 (for 250-300% poverty). People earning less than 150% of poverty did not face a penalty.

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 that plan’s total enrollment in that income group in June 2008.²⁵

Several pieces of evidence suggest that this enrollment spike was caused by the financial penalties. 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),²⁶ and there was no enrollment spike for individuals above 150% of poverty in December-March of other years. 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 specification, analogous to the graph in Figure 1. With one observation per month, t , and enrollee income group, g ($< 100\%$ poverty and 150-300% poverty), we estimate

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

where the dependent variable is new enrollment divided by that group’s enrollment in June

²⁵Enrollment, which had been steadily growing since the start of CommCare, stabilizes around June 2008, so we use June 2008 enrollment as an estimate of equilibrium market size.

²⁶People earning 100-150% of poverty are omitted from the control group 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.

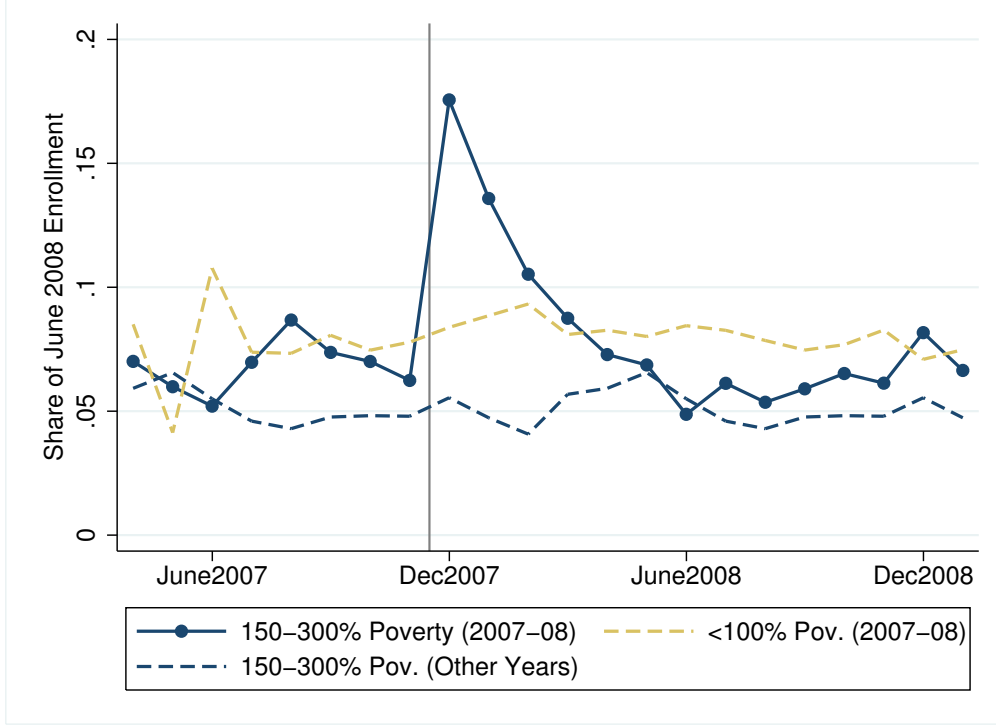


Figure 1: New Enrollees in Cheapest Plan by Month

NOTE: This figure shows 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.

2008, $MandIntro_t$ is a dummy for the mandate penalty introduction 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 year dummies (for CommCare’s market year, which runs from July-June), and $Treat_g$ is a dummy for the treatment group (150-300% poverty). 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. In Appendix B, we report full regression results, including robustness checks. Translating this increase into a semi-elasticity of demand, we find that coverage increases by 0.97% on average for each \$1 increase in the penalty. As we describe in Section 4, we match the 22.5% increase in enrollment as a moment in the structural model.

Our estimates are consistent with past work studying the introduction of the mandate in Massachusetts. Chandra et al. (2011) also study the CommCare market and find similar

results, though they focus on the effects of the mandate on adverse selection rather than the net increase in coverage. Hackmann et al. (2015) study the introduction of the mandate for the unsubsidized, higher-income population, who face a higher mandate penalty (\$83-105 per month). They find a slightly larger increase in coverage for this population (an increase of 26.5% points, or 37.6% relative to baseline coverage), but the implied semi-elasticity of demand is lower, as one would expect for a higher-income group.

Premium Decrease Experiment

As a robustness check for the effects measured from the introduction of the mandate penalty, we use an increase in subsidies in July 2007 that lowered the premiums of all plans for enrollees earning between 100-200% of poverty. Enrollees above 200% poverty, whose premiums were essentially unchanged at this time, serve as a control group. A decrease in all plans' premiums has an equivalent effect on *relative* prices as an increase in the mandate penalty, so this change gives us another way to estimate our statistic of theoretical interest. This approach also addresses a potential concern with the mandate penalty experiment – that the *introduction* of a mandate penalty may have a larger effect (per dollar) than a *marginal increase* in the relative price of uninsurance.

We present details and results for this experiment in Appendix B. The results are quite similar to the mandate penalty introduction. We again find that each \$1 increase in the relative price of uninsurance (i.e., \$1 decrease in plan premiums) raises insurance demand by about 1%. In particular, the estimated semi-elasticity for the 150-200% poverty group (the only group affected by both changes) is nearly identical for the two natural experiments. The similarity across two different changes gives us additional confidence in the validity of the results.

3.2 Pricing distortion

The difference in markups between fixed and price-linked subsidies depend on the semi-elasticity of demand with respect to the mandate penalty, (which we estimate to be .095) and the semi-elasticity of demand for the cheapest plan with respect to its own price. We consider a range of possible values for this latter number. Chan and Gruber (2010) analyze this market and report an own-price semi-elasticity of demand of 0.0154, but they do not allow for substitution to the outside option of uninsurance. When we adjust for that, their estimates imply an overall semi-elasticity of .020.²⁷ The structural analysis we do later gives an own-price semi-elasticity for Celticare of .0291. Plugging these into Equation (4), we get

²⁷They find a price coefficient of -.027. The in-market share of the cheapest plan is 47.3% and 54.5% of eligible enrollees buy insurance. Under a logit model (or any model with independence of irrelevant alternatives) we have $\eta_j = -\alpha \cdot (1 - s_j) = 0.027 \cdot (1 - 0.473 \cdot .545) = 0.0200$. If we use the reported elasticity of .0154, the estimated price distortion is \$47

that if the exchange’s risk-adjustment affectively addresses adverse selection, the difference in markups is between

$$\begin{aligned} Markup_{\underline{j}}^{\text{PLink}} - Markup_{\underline{j}}^{\text{F}} &= \frac{\eta_{\underline{j},M}}{\eta_{\underline{j}}(\eta_{\underline{j}} - \eta_{\underline{j},M})} = \frac{0.0095}{0.0291(0.0291 - 0.0095)} \approx \$17, \\ &= \frac{0.0095}{0.020 \cdot (0.020 - 0.0095)} \approx \$45. \end{aligned}$$

A price change of between \$17 and \$45 corresponds to a percent change in price of 4-11%. The percent change in markup is obviously much larger. Using Equation (5), we have,

$$\begin{aligned} \frac{Markup_{\underline{j}}^{\text{PLink}} - Markup_{\underline{j}}^{\text{F}}}{Markup_{\underline{j}}^{\text{F}}} &= \frac{\eta_{\underline{j},M}}{(\eta_{\underline{j}} - \eta_{\underline{j},M})} = \frac{0.0095}{(0.029 - 0.0095)} \approx 48\%, \\ &= \frac{0.0095}{(0.020 - 0.0095)} \approx 90\% \end{aligned}$$

suggesting price-linked subsidies generated markups that were 48%-90% higher than what they would have been under fixed subsidies. Adverse selection that is not accounted for by the exchange’s risk adjustment ($\frac{\partial c^{\text{Net}}}{\partial P_j} > 0$) would decrease the dollar differences, though not the percents.

4 Structural Model and Estimation

There are two limitations to the reduced-form estimation implemented above. The first is that it does not allow for strategic interactions across plans; by looking only at the change in incentives for the cheapest plan, it does not consider how firms may respond to the change in the cheapest plan’s price and how it may respond to their response. The second is that it only estimates a price difference, not the welfare effect of that difference. To address both of these issues, we estimate a cost function and demand system using the dataset of CommCare enrollees and eligible uninsured individuals described in Section 2. This model allows us to simulate the market under different subsidy policies – incorporating firms’ strategic interactions – and provides the inputs necessary for the welfare framework that we describe in the next section.

4.1 Demand

Model

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) in which 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

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

where P_{yj}^{cons} is the plan's post-subsidy premium for consumer i (which depends on income, y_i), $\xi_j(Z_i)$ is plan quality, and ϵ_{ij} is an i.i.d. type-I extreme value error giving demand its logit form. Price sensitivity varies with income and demographics: $\alpha(Z_i) = \alpha_y + \alpha_d$. Plan quality dummies 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

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

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 disutility 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 similar to those used by Berry et al. (2004): (group-specific) market shares and interactions of average chosen plan characteristics and consumer observables. These moments identify the coefficients in $\alpha(Z_i)$ and $\xi_j(Z_i)$, as well as the non-random coefficients in β .

We address the concern of correlation between price and unobserved plan quality by using *within-plan* premium variation created by exchange subsidy rules. (The the same variation as the reduced-form estimation in Section ??.) Since all plans are free for below-poverty enrollees while higher-income enrollees pay higher premiums for more expensive plans, there are differential premium *changes* over time, which we use for identification as in Shepard (2016).²⁸ For instance, when a plan increases its price between years, its

²⁸The discussion that follows closely follows that of Shepard (2016).

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 3.1. 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 B shows details of the method and formulas of all moments.

Estimated Parameters

The demand coefficients are summarized in Table 2.²⁹ The average price coefficient ranges from -0.046 for those just above the poverty line to -0.023 for those making 250%-300% FPL. Price sensitivity decreases with age, making the oldest group less than half as sensitive as the youngest. On average, consumers prefer Fallon to BMC to Network Health to NHP to CultiCare, with the biggest difference being between CultiCare and the other plans. The utility of uninsurance (β) is much lower for females and decreasing with age, consistent with the fact that females and older people are less likely to be uninsured. We allow the value of uninsurance to vary with observable and unobservables. Both generate large variance in the value of uninsurance, but unobservables play a larger role, leading to a standard deviation of 0.91 across individuals, relative to the 0.67 standard deviation driven 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. We find that each \$1/month increase in a plan’s consumer premium lowers its demand by an average of 2.4%, which increases over time because of the entry of CultiCare in 2010. This is quite large, reflecting the price-sensitivity of this low-income population.³⁰ Converting this semi-

²⁹To conserve space and ease interpretation, we report a summary of the coefficients in both the demand and cost models (see Section 4). A full set of coefficient estimates is available on request.

³⁰Our estimate is somewhat larger than the 1.5% estimate reported by Chan and Gruber (2010) for CommCare in an earlier period – even after adjusting their number to allow for substitution to uninsurance, which they do not consider. Our results may differ because we allow for heterogeneity in price coefficients

Table 2: Summary of Parameters in Demand Model

(a) Premium Coefficient			(b) Plan and dummies		
	Avg	S.E.		Avg	S.E.
100-150 % Pov	-0.046***	0.0054	CeltiCare	-0.857***	0.0564
150-200 % Pov	-0.029***	0.0061	NHP	-0.067***	0.0147
200-250 % Pov	-0.028***	0.0072	Network	0.126***	0.0108
250-300 % Pov	-0.023**	0.0010	BMC	0.169***	0.0312
5yr Age Bin	.0019***	0.0006	Fallon	0.209***	0.0563
Female	.0028**	0.0014			

(c) Uninsurance Value				
	Avg	S.E.	Total Standard Deviation	1.14
5yr Age Bin	-0.066***	0.019	From Observables	0.67
Female	-.814***	0.188	From Unobservables (σ)	0.91

Note: The table summarizes the utility coefficients entering our demand model. Panel (a) reports premium coefficients by income and how price-sensitivity varies with demographics. Panel (b) gives the average exchange plan dummies (with the average plan normalized to 0); these are weighted averages of the full set of plan-region-year and plan-region-income group dummies in the model. The bottom panel shows how the utility of uninsurance varies by demographics and how much of the variation is due to unobservables.

Table 3: Average Semi-elasticities

Own Price Semi-Elasticity				Semi-Elasticity of Insurance w.r.t. Mandate Penalty	
By plan		By Year		By Year	
CeltiCare	-2.91%	2008	-2.28%	2008	1.09%
NHP	-2.67%	2009	-2.27%	2009	1.23%
Network Health	-2.44%	2010	-2.48%	2010	1.28%
BMC	-2.14%	2011	-2.55%	2011	1.29%
Fallon	-2.69%	All	-2.41%	All	1.24%

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. Average semi-elasticities vary across years both because demand parameters vary (e.g., plan dummies) and because of changes over time in enrollee demographics, participating plans, and market shares.

elasticity into a rough “insurer-perspective” elasticity by multiplying times the average price (\$386/month) yields an elasticity of -9.3, which is larger than what has been found in employer-sponsored insurance (see discussion in Ho, 2006). However, multiplying times the much lower average consumer premium (\$48 for above-poverty enrollees) yields a more modest “consumer-perspective” elasticity of -1.2.

4.2 Costs

Model

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

$$E(c_{ijt}) = \exp(\mu X_{it} + \psi_{0,t} + \psi_{j,r,t}). \quad (6)$$

Costs vary with consumer income and demographics ($X_i = \{y_i, d_i\}$), the year ($\psi_{0,t}$), and a region-year specific plan effect, $\psi_{j,r,t}$, which are normalized to average zero each year. 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.³¹ 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 maximum likelihood estimation of Equation (6) is that estimates of $\psi_{j,r,t}$ will be biased by selection on unobserved sickness. This is particularly relevant because X_i 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 and a sample limited to new and re-enrollees. These estimates are identified based only on within-person cost variation when an individual leaves the market and later re-enrolls in a different plan (e.g., because plan prices have changed).³² We then adjust observed costs by removing the

by income and demographics, and we also use a control group (below-poverty enrollees, for whom plans are free) to deal with the endogeneity of prices to unobserved quality.

³¹The 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).

³²This method eliminates selection if individuals’ unobserved risk factors are stable over time or uncorrelated with plan changes. It will, however, miss selection on risk changes – e.g., if individuals who get sicker between enrollment spells select into different plans. We have experimented with using instrumental

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.³³

Estimated Parameters

The first part of Table 4 shows averages of plan cost effects (i.e., $\exp(\psi_{j,r,t}) - 1$), broken down by before and after 2010 when CultiCare entered. The numbers reported are normalized so that the share-weighted average cost effect in each year is 0. Prior to 2010, both Network Health and BMC had similar cost effects – about 7% below average, with other plans somewhat higher. When CultiCare entered, it became the clear low-cost plan – 32% below average. Cost effects of the other plans changed somewhat, but their ordering did not.

Table 4: Cost Parameters

	Plan Effects				
	CultiCare	Network Health	BMC	Fallon	NHP
2008-2010		-7.27%	-6.57%	+9.99%	+15.29%
2010-2011	-32.14%	-8.04%	-1.93%	+7.68%	+13.73%
	Percent of Federal Poverty Line				
	100-150	150-200	200-250	250-300	
Relative to < Poverty	-23.4%	-18.2%	-25.2%	-29.5%	

Note: The table shows average plan cost effects, which give the percent difference between the expected cost for a given consumer under that plan and the average plan. The plan parameters, $\psi_{j,r,t}$, are estimated as described in Section 4.2; the reported percentages are averages of $\exp \psi_{j,r,t} - 1$, normalized so that the share-weighted average is zero. The reported percentages for income groups 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 (18-30%) 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 sensible – costs increase with age and are higher for females at young ages and higher for males at older ages.

variables to deal with the selection problem but have not found an approach that works.

³³In addition to insurer medical costs captured by this model, insurers incur administrative costs for functions like claims processing. Using plan financial reports, we estimate variable administrative costs of approximately \$30 per member-month. We add this to our cost function, but, following Massachusetts rules, we assume that insurers are paid an equal-size (\$30) administrative fee (not subject to risk adjustment). Thus, administrative costs cancel out in the profit function.

4.3 Equilibrium

We use these demand and cost models to simulate equilibrium under different subsidy policies. For our simulation population, we limit our sample to consumers above poverty, which matches the criteria for subsidy eligibility in the ACA exchanges. We also use the mandate penalty and affordable amounts from the ACA, though we find similar results if we use the Massachusetts rules. However, we use the Massachusetts plans and price-linking rule (linking to the cheapest plan), since we do not have a way to estimate demand across gold/silver/bronze tiers as in the ACA.

To mitigate adverse selection, both Massachusetts and the ACA risk-adjust the payments to insurers. We incorporate this by estimating a risk score for each individual in each year, ϕ_{it} , that indicates how costly they are expected to be relative to the average enrollee. Dropping the t subscripts for convenience, a plan with price P_j receives $\phi_i P_j$ for enrolling consumer i . We assume that ϕ_i perfectly captures the individual (non-plan) portion of costs: $\phi_i = \exp(\hat{\mu}X_i + \psi_0) / \bar{c}$, where \bar{c} is the average of $\exp(\hat{\mu}X_i + \psi_0)$ across individuals who buy insurance.³⁴ Thus, risk adjustment fully individuals' expected costs in a proportional way: net revenue on consumer i ($= \phi_i P_j - E(c_{ij})$) simplifies to $\phi_i (P_j - \bar{c} \exp(\psi_{j,r}))$.³⁵ This form of risk-adjustment addresses selection on observables *across* plans in the exchange. However, it does not mitigate the effects of adverse selection *into the market*, since the exchange cannot transfer money away from people who never enroll.

Using our demand and cost models, insurer profits for a given year equal:

$$\pi_j = \sum_i \phi_i (P_j - \bar{c} \exp(\psi_{j,r})) \cdot Q_{ij} (P^{Cons}(P))$$

where $P^{Cons}(P)$ is the subsidy function mapping prices into consumer premiums. We assume that each insurer sets its price to maximize profits in static, full information, 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. For each year and subsidy policy, we simulate the 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 insurers' FOCs (as shown in Equation (1)) is a total derivative that incorporates any effect of changing P_j on the subsidy if plan j is pivotal (cheapest) under

³⁴In principle, we could also model imperfect risk adjustment by assuming that $\phi_i^{Imperfect} = (\phi_i)^\gamma$ 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.

³⁵We 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} .

price-linked 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 $\frac{dQ_j}{dP_j} = \frac{\partial Q_j}{\partial P_j}$). 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 at which the two plans may tie in equilibrium. In our simulations this occurs in 2009, so we show results for both the minimum and maximum prices consistent with this range of equilibria.

4.4 Simulations

We compare the equilibrium under price-linked subsidies to the outcome under a policy where the subsidy for each income group is fixed equal to the subsidy amount that emerged under the price-linked policy. We do this separately for 2009 and 2011, to illustrate the very different competitive dynamics 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 (the one we expect the regulator would have chosen). 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, monthly prices for the cheapest plans (BMC and Network Health) are between \$4 and \$26 (1-6%) higher under price-linked subsidies than under fixed subsidies. For CeltiCare in 2011, the price is \$24 (6%) higher under price-linked subsidies. These changes are slightly larger than the \$17 we calculated using the model’s semi-elasticity and the first-order approximation of the distortion. We think the mostly likely explanation is the interactions with other firms pricing incentives: their response to Celticare’s price increase causes Celticare to raise its price even more. The change in the average price is smaller than the change in the cheapest price because the other insurers do not change their prices as much between fixed and price-linked subsidies. Because prices are higher while (income-specific) subsidies are held constant,³⁶ fewer people buy insurance under price-linked subsidies.

³⁶Changes in the income composition of the insured cause the small change in the average subsidy.

Table 5: Equilibrium under Price-Linked and Fixed Subsidies

	Subsidy Type	Min Price	Average Price	Average Subsidy	Share Insured
2009	Min Price-linked	\$415.3	\$418.6	\$354.4	57.6%
	Max Price-linked	\$437.3	\$441.2	\$375.8	58.5%
	Fixed	\$411.7	\$416.6	\$354.2	58.3%
	Min P-link - Fixed	\$3.5	\$2.0	\$0.2	-0.7%
	Max P-link - Fixed	\$25.6	\$24.6	\$21.7	0.2%
2011	Price-Linked	\$403.8	\$427.3	\$345.7	40.0%
	Fixed	\$379.9	\$405.4	\$346.7	45.9%
	Price-link - Fixed	\$23.8	\$21.9	-\$1.0	-5.9%

Note: This is a comparison of the market equilibria under price-linked and fixed subsidies. In 2009, there are a range of equilibria under price-linked subsidies; this reports statistics for the equilibrium with the minimum and maximum cheapest price. The average price is weighted by plan share. The average subsidy is across income groups. Changes in the income composition of the insured cause the small change in the average subsidy in 2011 and the minimum equilibrium in 2009.

These simulations illustrate the rough magnitude of the pricing distortion using reasonable demand and cost parameters estimated from the Massachusetts setting. However, the actual distortion may have differed from our simulations for several reasons. The market may not have reached equilibrium by the period we consider. Limits on medical loss ratios (MLRs) could limit the price distortion under-price linked subsidy if setting its optimal price would give the cheapest plan too high markups. MLRs – calculated as the ratio of medical costs to revenue – for each simulation are reported in Appendix C. For most cases, MLRs are above the required floor of 80%, with the exception of BMC in the max-price equilibrium in 2009 (whose MLR is 77%). In addition to MLRs, the actual Massachusetts market differed from our simulations due to the presence of below-poverty (fully subsidized) consumers and because of other price regulations in place. We do not model these factors, which did not continue in the ACA exchanges. Despite these caveats, we think this analysis indicates the potential for substantial price distortions from price-linked subsidies.

A simple way to interpret the difference between the equilibria – without the assumptions needed for welfare analysis – is to ask how much money the government could save by switching to fixed subsidies, after adjusting the subsidy amount to *hold insurance coverage fixed*. By this test, we find that net expenditures would be 6.1% lower in 2011 and 0.7-7.8% lower in 2009 (for the range of equilibria). These would translate to substantial savings in programs the size of CommCare (about \$800 million in 2011) or the ACA exchanges (about \$40 billion in 2016). Alternatively, we can ask how much higher insurance coverage rates would be under fixed subsidies, *holding government spending fixed*. We find that coverage

(among the CommCare eligible population) would be 3.1 percentage points higher in 2011 and 0.2-2.6 percentage points higher in 2009.³⁷

Fewer Insurers

Many of the ACA exchanges have only two insurers. To see how this affects the distortion, we again simulate the market for price-linked and fixed subsidies using our model for 2011, but with the market limited to two available plans. Table 6 shows the results for each pair of the four main insurers (excluding Fallon, which is a smaller regional plan). The table shows the distortion – the increase in the cheapest price under price-linked relative to fixed subsidies – both as a dollar amount and as a percent of the price under fixed subsidies.

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 distortion amounts are larger than what we estimate for the full market with five insurers (\$23.8). More notable is how much larger the distortion is when there is a large cost difference between the two plans. For instance, the distortion is only slightly larger (\$28.0) when CeltiCare and Network Health, two low-cost plans, are competing. But it is more than twice as big (\$50.1) when CeltiCare competes against high-cost NHP. Thus, not only the number but the *type* of insurers matters critically for the distortionary effect of price-linking: it is less bad when the competing plans have more similar cost structures.

With only two insurers, medical loss ratio rules become more binding. CeltiCare’s cost-to-revenue ratio is below 80% under price-linked subsidies against any competitor and below 70% against NHP. (CeltiCare also has an MLR below 80% for fixed subsidies when competing with NHP). Network Health and BMC also have MLRs below 80% when competing against NHP. For these uncompetitive markets the distortion from price-linked subsidies may be smaller because of minimum MLRs.

5 Welfare and Cost Uncertainty

5.1 Welfare Framework

To assess the tradeoffs involved with price-linked subsidies, we need to specify a regulatory objective function. We assume the regulator values consumer surplus and dislikes government costs. Adverse selection implicitly gives the regulator a reason to subsidize health insurance as opposed to merely redistributing cash³⁸ (increasing subsidies can lower insurance prices,

³⁷These estimates are changes in the take-up rate for subsidized insurance among eligible individuals, not changes in the overall uninsured rate. The CommCare eligible population of about 300,000 was about 5% of the total state population.

³⁸There is a large literature about rationales for in-kind benefits; see Currie and Gahvari (2008) for a review.

Table 6: Price Distortion with Two Insurers in 2011

		(Lower cost)	→	(Highest cost)
		Network Health	BMC	NHP
(Lowest cost)	CeltiCare	\$28.00 (7.3%)	\$37.43 (9.4%)	\$50.13 (11.6%)
↓	Network Health		\$30.37 (6.7%)	\$36.78 (7.8%)
(Higher cost)	BMC			\$31.00-38.79 (6.0-7.5%)

(All 5 plans in the market: \$23.83(6.3%))

Note: This table shows the difference in the cheapest price under price-linked and fixed subsidies, both as a dollar amount and as a percent of the cheapest price under fixed subsidies, for markets with two insurers. The plans are listed in order of increasing costs: CeltiCare, Network Health, BMC, NHP. When BMC and NHP compete there are multiple equilibria, the range of distortions is given.

generating an overall social benefit), but we add the explicit rationale that the regulator cares about a social externality of people being uninsured. Together these form a consumer surplus standard (common in antitrust analysis) adjusted for government costs and social externalities; we call this “purchaser surplus” (PS). We think this standard is consistent with the public health care reform debate, which put little weight on insurer profits,³⁹ but we also allow the regulator value insurer profits. Profits get a weight of $0 \leq \chi \leq 1$, perhaps reflecting the excess burden of raising government funds.⁴⁰

These components of welfare depend on the subsidy and mandate levels (S and M) and market costs (C), which determine equilibrium prices ($P^{cons} = P - S$), which drive consumer demand. Formally, public surplus is

$$PS = \underbrace{CS(P^{cons}, M)}_{\text{Cons. Surplus}} + \underbrace{(1 - D_0(P^{cons}, M)) \bar{E}(P^{cons}, M, C)}_{\text{Avoided Externality of Uninsurance}} - \underbrace{(S - (S - M)D_0(P^{cons}, M))}_{\text{Net Govt. Cost}},$$

where $D_0(P^{cons}, M)$ is the share uninsured and $\bar{E}(P^{cons}, M, C)$ is the average externality avoided per insured consumer (which depends on the level of costs, C in the market). If we subtracted the externality of the uninsured instead of adding the “saved externality” of the insured, the large direct surplus loss from higher costs would make it hard to see the differences between subsidy regimes. The two are conceptually equivalent. Government costs equal subsidies paid for the insured minus mandate penalty revenue collected from the uninsured. We also look at public surplus plus firm profits, $PS + \chi \cdot \pi(P, S, M, C)$.

³⁹Indeed, policymakers seem eager to constrain insurer profits with policies like medical loss ratio limits.

⁴⁰For instance, if each \$1 of public funds involved \$0.30 of excess burden, the associated χ would be $1/1.30 = 0.77$. We do not downweight consumer surplus in the same way because the government can adjust the level of subsidies and the mandate penalty to transfer money to/from consumers until it is indifferent between \$1 lower consumer prices and \$1 of government funds. We do not model this optimal redistribution problem directly, though we consider the role of consumer risk aversion in our discussion below.

Externality and Cost shocks

The externality of uninsurance has two components. The first is the cost of the health care the uninsured receive, but do not pay for. Recent work has shown that the uninsured use substantial charity care (or “uncompensated care”), whose cost is borne by hospitals and public clinics who cannot or will not deny needed care (Mahoney, 2015; Garthwaite et al., 2015; Finkelstein et al., 2015). The second component is a pure (paternalistic) social disutility of people lacking insurance; this is consistent with much of the political language motivating the subsidies and rationalizes observed level of subsidies (which are higher than charity care costs).

Charity care costs are, unfortunately, quite difficult to measure. Instead of attempting to estimate them in our Massachusetts setting, we assume that charity care costs $C_i^{Charity}$ are proportional to an individual’s expected costs in the average exchange plan. We then treat the paternalistic part of the externality, E_0 , as a residual to rationalize a given baseline subsidy level as optimal. For each simulation specification we calibrate E_0 so that at baseline costs, 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 would be partly driven by the deviation of subsidies from their optimal level, absent a cost shock. Empirically, we find that a positive E_0 is needed to rationalize the ACA’s affordable amounts.⁴¹

The welfare effects of different subsidy policies under cost uncertainty depend on how the externality is affected by cost shocks. We assume that the paternalistic component is unaffected, but that charity care costs move with market level costs, according to a scaling parameter, $\lambda \leq 1$.⁴² We consider several values for λ . Finkelstein et al. (2015) find that when individuals are uninsured, third parties cover about 60% of what their costs would have been if insured by Medicaid (with individuals paying 20% and 20% being moral hazard), suggesting a value of $\lambda = 0.6$. Since price-linking is more attractive when the externality moves more with prices, we consider $\lambda = 1$ as a best-case scenario for price-linked subsidies. We also consider the intermediate case of $\lambda = .8$. Putting all this together, for a given cost shock Δ , the externality for consumer i is

$$E_i = C_i^{Charity} + E_0 = \lambda \cdot \underbrace{\exp(\hat{\mu}X_i + \psi_0)}_{\text{Baseline cost in avg. plan}} (1 + \Delta) + E_0, \quad \lambda = \{.6, .8, 1\}.$$

⁴¹With $E_0 = 0$ the implied over-subsidizing of insurance would mean that the additional people insured under fixed subsidies would be a loss to the government, where really that loss is due to an affordable amount that is too high given the externality.

⁴²More generally, we can think of E_0 as including any component of the externality that does not scale with costs and include any components that scale with cost in the term for charity care.

Price-linked vs Fixed Subsidies

Absent uncertainty – i.e., with full information about costs, demand, and pricing behavior – the regulator can predict prices and therefore the subsidy amount, S^E , that will emerge under a price-linked subsidy policy. The regulator could then replicate the same subsidy amount as a fixed subsidy, leading insurers to lower prices. These lower prices (with no change in the subsidy amount) lead to gains for consumers and more people getting insurance. If the level of M had been set optimally, the regulator would be indifferent between the costs and benefits of additional insurance purchases. But the gains to consumers (at insurers' expense) would be desirable to the regulator as long as $\chi < 1$. Appendix A formalizes this argument under special conditions that make the problem analytically tractable;⁴³ it holds empirically in our simulations.

Under uncertainty, this argument breaks down, since the regulator cannot predict the equilibrium prices and subsidy. Price-linked subsidies may be better, despite higher prices, if (1) regulators are uncertain about market costs (or prices), *and* (2) the *optimal* subsidies are higher in states of the world where prices are unexpectedly high. To see the importance of the second condition, consider the analogy of setting a tax on a pollution-producing good. If each unit of the good generates a negative externality of $\$X$ – regardless of the good's price or production cost – then the optimal tax is $\$X$. Linking the tax to prices makes little sense even if costs or prices are uncertain.

The externality of charity care creates this link between prices and optimal subsidies in the case of health insurance. If there are market-wide health care cost shocks – e.g., an expensive new treatment or an increase in nurses' wages – these may increase both insurers' costs (and prices) *and* the externality of charity care. Higher prices would therefore signal a larger externality and a larger optimal subsidy, creating a rationale for price-linked subsidies. Importantly, these shocks must be observed by insurers but not by the regulator; otherwise, prices would contain no additional information for regulators.

If consumers are risk averse,⁴⁴ price-linked subsidies also have the advantage of stabilizing post-subsidy consumer prices and transferring the cost risk to the government. Let consumers have concave utility over non-health insurance consumption, $u(\cdot)$, with a (local) coefficient of relative risk aversion γ . If Y is income and P is the premium for insurance, then the

⁴³The proof in Appendix A uses logit demand and assumes no adverse selection. Intuitively, the logic requires that prices be strategic complements (as in logit) or not too strongly substitutes, so that a higher price of the pivotal plan does not *lower* average prices.

⁴⁴This is perhaps reflected in the rhetoric around “affordability,” since forcing poor people to buy expensive insurance is a lump-sum tax.

change in utility from a premium increase of ΔP is

$$\begin{aligned} u(Y - P - \Delta P) - u(Y - P) &\approx -u'(Y - P)\Delta P + u''(Y - P)\frac{(\Delta P)^2}{2} \\ &= -\Delta P - \frac{\gamma}{Y - P}\frac{(\Delta P)^2}{2}. \end{aligned} \quad (7)$$

The second line follows from the fact that in order to combine consumer surplus and government spending, we implicitly assume that prior to any cost shocks, the transfer system was set “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 is normalized to one). When weighted by demand shares, the effect on risk neutral consumers included in our baseline calculation. The second term is the additional cost from risk aversion, which we include as a robustness check.⁴⁵

5.2 Baseline Welfare

For the simulated equilibria described in Section 4.4, we can calculate the different components of welfare under each subsidy policy. In each year, for each subsidy policy, 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 $CS_i = \frac{1}{-\alpha(Z_i)} \cdot \log\left(\sum_{j=0}^J \exp(\hat{u}_{ij})\right)$ and choice probabilities ($\widehat{Pr}_{ij} = \exp(\hat{u}_{ij}) / \sum_{j=0}^J \exp(\hat{u}_{ij})$). The government’s expected net expenditures on each individual are $S - (S + M) \cdot \widehat{Pr}_{i0}$. We calculate the externality of charity care and calibrate the paternalistic externality by assuming the government set the affordable amount optimally given the true externality.

Table 7 shows the impacts of price-linked and fixed subsidies on different components of welfare. The numbers reported are per *eligible* consumer (including the uninsured) so tend to be smaller in magnitude than the price changes discussed above. For instance, consumer surplus is about \$5 per person-month higher under fixed subsidies in 2011 – an effect driven by the \$24 fall in CultiCare’s price applied to about 20% of the eligible population who purchase CultiCare. We note that the overall level of consumer surplus is relatively small (and sometimes negative) because we compute it relative to a world without the program – i.e., no subsidized insurance *and* no mandate penalty. Because a sizable share of eligible people are uninsured (and therefore hurt by the mandate penalty), it is not surprising that

⁴⁵To calculate a consumer’s non-health insurance consumption ($Y - P$) 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 $(\Delta P)^2$, we calculate the premium change under fixed subsidies relative to baseline for each plan, square it, and weight by the individual’s probability of purchasing each plan. 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.

the program has little (or negative) effect on consumer surplus.

Table 7: Costs and Surplus under Price-Linked and Fixed Subsidies

		\$ Per Eligible Consumer per Month					
Subsidy Type		Cons Surplus ⁺	Saved Externality ⁺	Gov Costs ⁼	Public Surplus	Insurer Profits	PS+ Profits
2009	Min Price-linked	-0.3	246.2	-177.2	68.6	33.9	102.6
	Max Price-linked	0.9	249.7	-193.2	57.3	45.3	102.6
	Fixed	0.7	249.0	-179.9	69.8	33.4	103.3
	Min P-link - Fixed	-1.0	-2.9	2.7	-1.2	0.5	-0.7
	Max P-link - Fixed	0.1	0.7	-13.3	-12.5	11.9	-0.7
2011	Price-Linked	-14.4	165.5	-103.2	47.9	22.3	70.2
	Fixed	-9.3	188.2	-125.3	53.6	22.7	76.3
	Price-link - Fixed	-5.1	-22.7	22.1	-5.7	-0.4	-6.1

Note: This is a comparison of the market equilibria under price-linked and fixed subsidies. In 2009, there are a range of equilibria under price-linked subsidies; this reports statistics for the equilibrium with the minimum and maximum lowest price. ‘Per Eligible Consumer’ includes both insured and 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 is mandate revenue minus subsidy expenditures. Saved externality is the sum across consumers of the probability that they buy insurance times their externality, as defined in Section ???. Public surplus adds consumer surplus, the saved externality and government costs. The last column adds firm profits to public surplus.

Public surplus is the sum of consumer surplus, the saved externality, and government costs. We focus first on the welfare comparisons between price-linked and fixed subsidies where the subsidy amounts are held fixed – i.e., in 2011 and the minimum price-linked equilibrium in 2009. In these cases, the changes in government costs and the saved externality are both due to changes in the insurance coverage rate. These changes move in opposite directions by about the same amount by an envelope theorem argument – because fixed subsidies are set optimally, the government’s value of covering an additional person equals the additional public costs. As a result, the change in public surplus is approximately equal to the change in consumer surplus, which increases because insurer prices fall while subsidies do not change. In 2011, public surplus is \$6 lower per month per potential customer under price-linked subsidies. In the minimum price-linked equilibrium for 2009, public surplus is just \$1 lower, reflecting the fact that the minimum price distortion consistent with equilibrium is relatively small (see Table 5).

If we instead examine 2009’s maximum price equilibrium, public surplus is \$13 per month lower under price-linked than under fixed subsidies. We note that this larger decrease is driven both by the larger price distortion (as plans coordinate on their preferred equilibrium) and by the fact that government subsidies are higher (by about \$22). As a result, the fall in

public surplus is driven by an increase in government costs, with little change in consumer surplus or the saved externality.

The last two columns of Table 7 show insurer profits and the sum of public surplus and profits. Profits are higher under price-linked subsidies in 2009 (for both equilibria) but actually a bit lower in 2011. These different results illustrate the ambiguous effect of shifting from fixed to price-linked subsidies on total profits. While prices increase, the higher prices also decrease quantity insured, particularly among healthier, low-cost consumers.⁴⁶ As a result, we find that the sum of public surplus and profits also falls under price-linked subsidies in all of our simulations.

5.3 Cost Uncertainty

In allowing for cost uncertainty, it is key that the fixed-subsidy is set prior to the cost shock realization. We set it equal to the price-linked subsidy that emerges in equilibrium with baseline costs. We then simulate equilibrium under each subsidy policy when actual costs are scaled upward/downward by a shock Δ that ranges from -15% and +15%. This proportional adjustment applies to all plans' costs and to an individual's cost of charity care. We assume that firms observe this cost shock and price based on it but that it is unobserved by the regulator. We also assume that the shock is unobserved to consumers so their utility of insurance (as reflected in ξ_j and β) is unaffected.⁴⁷

To conserve space, we focus on simulations for 2011; results for 2009 are qualitatively similar and shown in Appendix C. Figure 2a shows the price of the cheapest plan and the average subsidy for 2011. Unsurprisingly, prices (shown in solid lines) increase with costs under both policies. They increase slightly faster under price-linked subsidies.⁴⁸ The figure illustrates how subsidies (shown in dashed lines) move quite differently under the two policies. Subsidies increase in tandem with the cost shock under price-linked subsidies, but are flat under the fixed policy (aside from small changes in the average, driven by changes in the income composition of insured consumers). We note that for the largest negative cost shocks we simulate, prices under fixed subsidies can fall slightly below the subsidy amount for some income groups. In these cases, we assume that insurers can charge “negative premiums” – i.e., rebate the difference to consumers.⁴⁹ If instead the government did not allow rebates, then

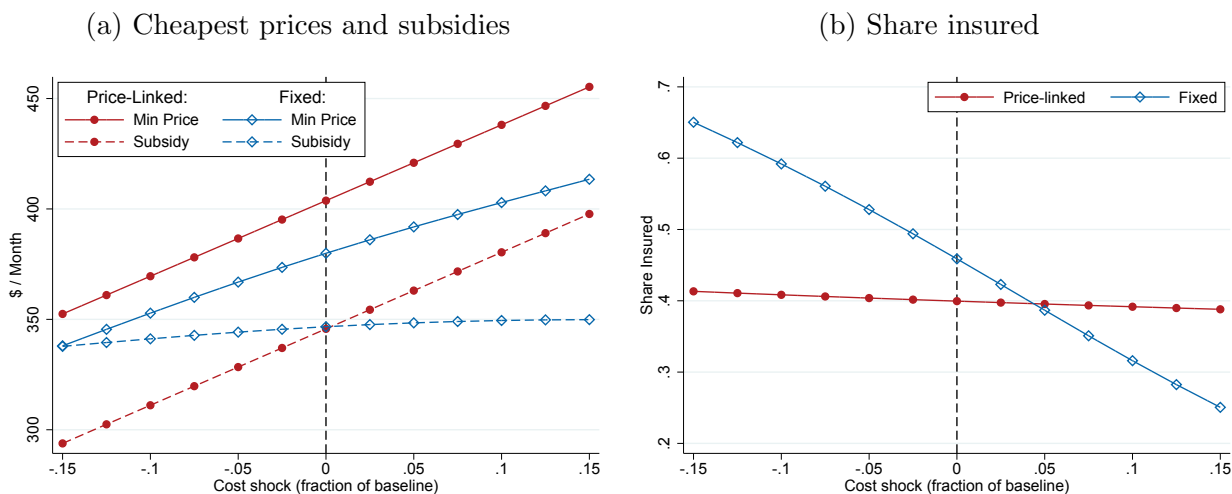
⁴⁶This does not imply that insurers would have been better off if they had colluded to lower prices because under the price-linked subsidies, that would have also lowered the subsidy, so they would not have gotten the resulting increase in demand and healthier consumers.

⁴⁷Even if consumers observed the cost shock, it is not obvious how their utility of insurance would be affected. Insurance against a larger risk is generally more valuable, but higher costs also mean the uninsured receive more charity care. Conceptually, if we allowed demand to shift with the cost shock, our simulated effects on coverage would change. But the basic welfare/policy implications – which are driven by the size of the subsidy and externality – would not clearly be different.

⁴⁸On average a 15% cost shock is about \$56, so prices increase a little less than one-for-one with costs.

⁴⁹The ACA does not allow negative premiums, but implementing these does not seem infeasible. Rebates

Figure 2: Equilibrium under cost shocks



Note: These figures show the equilibrium under price-linked and fixed subsidies for cost shocks of -15% to 15% of baseline. The left figure shows the price of the cheapest plan and the subsidy level. The right figure shows the share of consumers who purchase insurance.

fixed subsidies could also distort pricing incentives, since insurers would have no incentive to lower prices below the subsidy amount. Although not a major issue for the range of shocks we consider (where average subsidies are always less than the minimum price), this could become important for even larger negative shocks.

Figure 2b shows the impacts on the share of eligible people who purchase insurance. Price-linked subsidies stabilize the insurance take-up rate (at about 40%), regardless of the cost shock. This outcome is natural because price-linked subsidies intentionally hold fixed consumer premiums for the cheapest plan, regardless of whether insurer prices rise or fall. However, under fixed subsidies coverage varies substantially: it is much higher (up to 65%) with a negative cost shock and lower (down to 25%) with a positive shock. The ability of price-linked subsidies to stabilize coverage in the face of cost shocks is an important property. However, whether this property is desirable is less clear: it may be optimal for fewer people to buy insurance when it is more expensive.

Figure 3 shows public surplus under the different assumptions on how the externality scales with costs, again for cost shocks from -15% to +15%. Consistent with our earlier discussion, fixed subsidies (in blue) are preferred at zero cost shock in all cases.⁵⁰ But as costs diverge from expectations, the gap between fixed and price-linked subsidies narrows,

to consumers could be accomplished either via direct payments or via tax rebates (similar to the way the ACA's mandate penalty works).

⁵⁰The differing values of λ and the calibrated fixed externality E_0 under these three scenarios mean that the levels of public surplus do not line up across graphs. However, in each case, the units of public surplus are dollars per eligible person per month.

particularly for larger values of λ . This occurs because the optimal subsidy gets farther from the fixed subsidy level as costs – and therefore the externality of uninsurance – diverge from expectations. If the externality increases one-for-one with costs ($\lambda = 1$), price-linked subsidies do better for cost shocks $\geq 15\%$ or less than -12.5% . However, as the externality increases less with costs (lower λ), larger cost shocks are necessary to make price-linking better. Indeed, for $\lambda = 0.6$, we find fixed subsidies do better even for shocks of over 20% .

If we include profits in the welfare metric, and use the corresponding calibrated fixed externality,⁵¹ the results are conceptually similar and are shown in Figure 5 in Appendix C. For $\lambda = 0.8$ or 1.0 , the range of shocks under which fixed subsidies do better is somewhat smaller (e.g., Δ between -7.5% and $+10\%$ for $\lambda = 1.0$). But for $\lambda = .6$, fixed subsidies are still better for the full range of shocks we consider.

Appendix C also shows results for 2009, using public surplus (without profits) as the welfare metric. We show the minimum price-linked equilibrium, in which the pricing distortion is quite small. Again, the results are qualitatively similar: fixed subsidies do better at zero cost shock but worse as the cost shock grows. Not surprisingly, given the baseline results, the cost shocks necessary for price-linked subsidies to be better are smaller than in 2011. For $\lambda = 1$, price-linked subsidies are better for $|\Delta| \geq 7.5\%$. However, the maximum priced equilibrium (not shown) makes fixed subsidies look even better than they do in 2011.

Adding Risk Aversion

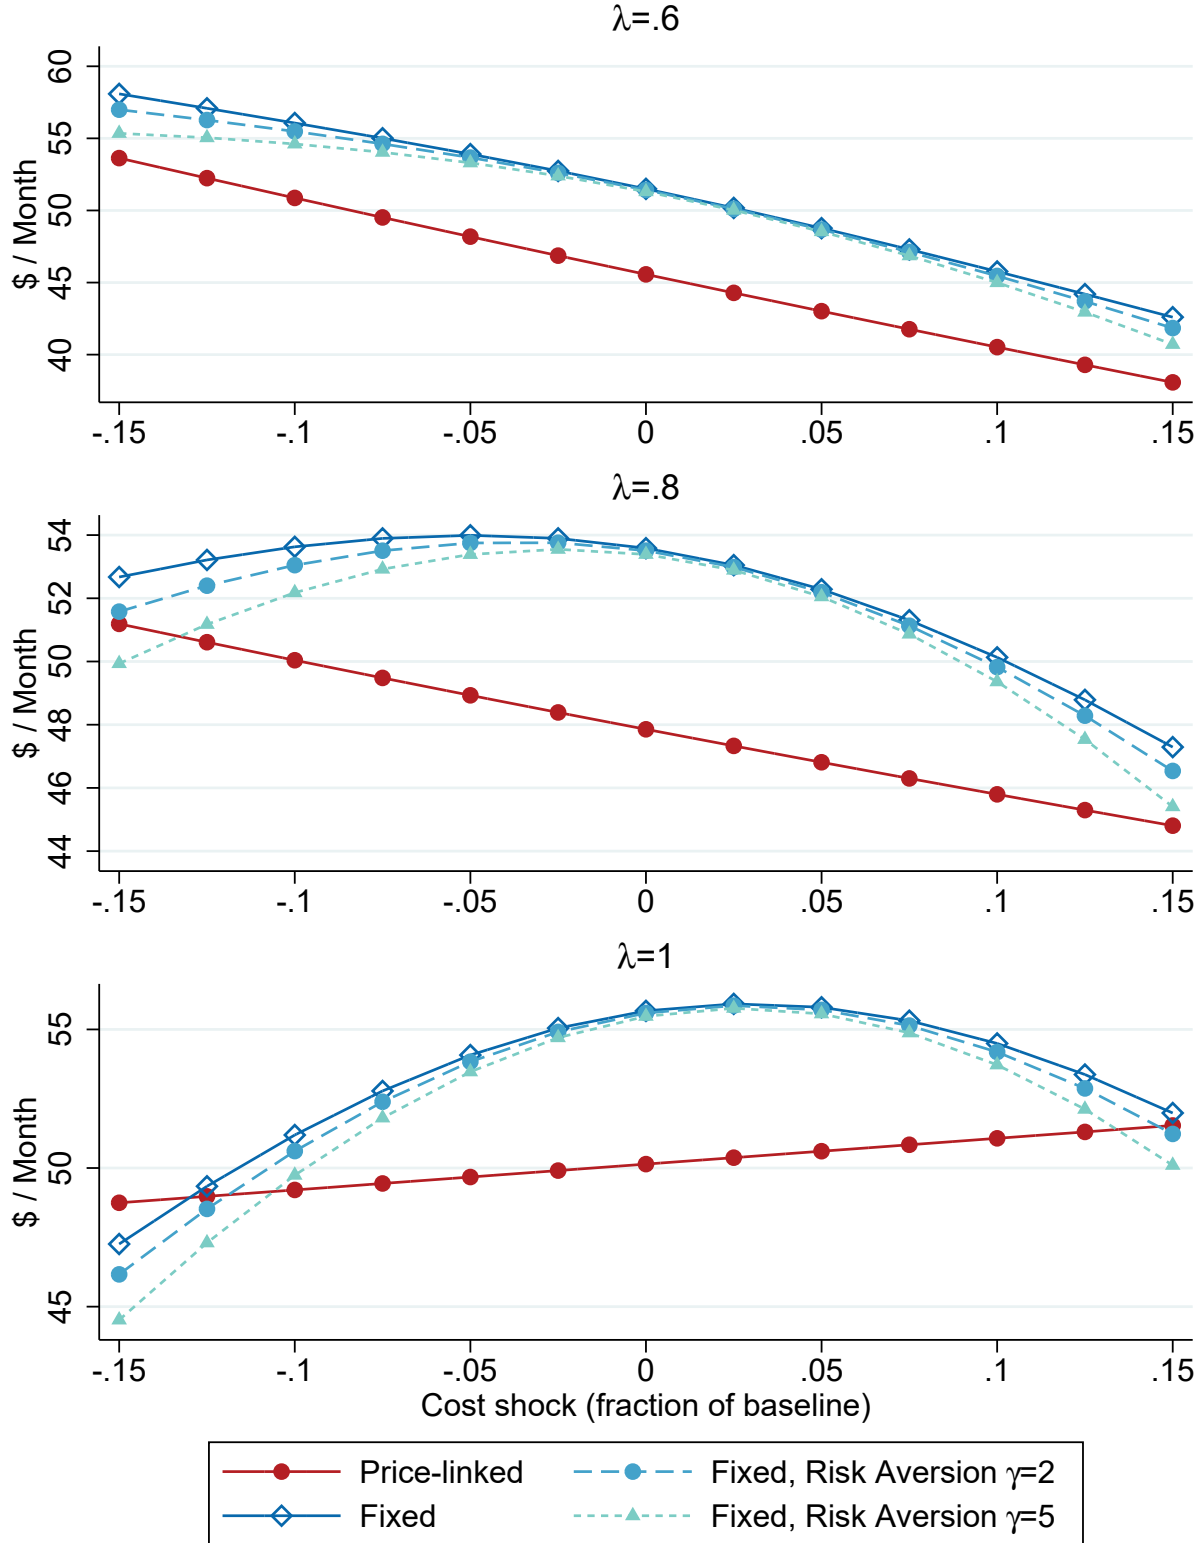
Fixed subsidies lead to more variation in consumer premiums from cost shocks than price-linked subsidies; this is costly when consumers are risk-averse – utility is concave. The dashed and dotted lines in figure 3 adjust the welfare under fixed subsidies to include the risk cost to consumers.⁵² Chetty (2006) argues that $\gamma \leq 2$; to consider the case most favorable to price-linked subsidies, we use $\gamma = 2$, but the effect is still small. Under $\lambda = 1$, fixed subsidies go from marginally better than price-linked to marginally worse for $\Delta = -12.5\%$ and $\Delta = 15\%$; the results are otherwise qualitatively unchanged. This is explained by the well-known fact that for reasonable risk aversion, utility is locally quite close to linear (Rabin, 2000). The premium variation we consider (about $\pm \$50$ per month) represents only about $\pm 3\%$ of average income, even for this low-income population.

Figure 3 also shows results for a higher $\gamma = 5$, to capture the idea that society might have a strong concern about “affordability” or to proxy for factors like consumption commitments that increase local risk aversion (Chetty and Szeidl, 2007). The effect is obviously larger, but the comparison to price-linked subsidies is still comparable to $\gamma = 2$. We thus conclude

⁵¹As we discuss in the Appendix, if the regulator is maximizing social surplus, then a different level of E_0 , the fixed externality, is needed to rationalize the baseline subsidy level as optimal.

⁵²See Equation (7).

Figure 3: Public surplus under cost shocks in 2011



Note: This figure shows public surplus (in dollars per month per eligible member) 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. The dashed and dotted lines are risk adjusted with factors of relative risk aversion of 2 and 5, respectively.

that in our framework, insuring consumers against price risk is a less important rationale for price-linked subsidies than the correlation between prices and the externality of charity care.

How Much Cost Uncertainty is Reasonable?

To get a sense of what size cost shocks are most relevant, we need to think about how much uncertainty a policymaker faces when setting a fixed subsidy. We use state-level average costs in the U.S. National Health Expenditures (NHE) data for 1991-2009 (the period over which state-level data are available) to get ballpark magnitudes for different scenarios.

First, consider the case of a mature market and a policymaker attempting to set fixed subsidies in a ‘smart’ way based on all available information. We assume that the policymaker can observe lagged cost data from insurers (e.g., in state insurance department rate filings) and other data sources (e.g., hospital cost reports, Medicare data). Thus, the relevant uncertainty is about how much costs will *grow* between the lagged data and the year for which subsidies are being set. In the NHE data, the standard deviation of state-level annual cost growth is 1.9%, and for three-year growth it is 4.8%. Thus, if a regulator can observe costs for the current year when setting subsidies for next year, cost shocks (i.e. deviations from an expected change) of $|\Delta| > 5\%$ occur less than 1% of the time. Cost shocks of $|\Delta| > 12.5\%$ – the minimum shock for which price-linked subsidies do better in our $\lambda = 1$ case – are an extreme tail outcome. If (more conservatively) the regulator can observe costs from two years ago when setting next year’s subsidies (a three-year lag), cost shocks of $|\Delta| > 12.5\%$ still occur less than 1% of the time. We conclude that, in this case, fixed subsidies do better across a broad range of “reasonable” cost uncertainty.

However, cost uncertainty may be larger than in this case for several reasons. Uncertainty is likely to be much larger in a new market, like the ACA in 2014, where past data is not available. Furthermore, there may be political economy constraints that prevent adjusting fixed subsidies in the way described above. For instance, Congress might set an initial level of fixed subsidies and index them based on an assumed rate of cost growth. Actual costs could diverge substantially from this assumed trend over an extended period. For instance, using the medical CPI + 1% as an index, after 10 years, costs in the NHE data would on *average* have been 9% higher than expected and would diverge from expectations by more than 15 percentage points about one-fourth of the time.

There are many other combinations of updating and indexing fixed subsidies that one could consider, but these examples show a general pattern. In the medium-to-long term, just indexing fixed subsidies will frequently lead to worse outcomes than price-linked subsidies. But ‘smart’ fixed subsidies that try to adjust for actual cost trends in an area will generally do better than price-linked subsidies, even if they are updated only every few years.

6 Discussion and Implications for Other Markets

We have presented the tradeoffs involved with price-linked subsidies and analyzed these empirically in a realistic model of a health insurance exchange. Of course, the relative costs and benefits of fixed and price-linked subsidies will vary from market to market. In general, the more uncertainty the government faces or the more competitive the market, the more likely that the benefits of price-linked subsidies will outweigh the costs. 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 that does not distort pricing. If varying the mandate penalty is less politically concerning than allowing prices to vary, then our alternate policy idea of avoiding the distortion by applying the subsidy to a higher baseline mandate penalty may be desirable. Alternatively, if letting prices vary is less politically concerning, fixed subsidies and mandate penalties may be preferable.

In more competitive markets, the government will want to use the information from market prices because doing so does not distort prices as much.⁵³ We have focused on the competitiveness of the insurer market, but monopoly power in the provider market is also problematic for price-linked subsidies. If insurers are competitive and price at marginal cost, a monopoly provider will take into account how much increasing its price (which is passed through to consumer prices) affects consumers' demand for insurance. With price-linked subsidies, demand is much less responsive to price, so providers will charge higher prices to insurers. Note that MLRs would not help in this scenario because the insurers are pricing at cost; it is the provider whose price is distorted and who gets a profit windfall.

6.1 Other Markets

ACA, Medicare, and Employer Insurance

The Massachusetts market has now been converted into an ACA exchange. The ACA exchanges differ in several respects: plans in multiple generosity tiers, multi-plan insurers, a slightly different risk adjustment formula, subsidies that depend on the second cheapest silver plan, and the inclusion of higher income groups (including unsubsidized consumers). The 10-20% of consumers who are unsubsidized could mitigate the distortionary effect of price-linked subsidies (which do not apply to their demand) if they do not all buy bronze-

⁵³In 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.

level plans. As we have discussed, medical loss ratio limits may have bite in limiting firms’ profit margins, but the flexibility of these rules – insurers can count vaguely defined “quality improvement” activities as medical costs – may limit their effectiveness.

Other factors may exacerbate the distortion in the ACA relative to Massachusetts. A *New York Times* analysis found that 58% of counties served by the federal exchange, had two or fewer insurers.⁵⁴ 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. If insurers weren’t required to offer a bronze plan and most consumers on the margin of uninsurance chose bronzes plans, the price distortion could be mitigated. Our uncertainty calibrations suggest that a fixed subsidy could be better *if* the level of the subsidy were not just indexed to an assumed growth rate (e.g., the medical CPI) but also updated regularly to reflect changes in area-level costs of health care.

The tradeoffs with price-linked subsidies that we have highlighted apply more broadly than the ACA. They apply in any market where (1) firms have market power⁵⁵ and (2) there is the possibility of substitution to an unsubsidized outside option. These conditions apply in Medicare Advantage, Medicare Part D, and employer-sponsored insurance programs.

In Medicare Advantage (MA), the distinction between price-linked and fixed subsidies is relevant for comparing “competitive bidding” and “premium support” reform proposals.⁵⁶ Both reforms propose explicitly linking subsidies to insurer prices. Under competitive bidding proposals, the price-linked subsidy applies only to MA plans, while the enrollee premium for traditional Medicare (the outside option) is held fixed. As we have shown, this distorts pricing incentives. Premium support applies the (price-linked) subsidy to all options, including traditional Medicare; this works like our alternate policy idea and avoids the pricing distortion.

Medicare Part D (the prescription drug program for the elderly) uses price-linked subsidies based on a national enrollment-weighted average of plan price bids. Because all plans’

⁵⁴Moreover, 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.

⁵⁵The 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).

⁵⁶MA’s current design has a combination of fixed and price-linked subsidies: benchmarks are set based on local traditional Medicare costs, but Medicare reduces the subsidy when a plan reduces its price below the benchmark. The distortion from these plan-specific price-linked subsidies may be significant.

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).⁵⁷

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.

Non-health Markets

The markets tend to be less centralized, but housing subsidies, school vouchers, and Pell Grants have many of the same properties of health insurance subsidies. If a city sets the value of a school voucher to ensure the affordability of at least one private school, rather than basing it on the cost of public education, it risks distorting upwards the private school prices. Most housing markets have more suppliers, so setting housing subsidies based on market level housing prices is likely to be less distortionary.⁵⁸ Again, price-linked subsidies would only be beneficial if the optimal subsidy increased with prices.

7 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 use 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 \$24 or 6% 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 \$24 increase in monthly subsidies would translate to \$46 million per year in government costs. For the ACA with 10.5 million subsidized enrollees for 2016 (ASPE, 2016), a \$24 per month subsidy increase would translate to over \$3 billion per year in federal spending.

⁵⁷Decarolis (2015) discusses why the distortion from the low-income subsidy is likely larger.

⁵⁸Linking an individual’s subsidy to the price of the specific apartment, as is implicitly done if the individual’s rent contribution is a fixed fraction of their income, would still be very distortionary if there were no other price regulations.

We do not view these numbers as a precise estimate of either the historical distortion in Massachusetts or the 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 become 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 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.
- Office of the Assistant Secretary for Planning and Evaluation, Department of Health and Human Services (2016). Health insurance marketplaces 2016 open enrollment period: Final enrollment report.
- Aizawa, N. (2016). Health insurance exchange design in an empirical equilibrium labor market model.
- Avalere Health (2015). Individual mandate penalty may be too low to attract middle-income individuals to enroll in exchanges.
- 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.
- Chetty, R. and Szeidl, A. (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics*, 122(2).
- 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.
- 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.
- 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. (2014). Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64.
- Ericson, 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.
- Ericson, K. M. M. and Starc, A. (2016). How product standardization affects choice: Evidence from the Massachusetts health insurance exchange. *Journal of Health Economics*, pages 71–85.
- Finkelstein, A., Hendren, N., and Luttmer, E. F. (2015). The value of Medicaid: Interpreting results from the Oregon health insurance experiment.
- Garthwaite, C., Gross, T., and Notowidigdo, M. J. (2015). Hospitals as insurers of last resort.
- 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.
- Ho, K. (2006). The welfare effects of restricted hospital choice in the US medical care market. *Journal of Applied Econometrics*, 21(7):1039–1079.
- Ho, K., Hogan, J., and Morton, F. S. (2015). The impact of consumer inattention on insurer pricing in the Medicare Part D program.
- Ho, K. and Lee, R. S. (2013). Insurer competition in health care markets.
- Liu, Y. and Jin, G. Z. (2015). Employer contribution and premium growth in health insurance. *Journal of Health Economics*, 39:228–247.
- Mahoney, N. (2015). Bankruptcy as implicit health insurance. *The American Economic Review*, 105(2):710–746.
- Polyakova, M. (2016). Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D. *American Economic Journal: Applied Economics*, 8(3).

- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292.
- Ruggles, S., Genadek, K., Goeken, R., Grover, J., and Sobek, M. (2015). Integrated public use microdata series: Version 6.0 machine-readable database. Minneapolis: University of Minnesota.
- Shepard, M. (2016). Hospital network competition and adverse selection: Evidence from the Massachusetts health insurance exchange.
- Shleifer, A. (1985). A theory of yardstick competition. *The RAND Journal of Economics*, 16(3):319–327.
- Starc, A. (2014). Insurer pricing and consumer welfare: Evidence from Medigap. *The RAND Journal of Economics*, 45(1):198–220.
- Tebaldi, P. (2016). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA.
- Town, R. and Liu, S. (2003). The welfare impact of Medicare HMOs. *RAND Journal of Economics*, pages 719–736.
- 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

Multi-Plan Insurers (ACA Case)

In the ACA insurers must offer a plan in each of multiple tiers – bronze, silver, gold, and platinum.⁵⁹ Subsidies are set equal to the price of the second-cheapest silver plan minus a pre-specified “affordable amount.”⁶⁰ 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 a 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 $l = \{(B)ronze, (S)ilver, (G)old, (P)latinum\}$. For notational simplicity, 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:

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

The second term reflects the standard effect for multi-product firms that when the insurer raises the price of one plan, it captures revenue from consumers who switch to its other plans. However, with price-linked subsidies, the markup for the subsidy-pivotal plan is:

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

⁵⁹Platinum 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.

⁶⁰The subsidy applies equally to all plans (though no premium can go below \$0), ensuring that at least two silver plans (and likely some bronze plans) cost low-income consumers less than the affordable amount.

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 6.

Optimality of Fixed Subsidies

For logit demand, we can show that, absent uncertainty, fixed subsidies are better than price-linked. Let firm 1 be the pivotal, lowest priced plan ($\underline{j} = 1$). Suppose the government sets both a fixed component of the subsidy (which could be zero) and decides to what extent the subsidy will depend on the price of cheapest plan. So the total subsidy is $\tilde{S} = S + \alpha P_1$. We analyze how α affects welfare, assuming S is set optimally. The optimality of S implies

$$\frac{dW}{dS} = \frac{\partial W}{\partial \tilde{S}} + \sum_j \frac{\partial P_j}{\partial S} \frac{\partial W}{\partial P_j} = 0. \quad (9)$$

For the effect of α on welfare we have

$$\begin{aligned} \frac{dW}{d\alpha} &= \frac{\partial W}{\partial \tilde{S}} \cdot P_1 + \sum_j \frac{\partial P_j}{\partial \alpha} \frac{\partial W}{\partial P_j} \\ &= \sum_j \frac{\partial W}{\partial P_j} \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_1 \right) = - \sum_j D_j \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_1 \right) \end{aligned} \quad (10)$$

where the second line uses Equation (9) to substitute for $\frac{\partial W}{\partial \tilde{S}}$.

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

$$\begin{aligned} (P_1 - c_1)(1 - D_1 - \alpha D_0) &= 1, \\ (P_j - c_j)(1 - D_j) &= 1 \quad j \neq 1. \end{aligned}$$

For the lowest price plan, differentiating with respect to S gives

$$\begin{aligned} \frac{\partial P_1}{\partial S} (1 - D_1 - \alpha D_0) &= (P_1 - c_1) \cdot \\ &\left((D_1 + \alpha D_0) \sum_k \frac{\partial p_k}{\partial S} D_k + \left(\frac{\partial D_1}{\partial P_1} - D_1^2 \right) \frac{\partial P_1}{\partial S} + D_1 D_0 - \alpha D_0 (1 - D_0) \right) \\ \frac{\partial P_1}{\partial S} ((1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)) &= (D_1 + \alpha D_0) \sum_k \frac{\partial p_k}{\partial S} D_k - D_0 (\alpha - \alpha D_0 - D_1) \end{aligned}$$

where the second line uses the first-order condition to substitute for $P_1 - c_1$. For the other

plans we have

$$\frac{\partial P_j}{\partial S} ((1 - D_j)^2 + D_j) = D_j \sum_k \frac{\partial p_k}{\partial S} D_k + D_j D_0 \quad j \neq 1.$$

Taking the weighted sum gives

$$\begin{aligned} \sum_j \frac{\partial P_j}{\partial S} D_j &= \sum_{j \neq 1} \frac{D_j}{(1 - D_j)^2 + D_j} \left(D_j \sum_k \frac{\partial p_k}{\partial S} D_k + D_j D_0 \right) \\ &\quad + D_1 \frac{(D_1 + \alpha D_0) \sum_k \frac{\partial p_k}{\partial S} D_k + D_0 (D_1 - \alpha (1 - D_0))}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)}. \quad (11) \end{aligned}$$

Differentiating the first-order conditions with respect to α gives

$$\begin{aligned} \frac{\partial P_1}{\partial \alpha} (1 - D_1 - \alpha D_0)^2 &= \frac{\partial D_1}{\partial \alpha} + \alpha \frac{\partial D_0}{\partial \alpha} + D_0 + \sum_k \frac{\partial p_k}{\partial \alpha} \left(\frac{\partial D_1}{\partial p_k} + \alpha \frac{\partial D_0}{\partial p_k} \right), \\ \frac{\partial P_1}{\partial \alpha} ((1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)) &= P_1 D_0 (D_1 - \alpha (1 - D_0)) + D_0 + (D_1 + \alpha D_0) \sum_k \frac{\partial p_k}{\partial \alpha} D_k, \\ \frac{\partial P_j}{\partial \alpha} ((1 - D_j)^2 + D_j) &= D_j \sum_k \frac{\partial p_k}{\partial \alpha} D_k + P_1 D_j D_0. \end{aligned}$$

Taking the weighted sum gives

$$\begin{aligned} \sum_j \frac{\partial P_j}{\partial \alpha} D_j &= \sum_{j \neq 1} \frac{D_j}{(1 - D_j)^2 + D_j} \left(D_j \sum_k \frac{\partial p_k}{\partial \alpha} D_k + P_1 D_j D_0 \right) \\ &\quad + \frac{D_1 (D_1 + \alpha D_0) \sum_k \frac{\partial p_k}{\partial \alpha} D_k}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)} + \frac{D_1 D_0 (1 + P_1 (D_1 - \alpha (1 - D_0)))}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)}. \quad (12) \end{aligned}$$

Defining

$$X = \left(1 - \sum_{j \neq 1} \frac{D_j}{(1 - D_j)^2 + D_j} - \frac{D_1 (D_1 + \alpha D_0)}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)} \right),$$

and rearranging Equations (11) and (12), we get

$$\begin{aligned} X \sum_j \frac{\partial P_j}{\partial S} D_j &= \left(\sum_{j \neq 1} \frac{D_j^2 D_0}{(1 - D_j)^2 + D_j} + \frac{D_1 D_0 (D_1 - \alpha (1 - D_0))}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)} \right), \\ X \sum_j \frac{\partial P_j}{\partial \alpha} D_j &= \sum_{j \neq 1} \frac{D_j^2 D_0}{(1 - D_j)^2 + D_j} P_1 + \frac{D_1 D_0 (1 + P_1 (D_1 - \alpha (1 - D_0)))}{(1 - D_1 - \alpha D_0)^2 + D_1 (1 - \alpha D_0)}. \end{aligned}$$

Returning to (10), we have

$$\frac{\partial W}{\partial \alpha} = - \sum_j D_j \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial S} \cdot P_1 \right) = - \frac{1}{X} \frac{D_1 D_0}{(1 - D_1 - \alpha D_0)^2 + D_1(1 - \alpha D_0)} < 0, \quad (13)$$

where the inequality follows from

$$\begin{aligned} 1 - \sum_{j \neq 1} \frac{D_j^2}{(1 - D_j)^2 + D_j} - \frac{D_1(D_1 + \alpha D_0)}{(1 - D_1 - \alpha D_0)^2 + D_1(1 - \alpha D_0)} \\ > 1 - \sum_{j \neq 1} \frac{D_j^2}{D_j} - \frac{D_1(D_1 + \alpha D_0)}{(1 - D_1 - \alpha D_0)^2 + D_1(1 - \alpha D_0)} \\ &> 1 - \sum_{j \neq 1} D_j - (D_1 + D_0) = 0. \end{aligned}$$

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

B Data and Estimation Details

ACS Data Construction

We use the American Community Survey (ACS) to estimate the size of the CommCare-eligible population that chooses uninsurance and to generate a micro dataset of these individuals for demand estimation and simulations. We start from the ACS for Massachusetts in 2009 and 2011, the two relevant years for our analysis. We restrict to individuals age 19-64, since seniors are in Medicare and low-income children (up to 300% of poverty) are in Medicaid in Massachusetts. We restrict to people with household income $\leq 300\%$ of the federal poverty level. 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 also exclude the uninsured who are eligible for Medicaid (rather than CommCare) – parents up to 133% poverty and disabled individuals (proxied by receiving SSI income). 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 focusing on new enrollees simplifies 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 adjust the ACS data weights so that the uninsured share in our final demand sample matches the overall eligible population uninsurance rate. Specifically, we first estimate the uninsurance rate in each year using *all individuals* in both datasets – the population estimate of the number of CommCare-eligible uninsured from the ACS and the number of covered individuals (in member-years) in the CommCare data (including both new and current enrollees). We

then rescale the ACS weights in our demand estimation sample – which contains only new enrollees for the CommCare data – so that the uninsured rate calculated in this sample matches the population uninsured rate.

Natural Experiments

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

Mandate Penalty Introduction

As described in the text, the mandate penalty went in to effect in December 2007. 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 that plan’s total enrollment for the income group in June 2008. We use data only up to June 2011 because of significant changes in the prices and availability of the cheapest plans that took effect in July 2011. We estimate an expanded version of the difference-in-difference specification shown in the text in Section 3.1:

$$\begin{aligned} NewEnroll_{g,t} = & \alpha_0 + \sum_{m \in DM} (\beta_{0m} + \gamma_{0m} \cdot MandIntro_t) \cdot 1_m + \delta_0 \cdot X_t \\ & + \left(\alpha_1 + \sum_{m \in DM} (\beta_{1m} + \gamma_{1m} MandIntro_t) \cdot 1_m + \delta_1 \cdot X_t \right) Treat_g + \varepsilon_{g,t}, \end{aligned}$$

where DM is the calendar months December through March, 1_m is an indicator for month m , $MandIntro_t$ is a dummy for the mandate penalty introduction period (Dec. 2007 to March 2008), X_t is a vector of time polynomials and CommCare-year dummies,⁶¹ and $Treat_g$ is a dummy for the treatment group. The difference-in-difference coefficients of interest are the γ_{1m} ’s. In our main specifications, we use two income groups: 150-300% as the treatment group and below 100% as the control group. We also break down estimates for the treatment group separately by 50% of poverty group.

Table 8 presents the regression results. Column (1) starts with a baseline single-difference specification (i.e., without the control group) that estimates the effect 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 difference-in-difference estimates. Finally, Column (3) adds dummies for December-March in all years, forming the triple difference specification (shown in the equation above) that nets

⁶¹The CommCare-year starts in July, so these dummies will not conflict with the treatment months of December to March.

Table 8: Introduction of the Mandate Penalty.

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment		All Treatment Groups			Income Group (% of Poverty Line)		
		(1)	(2)	(3)	150-200%	200-250%	250-300%
Group x Dec2007		0.112*** (0.003)	0.110*** (0.006)	0.103*** (0.007)	0.101*** (0.007)	0.113*** (0.007)	0.091*** (0.008)
x Jan2008		0.073*** (0.004)	0.067*** (0.006)	0.069*** (0.006)	0.061*** (0.007)	0.085*** (0.007)	0.086*** (0.007)
x Feb2008		0.043*** (0.005)	0.033*** (0.006)	0.033*** (0.006)	0.026*** (0.007)	0.045*** (0.006)	0.051*** (0.006)
x Mar2008		0.025** (0.006)	0.027*** (0.006)	0.020*** (0.007)	0.020** (0.008)	0.020*** (0.006)	0.022*** (0.007)
Total		0.253*** (0.017)	0.237*** (0.023)	0.225*** (0.024)	0.208*** (0.026)	0.263*** (0.022)	0.250*** (0.024)
Control Group (< 100% Poverty)			X	X	X	X	X
Dummies for Dec-March				X	X	X	X
Observations	51		102	102	102	102	102
R-Squared	0.969		0.923	0.925	0.927	0.922	0.920
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 and month (from April 2007 to June 2011). All specifications include CommCare-year dummy variables and fifth-order time polynomials, separately for the treatment and control group. (The CommCare-year starts in July, so these dummies will not conflict with the treatment months of December to March.) Columns (2) adds the < 10% poverty group as a control. Column (3) also includes dummy variables for all calendar months of December-March, separately for the treatment and control group, to perform the triple-difference. Columns (4)-(6) do the same regression as Column (3) separately by income group. See the note to Figure 1 for the definition of new enrollees and the cheapest plan.

out general trends for those months in other years. The final three columns take the final, triple-difference specification and breaks down the analysis by limiting the treatment group to narrower income groups (but keeping the control group, 100% poverty, unchanged).

Despite the relatively small number of group-month observations, all the relevant coefficients are statistically significant. In our preferred triple difference estimates in column (3), the mandate penalty increases enrollment in the cheapest plan by 22.5% of its steady state size. When we break these results down by income group, the coefficients are slightly larger for higher income groups – about 25% instead of 21% – who faced higher mandate penalties. The mandate penalties as of January 2008 were \$17.50 for the 150-200% poverty group, \$35 for the 200-250% poverty group, and \$52.50 for the 250-300% poverty group. Given these penalties, we can use our estimates to calculate semi-elasticities. The estimates imply that each \$1 increase in the mandate penalty raised demand by 1.19% for the 150-200% of poverty group, 0.75% for the 200-250% poverty group, and 0.48% for the 250-300% of poverty group, with a weighted average of 0.97%.

We have interpreted the increases in enrollment as being the result of the permanent \$17.50–\$52.50 monthly mandate penalty that went into effect in January 2008. But the 2007 uninsurance penalty – forfeiting the state tax personal exemption, with a value of \$219 – was assessed based on coverage status in December 2007, making that month’s effective penalty much larger. Technically, individuals who applied for CommCare in 2007 and were enrolled on January 1, 2008, did not owe the penalty for 2007. But since December 31, 2007, was the main advertised date for assessing the 2007 penalty, we want to make sure that the larger effective penalty for that month is not driving the results.

To test this, we note that 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 in January 2008. Figure 4a and 4b plot the probability of exiting the market within 1 month and 6 months of initial enrollment for each entering cohort of enrollees. The graph shows that exit probabilities are no higher for people entering CommCare in December 2007 than for nearby months. (Note that the large spike in one-month exits for the March 2008 cohort is due to an unrelated income verification program.)⁶² This analysis suggests that consumers were not enrolling for just December to avoid the larger penalty and leaving soon afterward.

Premium Decrease Experiment

Our second natural experiment addresses a potential 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 and also applied to the control group (again without financial enforcement). However to the extent there is a stigma specifically from paying a fine for non-coverage, this concern is valid.

Our second experiment is a decrease in the enrollee premium of all plans – via a reduction

⁶²The income verification program took effect in April 2008 for individuals above 150% of poverty. The program uncovered a large number of ineligible people, who were dis-enrolled in April 2008 and subsequent months. This also explains the upward trend in exits within 6 months leading up to April 2008.

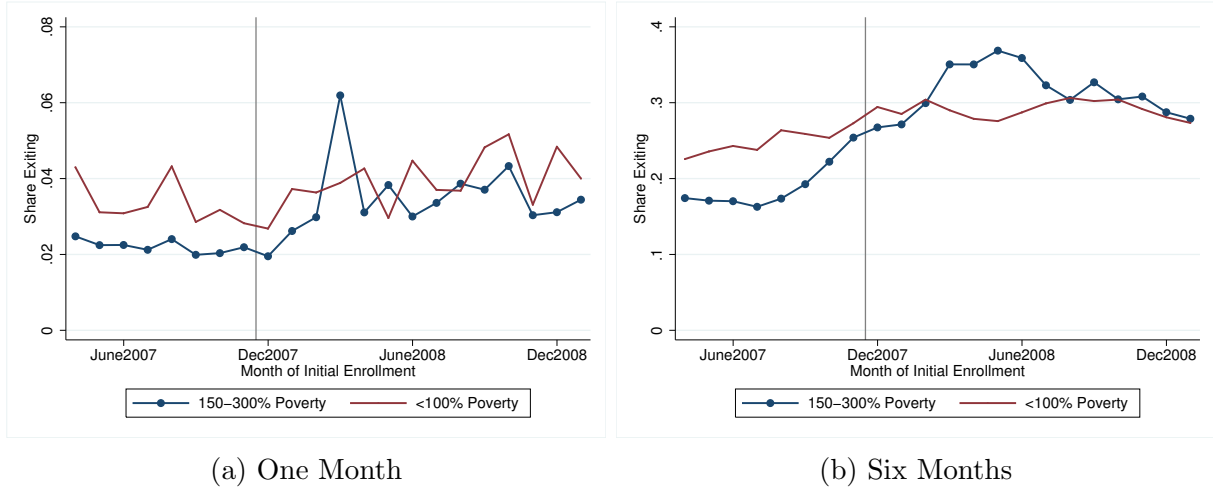


Figure 4: Share of New Enrollees Exiting Within the Specified Number of Months.

NOTE: These graphs show the rate of exiting CommCare coverage within (a) one month and (b) six months 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.

in the “affordable amount” that determines the post-subsidy premium of the cheapest plan – that occurred in July 2007. Although prices and contracts were fixed from the start of CommCare (in November 2006) until June 2008, the regulator decided to increase subsidies for certain groups in July 2007 to make insurance more affordable. For the 150-200% of poverty group, the affordable amount fell from \$40 to \$35 – which meant that the monthly premium of all plans fell by \$5. For consumers between 100-150% of poverty, the affordable amount was \$18 for the first half of 2007 and premiums were \$18-\$74.22. 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. So the aggregate effect of these changes is a lower bound on the effect of just lowering the affordable amount.

As a control group, we use the 200-300% poverty group, whose affordable amounts were essentially unchanged in July 2007.⁶³ We exclude the below 100% poverty group from our controls 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.

Table 9 presents the regression results. In Columns (1) and (4) we look at the single difference for the treatment group relative to trend (i.e., without a control group). Column

⁶³The affordable amount for 200-250% poverty was unchanged and that 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 bias our estimates downward.

Table 9: Decrease in the Affordable Amount.

Effect on New Enrollees in Cheapest Plan / June 2008 Enrollment		100-150% Poverty			150-200% Poverty		
		(1)	(2)	(3)	(4)	(5)	(6)
Treatment Group x July2007		0.064*** (0.005)	0.060*** (0.008)	0.063*** (0.008)	0.013*** (0.005)	0.017** (0.007)	0.018** (0.008)
x Aug2007		0.052*** (0.005)	0.031*** (0.008)	0.035*** (0.009)	0.029*** (0.005)	0.016** (0.007)	0.022*** (0.008)
x Sep2007		0.022*** (0.005)	0.022 (0.008)	0.016* (0.008)	0.013** (0.005)	0.011* (0.006)	0.015** (0.007)
x Oct2007		0.062** (0.005)	0.054*** (0.008)	0.055*** (0.008)	0.008 (0.005)	0.009 (0.006)	0.009 (0.006)
Total		0.200*** (0.021)	0.156*** (0.031)	0.169*** (0.032)	0.063*** (0.019)	0.052** (0.025)	0.063** (0.027)
Control Group (200 – 300% Poverty)			X	X		X	X
Dummies for July-October				X			X
Observations	52	104	104	104	52	104	104
R-Squared	0.988	0.981	0.981	0.981	0.960	0.957	0.958
Robust s.e. in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$							

NOTE: This table reports difference-in-difference regressions of the effect of the decrease in the affordable amount. The dependent variable is the number of new CommCare enrollees who chose the cheapest plan, scaled by total group enrollment in that plan in June 2008. There is one observation per income group 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.) Columns (2) and (5) add the 200-300% poverty control group. Columns (3) and (6) add controls for July-October of other years, separately by treatment and control group. 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 mandate penalty introduction. See the note to Figure 1 for the definition of new enrollees and the cheapest plan.

(2) and (5) then add the 200-300% poverty control group to form the difference-in-difference estimates. Columns (3) and (6) do a triple-difference, further netting out changes in July-October of other years. In this triple difference specification, for enrollees 100-150% of poverty, we find a 16.9% increase in the cheapest plan's demand. Dividing by the \$18 reduction in its premium implies a semi-elasticity of 0.94%. For the 150-200% poverty group, we find a 6.3% increase in demand. Dividing by its \$5 premium reduction yields a semi-elasticity of 1.26%. We note that this semi-elasticity is nearly identical to the 1.19% semi-elasticity for the 150-200% poverty group found in the mandate penalty introduction experiment.

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 associated Z_i) we draw a ν_i .

For each ξ (the CommCare plan utility coefficients) 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 (in the utility of uninsurance) there is also a corresponding group, h – income, demographic, region, or year. We use share moments analogous to those above in F^1 , but 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-sensitivity 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 3.1. 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 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 the observed 22.5% change in demand for the mandate penalty introduction experiment:

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

where the first probability in the equation is based on the pre-period mandate penalty (M_i^{Pre}) – which is zero – and the second probability is based on the post-period mandate penalty (M_i^{Post}) that applied from January 2008 on. We do not use the premium decrease experiment estimates because it occurred before the start of our demand estimation period (January 2008). However, because the semi-elasticity estimated from it is so similar, including it as a moment would be unlikely to affect our results.

C Additional Results

Medical Loss Ratios

Regulatory limits on the minimum fraction of premiums that must be used to cover medical care could limit the distortionary impact of price-linked subsidies. The ACA, sets a minimum medical loss ratio (MLR) of 80%. Table 10 shows estimated costs as a fraction of revenue for each firm for each policy simulation. BMC’s MLR is below the 80% threshold under maximum-price equilibrium for price-linked subsidies in 2009. So MLRs may have had some bite for avoiding the more expensive equilibrium in 2009, but would not have limited the distortion 2011.

Table 10: Medical Loss Ratios

		BMC	Celticare	Fallon	NHP	Network Health
2009	Min Price-linked	81.5%		91.2%	91.4%	89.9%
	Max Price-linked	77.2%		92.1%	92.1%	85.3%
	Fixed	82.3%		91.2%	91.3%	90.1%
2011	Price-Linked	90.2%	80.0%	89.4%	90.3%	91.5%
	Fixed	89.9%	85.0%	89.1%	89.7%	91.4%

NOTE: This table shows the estimated medical loss ratio (100*costs/revenue) for each firm under each subsidy type. The 80% threshold is exceeded only in 2009 for one firm under price-linked equilibrium that is most favorable for the firms.

When there are only two insurers, profits are higher and MLRs become more binding. Table 11 shows the estimated MLRs if there were only two insurers in the market in 2011. CeltiCare’s MLR is below 80% for price-linked subsidies whenever it has only one competitor; Network Health and BMC also have MLRs below 80% when competing against NHP (the highest cost insurer). It seems MLRs are more relevant for extremely uncompetitive markets.

Table 11: Medical Loss Ratios with Two Insurers in 2011

		When Competing Against			
		CeltiCare	Network Health	BMC	NHP
MLR of:					
(Lowest cost)	CeltiCare		75%(81%)	78%(84%)	67%(74%)
↓	Network Health	91%(91%)		82%(87%)	75%(81%)
	BMC	89%(89%)	87%(87%)		75-76% (82%)
(Higher cost)	NHP	92%(91%)	91%(92%)	88-90%(91%)	

Note: This table shows the estimated medical loss ratio ($100 \times \text{costs}/\text{revenue}$) for the insurer of a given row, when competing only against the insurer of a given column. The plans are listed in order of increasing costs: CeltiCare, Network Health, BMC, NHP. The first number in each cell is the MLR under price-linked subsidies; the number in parentheses is the MLR under fixed subsidies. There are multiple equilibria under price-linked subsidies when only BMC and NHP are in the market, so there is a (small) range of MLRs.

Additional results under uncertainty

Including insurer profits in welfare

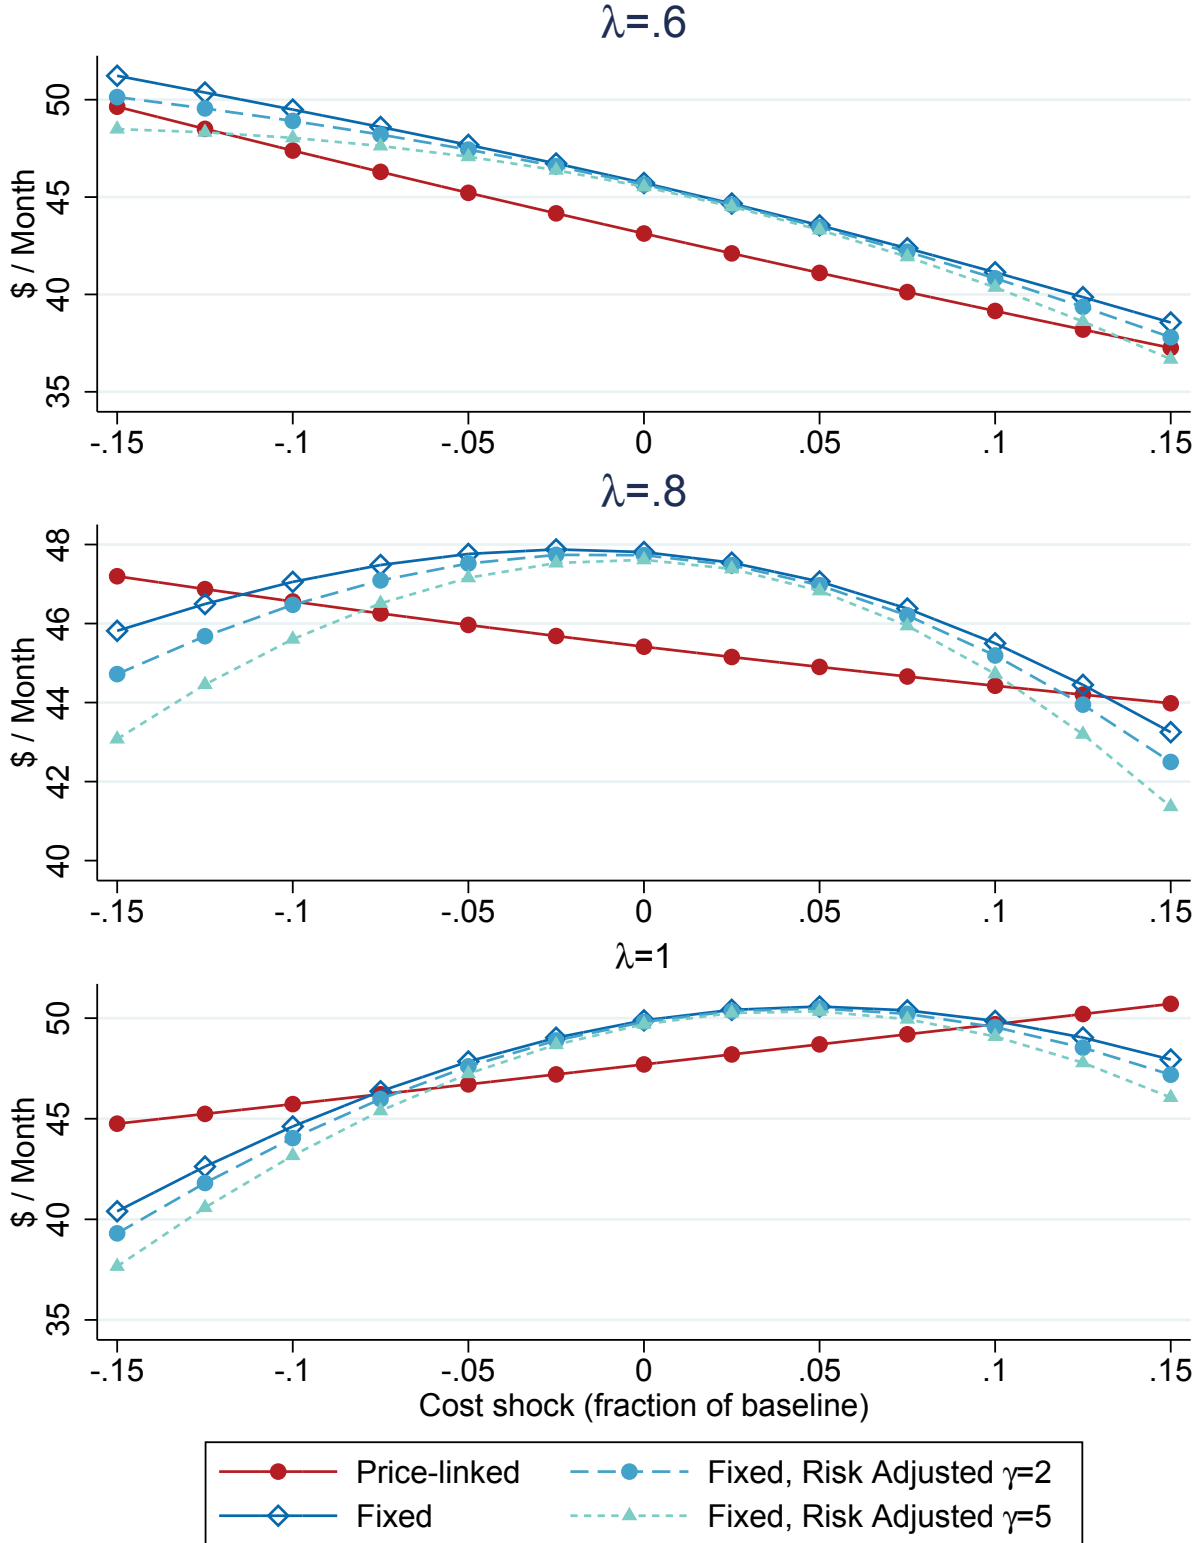
In order to properly understand the effects of uncertainty, we assumed that the fixed subsidy was set optimally for zero cost shock – and calibrated the fixed component of the externality of uninsurance, E_0 , accordingly to ensure that this optimality was true. If we include profits in the regulator’s objective function, we need to recalibrate E_0 assuming that at baseline costs the fixed subsidy maximized the new objective function. We find in practice that the calibrated E_0 is often negative. This occurs because if the regulator cares about profits, they will want to subsidize purchases even if there is no externality because price (willingness to pay) is above cost, so there are social gains to the marginal purchase. Although a negative E_0 seems non-intuitive, we want to ensure that the baseline subsidies are optimal at zero cost shock, so we proceed with the estimated values.

Figure 5 shows welfare including profits under fixed and price-linked subsidies for different cost shocks. Fixed subsidies do not do quite as well as without profits, but for $\lambda = .8$ fixed subsidies are better than price-linked for cost shocks $\leq 12.5\%$ and $\geq -10\%$.

Results for 2009

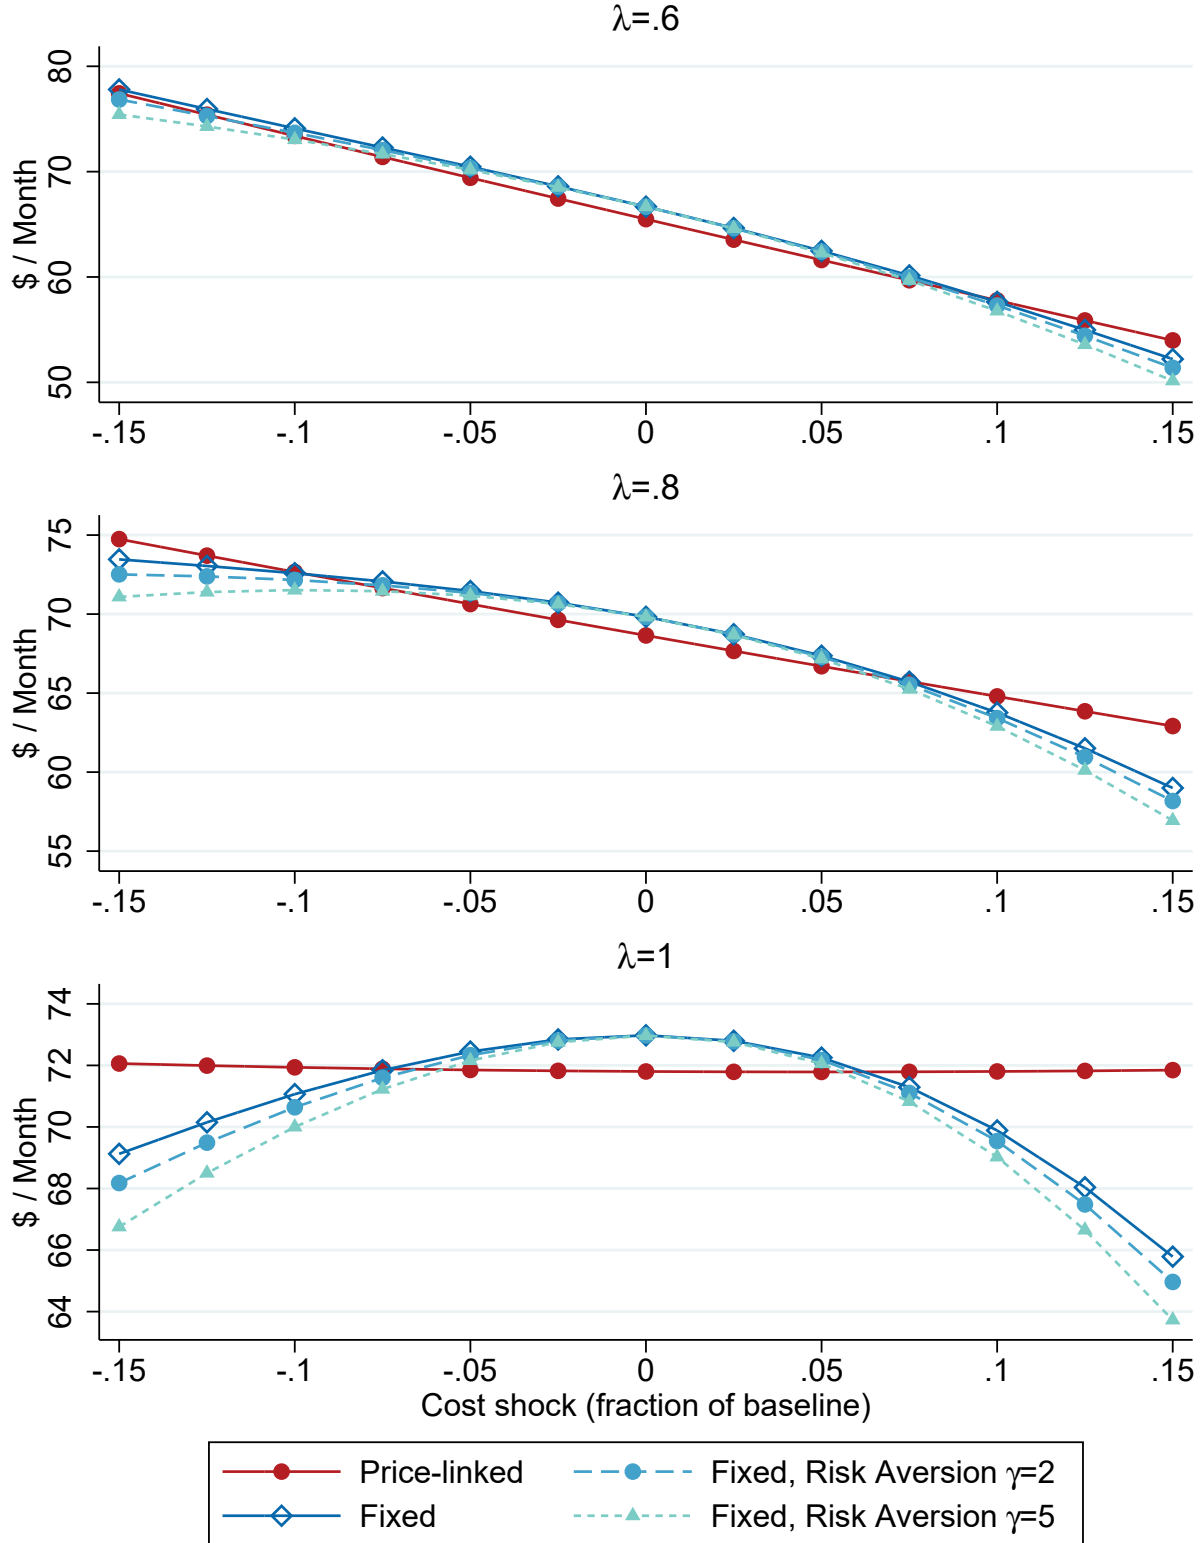
Fig 6 shows public surplus under fixed subsidies and under the minimum price-linked subsidy equilibrium in 2009, for cost shocks of different sizes. The results are qualitatively the same as 2011, but since the initial difference between the two subsidies is smaller, a smaller cost shock is required to make price-link subsidies better than fixed subsidies.

Figure 5: Public surplus plus profits under cost shocks in 2011



Note: This figure shows the sum of public surplus and profits 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. The dashed and dotted lines are risk adjusted with factors of relative risk aversion of 2 and 5, respectively.

Figure 6: Public surplus under cost shocks in 2009



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. The dashed and dotted lines are risk adjusted with factors of relative risk aversion of 2 and 5, respectively.