

Price-Linked Subsidies and Imperfect Competition in Health Insurance*

Sonia Jaffe[†]

Mark Shepard[‡]

March 19, 2019

Abstract

Policymakers subsidizing health insurance often face uncertainty about future market prices. We study the implications of one policy response: *linking* subsidies to prices, to target a given post-subsidy premium. We show that these price-linked subsidies weaken competition, raising prices for the government and/or consumers. However, price-linking also ties subsidies to health care cost shocks, which may be desirable. Evaluating this tradeoff empirically, using a model estimated with Massachusetts insurance exchange data, we find that price-linking increases prices 1-6%, and much more in less competitive markets. For cost uncertainty reasonable in a mature market, these losses outweigh the benefits of price-linking.

*The authors would like to thank Amitabh Chandra, David Cutler, Leemore Dafny, David Dranove, Keith Ericson, Jerry Green, Jon Gruber, Kate Ho, Scott Kominers, Tim Layton, Robin Lee, Neale Mahoney, Luca Maini, Tom McGuire, Ariel Pakes, Daria Pelech, Maria Polyakova, Jonathan Skinner, Amanda Starc, Pietro Tebaldi, seminar and conference participants at Boston University, Harvard economics department, Harvard Medical School, University of Chicago economics department, the NBER Health Care Program, ASHEcon (2016), IIOC (2017), and the Healthcare Markets Conference (2017), the editor and several anonymous referees 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.

An earlier version of this paper circulated under the title “Price-Linked Subsidies and Health Insurance Markups”

[†]Microsoft, 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 that offer a choice among subsidized private plans. Long used in Medicare’s private insurance option (Medicare Advantage), this approach has also been adopted for Medicare’s drug insurance program (Part D), the insurance exchanges created by the Affordable Care Act (ACA), and national insurance programs in, for example, the Netherlands and Switzerland. These programs aim to leverage the benefits of choice and competition while using subsidies to make insurance more affordable and encourage enrollee participation.

We study the implications of a key subsidy design choice that arises in these market-based programs: whether to *link* subsidies to prices set by insurers. Many programs take this “price-linked” subsidy approach. For instance, Medicare Part D links subsidies to market average prices, and the ACA exchanges link subsidies to the second-cheapest “silver” tier plan. Other programs choose to set subsidies at specific levels or based on external benchmarks not controlled by insurers – an approach we call “fixed” subsidies. Medicare Advantage, for example, sets subsidy benchmarks based on an area’s lagged costs in traditional Medicare. Shifting towards fixed, inflation-indexed subsidies was also a key feature in proposed alternatives to the ACA (e.g., the plan by former Speaker Paul Ryan (Ryan, 2017)). This choice between subsidy types is also relevant to some non-insurance markets where the government subsidizes individuals’ purchases from private providers, such as school vouchers and housing subsidies.

While linking a market-wide subsidy to prices is convenient when there is uncertainty about market prices, it also raises concerns about competition. Despite the prevalence of these price-linked subsidies, the economic incentives and tradeoffs involved are not well understood. In this paper, we argue that price-linked subsidies should be thought of as involving a basic tradeoff. On the one hand, they weaken price competition in imperfectly competitive markets, leading to higher prices and government costs than under fixed subsidies. 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 or political economy constraints. We use a simple theoretical model to show the intuition for these effects and to formalize the conditions under which they occur. We then use administrative data from Massachusetts’ pre-ACA insurance exchange for low-income adults to analyze the empirical magnitude of these effects using both a sufficient statistics approach and a structural model.

We first use a simple model to formalize the competitive implications of price-linked subsidies. On one level, the competitive concern is obvious: if an insurer can get a higher subsidy by setting a higher price, it is incentivized to raise its price. While this logic clearly applies to designs like percentage subsidies in which each plan gets a *different subsidy dollar amount* determined by its own price, its application is less obvious in settings like the ACA

that purposely give all plans the *same subsidy amount* – a key “managed competition” principle emphasized by Enthoven (1988). In this case, unless there is a monopolist (or firms collude), the logic is more complicated. Even if an insurer knows its price will be pivotal for the subsidy, its price affects both its own *and* its competitors’ subsidy. Indeed, in a market with an equal subsidy applying to *all* relevant options, price-linked subsidies would *not* distort price competition. We show that the competitive effect of price-linking depends on the presence of an “outside option” to which the subsidy does not apply. In the ACA, a higher subsidy decreases the cost of buying a market plan relative to the outside option of 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 change the *slope* of the demand curve in this way that clearly distorts prices upwards.¹

We derive a simple first-order approximation to the effect of price-linked subsidies on the equilibrium price of the subsidy-pivotal plan. Price-linked subsidies effectively remove the outside option as a price competitor for this plan: as the plan raises its price, it becomes more expensive relative to other plans, but its price relative to the outside option is unchanged. As a result, the distortion is larger when the pivotal plan competes more directly with the outside option (a larger cross-elasticity with respect to the price of the outside option – e.g., the mandate penalty). The distortion is also larger when the plan has more market power (a smaller own-price elasticity of demand). These conditions seem particularly applicable to the exchanges set up by the ACA. The outside option of uninsurance is likely to be important, since a large share of the subsidy-eligible population remain uninsured in both Massachusetts (as we show) and the ACA (Avalere Health, 2015). Furthermore, the ACA exchanges are highly concentrated, with about half of consumers living in markets with just one or two insurers in 2018 (Kaiser Family Foundation, 2018).

Next, having shown how price-linked subsidies distort prices, we ask whether price-linked subsidies might nevertheless be the right choice under certain circumstances. We argue that the case for price-linking relies on regulators facing at least one of two forms of *constraints*: informational constraints (uncertainty) or political economy constraints. Absent these limitations, a simple argument shows why fixed subsidies dominate: regulators could predict equilibrium prices – and therefore the price-linked subsidy amount – and replicate it with a fixed subsidy of equal size. This fixed subsidy would result in lower prices, greater coverage,

¹Our analysis suggests a simple alternate policy that could preserve price-linked subsidies while mitigating the distortion: link the mandate penalty (for uninsurance) to prices in the *opposite* direction, thus making the *total* incentive to buy insurance (subsidy plus penalty) invariant to prices. We discuss this policy in Section 1.3, though we note that 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 competitive distortion.

and a pure gain for consumers. The two types of constraints short-circuit this simple argument in different ways. Uncertainty limits the regulator’s ability to predict prices, while political economy constraints, such as regulatory capture, can prevent regulators from optimally using their information to adjust subsidies.

We discuss these rationales for price-linked subsidies in Section 5. We also pay special attention to formalizing and evaluating the uncertainty rationale. On its own, cost or price uncertainty is not sufficient to justify the higher prices from price-linked subsidies; the cost uncertainty must translate into uncertainty about the *optimal* subsidy level. Specifically, the optimal subsidies must be larger in states of the world where prices are unexpectedly high. We propose (and include in our model) two reasons why the regulator may want subsidies to be higher when prices are higher. First, the government may wish to insure low-income enrollees against the risk of price shocks. This relates to an explicit rationale for the ACA’s subsidy design: ensuring that post-subsidy premiums are “affordable” regardless of insurance prices. Second, higher prices may reflect higher underlying health care costs, which in turn relate to a major rationale for subsidizing insurance: the costs of uncompensated care received by the uninsured. These costs of care, borne by hospitals and clinics, are a negative externality of uninsurance: when deciding whether to forgo insurance, the uninsured do not consider the social cost of the care they receive.² To make consumers (partially) internalize this externality, the regulator wants insurance subsidies (or the penalty for uninsurance) to track this externality. Price-linking provides an indirect way to achieve this if both uncompensated care costs and insurance prices generally rise with market-level health care cost shocks. Importantly, this rationale requires that insurers have information on costs beyond what regulators have (or can use) when setting subsidies; otherwise, prices have no additional signal value.

Our model of insurer competition provides a framework for studying the pricing distortion and welfare tradeoffs of price-linked subsidies empirically. To do so, we draw on administrative plan enrollment and claims data from Massachusetts’ pre-ACA subsidized insurance exchange, supplemented with data on the uninsured from the American Community Survey. An important precursor to the ACA, the Massachusetts market lets us observe insurance demand and costs for a similar setting and low-income population. Like the ACA, Massachusetts used price-linked subsidies throughout its history, with subsidies set based on the cheapest plan’s price. We are unaware of any public market where there has been a policy transition between price-linked and fixed subsidies that would let us observe coun-

²There is growing evidence on the importance of uncompensated care for low-income people (Garthwaite et al., 2017; 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.

terfactual market outcomes to measure the policy effect directly. Instead, we use our model of insurer competition to estimate the effects and tradeoffs of price-linked subsidies based on parameters estimated from the Massachusetts data. Although the model is necessarily stylized, this method provides a way to evaluate the empirical importance of our theoretical points.

We use two methods for this empirical exercise. First, we use a sufficient statistics approach (Chetty, 2009), drawing on natural experiments in Massachusetts to estimate the key statistics that enter our first-order approximation to the pricing distortion. The main natural experiment is the introduction of the mandate penalty in December 2007. Using income groups exempt from the penalty as a control group, we estimate that each \$1 increase in the relative monthly price of uninsurance raised demand for the cheapest plan by about 1%. We also use a difference-in-differences approach based on within-plan differential price changes to estimate an own-price semi-elasticity of demand of -2.16%.

Our second empirical method is to estimate a full structural model of demand and cost and use the estimated parameters to simulate equilibrium in the insurance exchange under price-linked and fixed subsidies. While this approach necessarily involves more assumptions, it lets us go beyond price effects to estimate welfare impacts and simulate the tradeoffs involved in the presence of cost uncertainty and under different market structures. It also takes into account strategic interactions and adverse selection, which are not in the reduced form approximation. An important strength of our structural estimates is that we use difference-in-difference estimates from the natural experiments to identify the variance in the random component of individuals' utility of insurance, which is key to determining the relevant demand substitution patterns.

Across both methods, we find three sets of results. First, we estimate that price-linked subsidies raise prices by a non-trivial amount. Using the sufficient statistics method, we find that price-linked subsidies would raise the pivotal plan's price by about 9% (\$36 per month) in a market with similar market primitives as the Massachusetts exchange. The estimates from the structural model simulations are somewhat smaller (1-6% of prices under fixed subsidies, or \$4-26 per month) but generally in the same range. Based on our 2011 estimates, switching from price-linked to fixed subsidies could achieve either the same insured rate at 6.1% lower subsidy cost, or 1.3% greater insurance coverage at the same cost. Although modest, these effects imply meaningful increases in government costs. For instance, the \$24/month subsidy distortion (from our structural model simulations for 2011) would translate into \$46 million in annual subsidy costs for Massachusetts and over \$3 billion if extrapolated nationally to the ACA. Of course, the actual distortion for the ACA might be higher or lower than this figure because of the various ways ACA markets and policies differ from Massachusetts.

Second, we show that the pricing distortion depends critically on the intensity of market competition, particularly at the low-cost end of the market (since subsidies are set based on the lowest price in Massachusetts, and the second-lowest silver plan price in the ACA). We find that the effect of price-linking is about twice as large (6-12% of baseline price) when we simulate markets with just two competitors, as in many ACA markets. The largest distortions occur when the gap between the costs of the two insurers is large – i.e., if a low-cost plan competes against a high-cost plan. We also find that the cheapest plan’s price with fixed subsidies and one competitor is similar to its price with price-linked subsidies in the baseline market. This suggests that with four major insurers (as was the case in Massachusetts), switching to price-linked subsidies has a comparable effect on the price of the cheapest plan as removing all but one of the plan’s competitors.

Third, we evaluate market-level cost uncertainty as a rationale for price-linked subsidies, assuming an optimizing regulator who can flexibly set subsidies in each market. We find that cost uncertainty must be fairly high for the benefits of price-linked subsidies to outweigh the losses from higher prices. Even under the assumptions most favorable to price-linking, fixed subsidies do better for all cost shocks between -12.5% and +15%. 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 percentage points for annual growth (or 4.8% points over three years), making cost shocks larger than 12.5% unlikely. However, in new markets where regulators cannot rely on lagged cost information reported by insurers or enrollment has not reached equilibrium levels, larger shocks are more likely.

This analysis casts doubt on cost uncertainty as a sufficient explanation for price-linked subsidies. We therefore conclude that a rationalization of price-linking would rest on political economy factors, which we discuss in Section 5. These factors could include: limits on regulators’ bandwidth to optimize fixed subsidies based on local health care costs, an extreme concern that post-subsidy prices always be “affordable,” and the potential for regulatory capture in a fixed subsidy system. In practice, an important political economy issue for the ACA has been the presence of a federal regulator opposed to the law’s original objectives. In the face of regulatory action that contributed to sharply higher prices, price-linked subsidies have stabilized post-subsidy premiums and mitigated what might have been an adverse selection death spiral.

Related Literature Our paper is related to the literature on the rationale for and effects of health insurance subsidies and mandate penalties. Two prominent rationales 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); see Gruber (2008) for a review. It is most closely connected to a small but growing body of research studying the (often unintended) competitive implications of subsidy policies – including Decarolis (2015) and Decarolis et al. (2015) studying Medicare Part D, Curto et al. (2014) studying Medicare Advantage, and Liu and Jin (2015) studying the Federal Employees Health Benefits Program. Cutler and Reber (1998) consider the adverse selection cost of increasing competition in ‘markets’ for employer-provided health insurance. 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 is also part of broader literature estimating equilibrium under imperfect competition in health insurance markets. This includes work on the Massachusetts insurance exchanges (Ericson and Starc, 2015, 2016), Medicare insurance markets (e.g., Town and Liu, 2003; Starc, 2014; Curto et al., 2014; Polyakova, 2016; Ho et al., 2017), and other settings (e.g., Handel, 2013; Handel et al., 2015; Ho and Lee, 2017). Recent work by Shepard (2016) and Finkelstein et al. (2017) use the same CommCare setting and data to estimate insurance demand and cost primitives. But they are focused on different research questions: Shepard (2016) focuses on hospital networks and adverse selection, while Finkelstein et al. (2017) is focused on quantifying willingness to pay for insurance.

The remainder of the paper is structured as follows. Section 1 uses a simple model to show our theoretical analysis of price-linked subsidies; it also presents a welfare framework for thinking about the tradeoffs between subsidy policies. 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 effect of linking subsidies to prices. Section 4 describes the structural model, presents the estimated demand and cost parameters, and simulates the pricing equilibria and resulting welfare under each subsidy policy. Section 5 looks at how uncertainty and political considerations may make price-linked subsidies more attractive and discusses the applicability of our framework to other markets. Section 6 concludes.

1 Theory

We adapt a standard discrete choice model of demand to allow for a mandate penalty and 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, in which each insurer offers a single plan. In Appendix A.1, we show that the basic logic of the distortion carries through to a more general model allowing for multi-plan insurers (as in the ACA).

Insurers $j = 1, \dots, J$ each offer a differentiated plan 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.³ 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 must pay a 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 post-subsidy 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. Handel et al., 2015), we assume fixed plan attributes and focus instead on pricing incentives conditional on plan design. For simplicity, we model the case where the exchange’s risk adjustment completely accounts for adverse selection, so each insurer has a net-of-risk-adjustment marginal cost c_j that does not depend on prices. In Appendix A.2 we show how adverse selection (beyond what is adjusted for by the exchange) interacts with the subsidy structure, but the basic intuition is the same.

The insurer profit function is:

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

A necessary condition for Nash equilibrium is that each firm’s first-order condition is satisfied:⁴

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

This differs from standard oligopoly pricing conditions in that the firm’s price P_j enters

³The 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 include income-specific subsidies in our empirical model.

⁴These 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 networks. Thus, our main theoretical point about price-linked subsidies holds when quality is endogenous, though of course there may also be effects on quality and cost levels not captured in our model.

consumer demand *indirectly*, through the subsidized premiums, $P_j^{cons} = P_j - S(P)$. As a result, the term dQ_j/dP_j (a total derivative) combines consumer responses to premium changes and any indirect effects on demand if P_j affects the subsidy or mandate penalty (via the regulatory formula). The total effect on demand of raising P_j 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 Equation (2) by imposing an assumption that is standard in most discrete choice models: that (at least locally) price enters the utility function linearly.⁵ In our context, this also assumes that the uninsured pay or expect to pay the mandate penalty; this latter assumption is supported by the empirical finding in Section 3.1 that demand responds similarly to an increase in the mandate penalty and a decrease in all premiums. These assumptions imply that only *price differences*, not levels, matter for demand. Thus, raising all prices (and 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 either S or M is linked to prices.

How much the price-responsiveness of demand is attenuated by the price-linking of the subsidy or mandate penalty depends on the magnitude of $\partial Q_j/\partial M$. Intuitively, this is because neither S nor M affects 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; if substitution is high, the effect will be large. Thus, a key goal of our empirical work is to estimate $\partial Q_j/\partial M$.

⁵This 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 approximation to a CARA utility function, in which risk aversion is constant with income.

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

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

$$Markup_j^F \equiv P_j - c_j = \frac{1}{\eta_j} \quad \forall j,$$

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 mandate penalty). 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:

$$Mkup_{\underline{j}}^{\text{PLink}} \equiv P_{\underline{j}} - c_{\underline{j}} = \frac{1}{\eta_{\underline{j}} - \eta_{\underline{j},M}},$$

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. The first-order condition for the non-pivotal firms is the same as in the fixed-subsidy case.

If the decrease in demand elasticity for the pivotal plan is large enough that it wants to price so high that it would no longer be pivotal, the equilibrium conditions are more complicated. In this case there will generally be a range of possible equilibria, each with a tie among multiple cheapest plans, an issue we address in our simulations.

Comparing Fixed and Price-Linked Subsidies

Given the equations for the markup of the pivotal plan under each subsidy framework, we can look at the difference between the two. 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 absolute and percent increase in markups between fixed and price-linked subsidies:

$$Mkup_{\underline{j}}^{\text{PLink}} - Mkup_{\underline{j}}^{\text{F}} = \frac{\eta_{\underline{j},M}}{\eta_{\underline{j}} (\eta_{\underline{j}} - \eta_{\underline{j},M})} > 0, \quad (4)$$

$$\frac{Mkup_{\underline{j}}^{\text{PLink}} - Mkup_{\underline{j}}^{\text{F}}}{Mkup_{\underline{j}}^{\text{F}}} = \frac{\eta_{\underline{j},M}}{(\eta_{\underline{j}} - \eta_{\underline{j},M})} > 0, \quad (5)$$

which are generally positive because $\eta_{\underline{j},M} > 0$ and $\eta_{\underline{j}} > 0$ under standard demand assumptions, and $\eta_{\underline{j}} > \eta_{\underline{j},M}$ as long as there is at least one other in-market option besides \underline{j} . 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.⁶

Price-linked subsidies lower the effective price sensitivity faced by the pivotal (cheapest) plan, leading to a higher equilibrium markup than under fixed subsidies. 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 pivotal. In this case, the distortion directly effects 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

⁶This 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.

about others’ costs), the distortionary term $\eta_{j,M}$ 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 direct effects on other plans’ prices.

1.2 Welfare

To move from prices to welfare, we need to specify a regulator’s objective function, including a rationale for subsidizing insurance instead of just providing cash transfers.⁷ We base our analysis on a consumer surplus standard (common in antitrust analysis) adjusted for government spending (subsidy cost net of mandate revenue) and for an externality of uninsurance. Together, these form a regulatory objective that we refer to as “public surplus.” This objective captures two rationales for subsidizing insurance: adverse selection and negative externalities of individuals being uninsured.

We allow for two types of externalities of uninsurance. 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., 2017; Finkelstein et al., 2015). The second type of externality is a pure (paternalistic) social disutility of people lacking insurance, consistent with much of the political language motivating subsidies. In practice, we treat this social disutility as a free parameter that we calibrate to make the observed subsidy levels (or post-subsidy consumer premiums) optimal.

The aggregate externality and government spending are determined by consumer demand; consumer demand depends on the mandate penalty M , and the subsidy, S , which determine equilibrium prices and premiums ($P^{cons} = P - S$). The expected public surplus from individual i is

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

where D_0^i is the probability the individual does not chose insurance and E^i is the externality of that individual being uninsured, which is avoided when they purchase insurance. We think this public surplus standard is consistent with the public health care reform debate, which put little weight on insurer profits – indeed, policymakers seemed eager to constrain insurer profits with policies like medical loss ratio limits – but we also consider an alternative where the regulator values insurer profits, maximizing $PS + \sum_j \pi_j$.

⁷See Currie and Gahvari (2008) for a review of the large literature on rationales for in-kind benefits.

Role of Uncertainty

Absent uncertainty – i.e., with full information about costs, demand, and pricing behavior – the regulator can predict prices and therefore the subsidy amount that will emerge under a price-linked subsidy policy. The regulator could then set a fixed subsidy to replicate the amount of the price-linked subsidy, leading insurers to lower prices. The resulting lower premiums (since the subsidy is unchanged) lead to gains for consumers and more people buying insurance. If the mandate penalty had been set optimally, the regulator would be indifferent between the costs and benefits of additional insurance purchases. The gains to consumers (at insurers' expense) would be desirable as long as the regulator cares less about insurer profits than about consumer surplus; moreover, insurer profits may actually be higher under fixed subsidies due to the higher enrollment due to lower consumer premiums (a case we find empirically in some simulations).

Absent uncertainty, the following three conditions are jointly sufficient to ensure welfare is higher under fixed subsidies than under price-linked subsidies:

- (i) *Profits of the pivotal plan are increasing in the level of the subsidy.* This is a weak condition. It holds under either adverse selection or no selection. This is all that is necessary to ensure that price-linking increases the price of the cheapest plan.
- (ii) *Each firm's optimal price is increasing in the marginal cost of the subsidy-pivotal firm.* This is related to prices being strategic complements. (It holds trivially with logit demand.)
- (iii) *Welfare is decreasing in each plan's price.* This certainly holds if the regulator puts sufficiently little weight on firm profits and can hold more generally.

See Appendix A.3 for a more detailed discussion.

Under uncertainty, the regulator cannot perfectly predict equilibrium prices, so it cannot replicate the price-linked subsidy amount using a fixed subsidy. Price-linked subsidies may be better, despite the 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. 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 will likely 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.

With cost uncertainty, price-linked subsidies also have the benefit of stabilizing post-subsidy consumer prices, transferring the risk of cost shocks to the government; we think this is reflected in the political rhetoric around “affordability.” The first-order effect on consumers when premiums are higher (or lower) than expected is captured in the consumer surplus discussed above. When consumers are risk averse, there is an additional cost from the change in the marginal utility of consumption, which is approximately

$$\frac{\gamma (\Delta P^{cons})^2}{x \cdot 2}, \quad (6)$$

where, ΔP^{cons} is the premium difference, γ is the coefficient of relative risk aversion and x is consumption. (See Appendix A.4 for a derivation.) We allow for both of these benefits of price-linked subsidies in our empirical simulations of welfare under uncertainty.

1.3 Alternate Policy: Price-Linked Subsidies and Mandate Penalty

Our model suggests an alternative subsidy structure to eliminate the price distortion while still guaranteeing affordability of post-subsidy premiums, as with price-linked subsidies. 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, the subsidy policy does not diminish the effective slope of the demand curve. The government could set M_0 so that in expectation, $M(P)$ would equal 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. (Fixed subsidies do the same by holding *both* S and M fixed with prices.) Since plans’ pricing cannot impact this net incentive for insurance, the distortion of price-linking goes away. However, this policy loses most of the benefits of price-linked subsidies discussed above. The net incentive to buy insurance is not linked to the externality of uninsurance (since it does not vary with prices). The policy removes the costly variation in consumer premiums that fixed subsidies have, but the variation is transferred to the mandate penalty, which now varies with prices; the price-risk is transferred from consumers who buy insurance to those who do not. Of course, this policy only eliminates the distortion if the net incentive, $S + M$, is what matters for insurance demand, not the level of S and M individually. If 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, therefore, we focus on comparing fixed and price-linked subsidies, rather than this alternate policy.

2 Setting and Data

2.1 CommCare Setting

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, known as Commonwealth Care (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 level (FPL) and lacking 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 suitable for a model of imperfect competition.

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, such as customer service. Importantly, these flexible quality attributes apply equally to enrollees in all income groups, a fact we use in estimating demand.

In CommCare, subsidies are linked to the price of the cheapest plan so that this plan costs an income-specific “affordable amount,” which varies between \$0 and \$116. A consumer’s premium for a plan is the plan’s price (set by the insurer) minus the subsidy for that consumer’s income group. In addition (and unlike the ACA), CommCare applied special subsidies for the below-100% of poverty group that made all plans free, regardless of their pre-subsidy price. We use this fact to aid demand estimation – since this group can purchase the same plans but faces different relative prices than other income groups.

Since CommCare’s eligibility criteria exclude people with access to other sources of health insurance, eligible individuals’ relevant outside option is uninsurance.⁸ The price of uninsurance is the mandate penalty after its introduction in late 2007 (as discussed in Section 3.1). 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).

⁸In 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.

People regularly move in and out of eligibility for insurance through CommCare based on factors like losing/getting a job or a change in income (as a result, the median duration per enrollment spell is about 13 months). In our model, we treat eligibility as exogenous. When someone becomes eligible, they choose whether to enroll in CommCare and if so, which plan to choose. While enrolled, consumers can switch plans once a year during open enrollment; in practice, however switching rates are quite low (about 5%). Enrollees remain in CommCare until they either lose eligibility or choose to leave the market to become uninsured; we do not observe why they leave.

2.2 Data and Samples for Structural Model

Administrative data from the CommCare program⁹ let us observe (on a monthly basis) the set of participating members, their demographics, the plans and premiums available to them, 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 insurance exchanges with plan choice and cost data linked at the individual level.

We supplement the CommCare data with data on the uninsured from the American Community Survey (ACS) in order to get a dataset of CommCare-eligible individuals, whether or not they chose to purchase insurance.¹⁰ For the ACS data, we restrict the sample to Massachusetts residents who are uninsured and satisfy CommCare’s eligibility criteria based on age, income, and U.S. citizenship. We use the ACS’s weights to scale up to a population estimate of the uninsured. For the structural model, we use CommCare data from January 2008, when the individual mandate is fully phased in, to June 2011, just prior to the start of CommCare year 2012,¹¹ when plan choice rules and market dynamics shifted considerably (see Shepard, 2016). In addition, we use CommCare data from 2007 for our reduced form estimates using natural experiments. See Appendix B.1 for more information on the data and sample construction.

We use these datasets to construct several samples to estimate and simulate our structural model of insurance demand and costs. Our sample choices are guided by the varying strengths and limitations of the data and the purposes of our empirical exercise. At a high level, our goal is to use estimates of static demand and costs for the CommCare enrollees most closely resembling the ACA’s subsidized population (those above 100% of poverty) to simulate the

⁹This 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.

¹⁰We obtained ACS data from the IPUMS-USA website, Ruggles et al. (2015), which we gratefully acknowledge.

¹¹Because 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.

competitive effects of price-linked subsidies.

First, to estimate demand, we generate a sample (and associated market shares) based on active plan choices made by new enrollees (and newly re-enrolled customers). This lets us estimate consumer preferences for plans while abstracting from inertia known to affect plan switching decisions (Handel, 2013; Ericson, 2014). Because plan switching rates are quite low, initial plan choices by new enrollees are the primary driver of market shares. Although an approximation, this simplification has been used in structural work on insurance markets (e.g. Ericson and Starc, 2015) as a way of abstracting from the complex dynamics that inertia creates.

We cannot identify the ‘newly uninsured’ in the ACS to parallel the new enrollees in CommCare. Instead, to make the data comparable, we re-weight observations in the ACS to preserve the *overall* uninsured rate calculated from the full sample of CommCare enrollees (new and existing) plus the ACS uninsured (see Appendix B.1 for details).

Second, to estimate costs, we use the full CommCare sample, both new and current enrollees. We do so both to improve precision and to ensure we match overall average costs in the market. We have explored limiting to just new enrollees for cost estimation. However, we found that new enrollees tend to incur higher costs early in their spells, making estimates from them alone not representative of overall average costs.

Finally, for our simulations of market equilibrium, we use both CommCare and ACS observations (i.e., all potential enrollees) but limit the sample to people above 100% of poverty. We do so both because this matches the subsidy-eligible population in the ACA and because below-poverty CommCare enrollees do not pay premiums, so we cannot estimate their price sensitivity of demand.

Table 1 shows summary statistics for the three samples: the full sample used for cost estimation; the new enrollees used for demand estimation; and the simulation sample. The raw sample includes 455,556 unique CommCare enrollees and 4,562 observations of uninsured individuals in the ACS. The third row of Table 1 shows the average monthly size of each group, after weighting the ACS data to scale up to a population estimate. The population is quite poor, with about half having family income below the poverty line. Consumers’ ages range from 19-64; the uninsured are slightly younger than the insured.

The bottom half of Table 1 shows additional statistics for the simulation sample. For this group, 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; the uninsured rate is below 5% in the full ACS data for Massachusetts. 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 per month; the average

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 In-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 for the simulation and costs sample. “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 for if they were to enroll in the average plan.

cost for 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 (from our cost model described in Section 4.2) 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 Estimates of Key Statistics

The first-order approximation of the price effect of price-linked subsidies depends on two key statistics: the semi-elasticity of demand for the cheapest plan with respect to the mandate penalty and its semi-elasticity with respect to its own price (see Equation 4). We use natural experiments in the Massachusetts market to estimate these two numbers and calculate the approximate price effect. Details of the estimation, additional analyses and robustness checks are in Appendix B.2.

3.1 Response to the Mandate Penalty

We use two sources of exogenous variation in the relative price of uninsurance to estimate the first key theoretical statistic: the responsiveness of demand for the cheapest plan to the price of the outside option.

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 income group’s total enrollment in the relevant plan in June 2008, when CommCare reached a steady-state size.

Several pieces of evidence suggest that the enrollment spike was caused by the financial penalties. There were no changes in plan prices or other obvious demand factors for this




Figure 1: New Enrollees in Cheapest Plan around Mandate Penalty Introduction (Dec. 2007)

Note: This graph 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 (solid blue). The <100% poverty group (dashed yellow) is a control group not subject to the penalty. 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.

group at this time. As Figure 1 shows, there was no concurrent spike for people earning less than poverty (for whom penalties did not apply), and there was no enrollment spike for individuals above 150% of poverty in December-March of other years. 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, y (< 100% poverty and 150-300% poverty), we estimate

$$NewEnroll_{y,t} = (\alpha_0 + \alpha_1 \cdot Treat_y) \cdot MandIntro_t + \xi_y + \zeta_y \cdot DM_t + \delta_y \cdot X_t + \varepsilon_{y,t},$$

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

2008, $Treat_y$ is a dummy for the treatment group (150-300% poverty), $MandIntro_t$ 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, and X_t is a vector of time polynomials and year dummies (for CommCare’s market year, which runs from July-June). The coefficient of interest is α_1 .

We estimate that the mandate penalty caused a 22.2% increase in enrollment in the cheapest plan relative to its enrollment in June 2008 (when the overall market size had stabilized). Translating this increase into a semi-elasticity of demand, we find that enrollment in the cheapest plan increases by 0.95% on average for each \$1 increase in the penalty. As we describe in Section 4, we match the estimated 22.2% increase in enrollment as a moment in the structural model.

One difference between Massachusetts and the ACA is that the ACA links prices to the second-cheapest (silver) plan, rather than the cheapest. Interestingly, if we re-estimate these regressions using enrollment in the *second-cheapest* CommCare plan as the outcome, we find very similar effects (an enrollment increase of 21.1%). This provides some evidence that substitution with the outside option is relevant for low-price plans more broadly, not just for the cheapest plan.

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 that population (an increase of 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 look at an increase in subsidies in July 2007 that lowered the premiums of all plans for enrollees earning between 100-200% of poverty. Enrollees between 200-300% of 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 responsiveness to the mandate penalty. 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.2. 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 demand for the cheapest plan 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 Own-price Semi-Elasticity

The second key statistic affecting the price distortion from price-linked subsidies is the own price elasticity of demand. Following Shepard (2016), we estimate own-price elasticity using *within-plan* variation in consumer premiums created by the exchange’s subsidy rules. Subsidies make all plans free for below-poverty enrollees, while higher-income enrollees pay higher premiums for more expensive plans. This structure also creates differential premium *changes* over time, which we use for identification. For instance, when a plan increases its price between years, its (relative) premium increases for higher income groups, but there is no premium change for below-poverty enrollees (since it remains \$0).

Figure 2 illustrates the identification strategy. It shows the average monthly market shares among new enrollees for plans that decrease prices at time zero; an analogous figure for price increases is shown in Appendix B.2. Demand is relatively stable in all months before and after the price change but jumps up at time zero for price-paying (above-poverty) enrollees. Demand among “zero-price” below-poverty enrollees is unchanged through time zero, consistent with plans having relatively little change in quality (at least on average).

To run the difference-in-difference specification analogous to this graph, we collapse the data into new enrollees in each plan (j) for each income group (y) in each region (r) in each month (t). We regress

$$\ln(\text{NewEnroll}_{j,y,r,t}) = \alpha \cdot P_{j,y,r,t}^{\text{cons}} + \xi_{j,r,t} + \xi_{j,r,y} + \epsilon_{j,y,r,t}. \quad (7)$$

where P^{cons} refers to the premium the individual pays (not the price the plan receives). It is zero for the below poverty group for all plans (in all years and regions). The plan-region-year dummies absorb changes in quality of plans over time and the plan-region-income dummies account for the fact that some income groups may differentially like certain plans.

Since we use log enrollment, α corresponds to the semi-elasticity of demand with respect to own price. We get a semi-elasticity of 2.16% when weighted by average new enrollment in the region. We can compare these estimates to Chan and Gruber (2010), who study price sensitivity of demand in this market using a different identification strategy. Though they report an own-price semi-elasticity of demand of 1.54%, they do not allow for substitution




Figure 2: Identifying the Own-Price Semi-elasticity: Market share around price decrease
 Note: This graph shows the source of identification for the own-price semi-elasticity of demand. It shows average monthly plan market shares among new enrollees for plans that decreased their prices at event time 0. The identification comes from comparing demand changes for above-poverty price-paying (new) enrollees (for whom premium changes at time 0) versus below-poverty zero-price enrollees (for whom premiums are always \$0). The sample is limited to 2008-2011, the fiscal years we use for demand estimation.

to the outside option of uninsurance. When we adjust for that, their estimates imply a semi-elasticity for the cheapest plan of 2.17%.¹²

3.3 Estimate of Pricing Distortion

These natural experiments show that insurers have some market power and there is substitution to uninsurance based on its relative price, suggesting the potential for price-linked subsidies to distort prices. Before using a structural model of insurer competition to analyze this distortion and the welfare tradeoffs between subsidy structures, we use the formulas from Section 1.1 to get an approximation of the price effect. The semi-elasticity estimates of $\eta_{j,M} = 0.95\%$ (Section 3.1) and $\eta_j = 2.16\%$ (Section 3.2) suggest that price-

¹²The in-market share of the cheapest plan is 47.3% and 54.5% of eligible enrollees buy insurance. Any demand system with independence of irrelevant alternatives gives that the substitution is proportional to shares, so to convert the in-market semi-elasticity to the overall one, we multiply by $\frac{1-.473*.545}{1-.473} = 1.41$, the ratio of the overall share of people not choosing the cheapest plan to the in-market share of people not choosing it. This gives a semi-elasticity of $1.54\% \cdot 1.41 = 2.17\%$

linked subsidies increase the price of the cheapest plan by \$36 (with a std. error of \$6.1)¹³ per month (Equation (4)). This is about a 9% price difference.

This \$36 estimate suggests that the incentive distortion of price-linked subsidies leads to an important effect on markups, but there are a variety of reasons that it is imprecise. First, the equation is a linear approximation; if the demand semi-elasticities are not constant, it will be less accurate for non-marginal changes. Relatedly, converting the \$36 incentive change to a price change implicitly assumes a pass-through rate of 1. If pass-through is less than 1, the price distortion will be smaller. Second, this estimate does not allow for adverse selection into the market, which would: (1) reduce markups and therefore the dollar value of the main price distortion, and (2) add an additional term to the formula for the distortion since a higher subsidy (from a higher price of the cheapest plan) would also bring healthier people into the market, decreasing costs. The net effect on the distortion is ambiguous (see Appendix A.2 where we derive a distortion formula allowing for selection). Lastly, there may be strategic interactions with other plans' prices and the distortion could be smaller if capped by the second-cheapest plan's price. All of these factors are accounted for in the structural model, which we turn to next.

4 Structural Model and Estimation

The reduced form evidence above suggests that price-linked subsidies are likely to increase health insurance markups. To account for some of the market factors not captured by the reduced form analysis and to be able to calculate welfare effects in addition to price effects, we turn to a structural model. In this section, we present the model, estimate it using the CommCare data described in Section 2, and simulate equilibria under different subsidy policies and market structures.

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 the region, t is the time period (year) in which the choice is made, y is income group, and d is the demographic group. For demographic groups we use gender crossed with 5-year age bins because, even though we have detailed information about enrollees, gender and age are the only demographic information available for the uninsured. We suppress the i subscript when the attribute is itself a subscript.

¹³The standard error is calculated via the delta method, assuming zero covariance between the two estimates (since they are estimated based on different time periods).

The utility for consumer i of plan j equals

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

where $P_{j,i}^{cons}$ is the plan's post-subsidy premium for consumer 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 is captured by plan dummies that 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 to capture variation across areas and years, like differing provider networks.

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 relative 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 relative 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 demand for the cheapest plan with respect to the mandate penalty.

Estimation and Identification

We estimate the model by method of simulated moments (MSM). Details of the method and formulas of all moments are in Appendix C.1. We do not use firms' pricing first-order conditions as moments; instead we use micro moments of plan market shares for various groups of consumers, as in Berry et al. (2004). We match plan shares for consumers in each region-year and region-income group combination; this identifies $\xi_j(Z_i)$. Similarly, we match the share of individuals uninsured in each region, year, income and demographic group, which identifies the non-random coefficients in β .

To identify the price-sensitivity parameters, $\alpha(Z_i)$'s, we match the covariance of consumer observables and the price of chosen plans, again as in Berry et al. (2004). A standard concern in identifying price-sensitivity is that prices may be correlated with unobserved plan quality. We address this by using *within-plan* premium variation created by exchange subsidy rules. (This is the same variation as was used for the reduced form estimates in Section 3.2.) Subsidies make all plans free for below-poverty enrollees, while higher-income enrollees pay more for more expensive plans. These rules imply that different income

groups face different relative premiums for the same underlying plans. Moreover, they also create differential premium *changes* over time. For instance, if a plan increases its price between years, its premium increases for higher income groups but is unchanged (at \$0) for below-poverty enrollees. Econometrically, the rich set of plan dummies absorb all price variation except for these differential changes across incomes, just as in a difference-in-differences model. Plan-region-year dummies ($\xi_{j,r,t}$) absorb premium differences arising from plan pricing (which occurs at the region-year level), and plan-region-income dummies ($\xi_{j,r,y}$) absorb persistent taste differences across incomes.

To assess this identification strategy, recall that a plan j 's quality attributes (networks, customer service) are identical across income groups: the same plan is simply being sold at different premiums due to the exchange's subsidy rules. The setup allows income groups (including below-poverty consumers) to have persistent preference differences – e.g., due to differing perceptions of insurer reputation (this is captured by the $\xi_{j,r,y}$ terms). What it assumes is that any *changes* in plan quality apply equally to all incomes within a region (i.e., $\xi_{j,r,t}$ is a single number, not varying across y). This seems particularly reasonable because plan networks, the main aspect of quality, were quite stable over 2007-2011.

Because our identification strategy is analogous to diff-in-diff, another way to assess it is to look at trends in plan shares for different income groups before price changes (i.e., a test for parallel pre-trends). Figures 2 and A3 are strongly consistent with parallel (and basically zero) trends in market shares for the below- vs. above-poverty groups prior to a price change. If quality were varying over time in ways that differentially affected income groups, we would expect these graphs to display differential trends.

To estimate the variance of the random coefficient on uninsurance (σ), we employ a novel approach: we use the change in enrollment in the cheapest plan around the natural experiment of the introduction of a mandate penalty, as described in Section 3.1. Specifically, we match the estimated 22.2% 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 relative utility of uninsurance, the uninsured will tend to be people with lower utility 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.

The model is estimated by MSM, using a random draw of ν_i from the $N(0,1)$ distribution for each consumer in the demand estimation sample (using both from the CommCare and ACS data; see Table 1). The model is just-identified, and our parameter search exactly matches the empirical moments.

Estimated Parameters

The demand coefficients are summarized in Table 2. (A full set of coefficient estimates is in Appendix C.2.) The baseline price coefficients by income (which apply to the omitted demographic group of males age 40-44) range from -0.046 for those just above the poverty line to -0.023 for those making 250%-300% FPL. Females and older consumers are less price sensitive; on average, the premium coefficient increases by +.0028 for women and +.0019 for each 5 years of age. The oldest group is less than half as price-sensitive as the youngest. Panel (b) shows that on average, consumers prefer Fallon to BMC to Network Health to NHP to CeltiCare, with the biggest difference being between CeltiCare and the other plans. Panel (c) shows the utility of uninsurance (β), which is relative to the utility of the average plan (which we normalized to 0). For the average potential consumer, this relative utility of uninsurance is positive (+0.261); this is necessary to account for the low insurance take-up rates in our data despite very low (and often zero) subsidized premiums. However, there is substantial variation around this mean. Females and older people dislike uninsurance more. In addition, unobservable variation (σ) is substantial; it accounts for a standard deviation of 0.92 across individuals, relative to the 0.67 standard deviation driven by observables.

Table 2: Summary of Parameter Estimates in Demand Model

(a) Premium Coefficient			(b) Plan dummies		
	Coef	S.E.		Coef	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	-0.027***	0.0072	Network	0.126***	0.0108
250-300 % Pov	-0.023**	0.0010	BMC	0.169***	0.0312
x 5yr Age Bin	.0019***	0.0006	Fallon	0.209***	0.0563
x Female	.0028**	0.0014			

(c) Relative Utility of Uninsurance					
	Coef	S.E.			
Avg for > 100% Pov	0.261	0.271	Total Standard Deviation		1.14
x 5yr Age Bin	-0.066***	0.019	From Observables		0.67
x Female	-.814***	0.188	From Unobservables (σ)		0.92

Note: The table summarizes the utility coefficients entering our demand model. Panel (a) reports premium coefficients by income for the omitted demographic group of 40-44 year old males and how the coefficients vary 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. Panel (c) shows the average relative utility of uninsurance, how it 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 averages 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.

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%.¹⁴ This implies that CommCare enrollees are quite price sensitive, consistent with their being a low-income population. Converting this semi-elasticity into a rough “insurer-perspective” elasticity by multiplying by the average price (\$386/month) yields an elasticity of -9.3, which is larger than the typical range of -1 to -5 found for 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. With the standard Lerner formula, the insurer-perspective elasticity implies a markup of 11% over marginal cost, but adverse selection would imply a smaller markup over average costs. Perhaps consistent with this, the average plan’s actual markup in the data was $\$399 - \$372 = \$27$ (see Table 1) or 7% of average revenue.

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 estimate raw (not risk-adjusted) costs because, as we detail in Section

¹⁴Our 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 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.

4.3, risk-adjustment works by adjusting plan revenues by multiplying price times a consumer risk score. We assume that costs are generated by a generalized linear model (GLM) with expected costs for consumer i in plan j in year t of

$$E(c_{ijt}) = \exp(\mu X_{it} + \psi_{0,t} + \psi_{j,r,t}). \quad (8)$$

Costs vary with consumer income and demographics ($X_{i,t} = \{y_{i,t}, d_{i,t}\}$), the year ($\psi_{0,t}$), and a region-year specific plan effect, $\psi_{j,r,t}$, which are normalized to average to zero across plans each year. This functional form assumes that, for each region-year, a plan has a constant proportional effect on costs across all consumer types ($X_{i,t}$).

Although our claims data include a rich set of consumer observables, our inclusion of the uninsured population limits us to including in $X_{i,t}$ 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.¹⁵ 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 (8) 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 partially address this issue, we estimate the $\psi_{j,r,t}$ parameters in a separate version of Equation (8) 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). This method would fully eliminate selection if individuals' unobserved risk factors are stable over time or uncorrelated with plan changes. However, it would not address any selection on risk changes – e.g., if individuals who get sicker between enrollment spells systematically select certain plans.¹⁶

Lastly, we adjust observed costs by removing the estimated plan component and estimate the coefficients on individual characteristics from cross-person variation with all enrollees. We use the resulting predicted values of $E(c_{ijt})$ as our estimates of costs for each enrollee-plan possibility.¹⁷

¹⁵The 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).

¹⁶Our sense is that controlling for individual fixed effects addresses most of the selection problem, especially since our panel is over a short period (about four years). Consistent with this, we find that relative to a basic ML estimator, the $\psi_{j,r,t}$ estimates for CeltiCare (the low-quality plan, which we expect to be favorably selected) are shifted upwards towards zero (i.e., less negative) in our estimates.

¹⁷In addition to insurer medical costs captured by this model, insurers incur administrative costs for

Estimated Parameters

Table 4 shows averages of plan cost effects, separately for before and after 2010 when CultiCare entered. The numbers reported are percent effects on costs (i.e., $\exp(\psi_{j,r,t}) - 1$) and 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 Parameter Estimates

	Plan Effects				
	CultiCare	Network Health	BMC	Fallon	NHP
2008-2009		-7.27%	-6.57%	+9.99%	+15.29%
2010-2011	-32.14%	-8.04%	-1.93%	+7.68%	+13.73%

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 reported percentages are averages of $\exp(\psi_{j,r,t}) - 1$, normalized so that the share-weighted average is zero in each year.

The other cost parameters are reported in Appendix C.2. They are as one would expect – costs increase with age and are higher for females at young ages and higher for males at older ages.

4.3 Simulations

We use these demand and cost models and estimates to simulate pricing equilibria under different subsidy policies. 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. As is standard practice, we simulate Nash equilibrium under both policies and compare simulations to simulations, rather than comparing the observed outcome to a simulated outcome.

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. 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

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, and following CommCare rules, insurers are also paid an equal-size (\$30) administrative fee (not subject to risk adjustment). Thus, administrative costs cancel out in the profit function and do not impact our results.

tiers as in the ACA.¹⁸ We do our analysis separately for 2009 and 2011, to illustrate the very different competitive dynamics before and after CeltiCare’s entry in 2010.

Our overall goal is to simulate the effects of price-linked subsidies in a simple insurance exchange setting (like what we model in Section 1) using demand and cost parameters estimated from Massachusetts. For this reason, we do not model a variety of complex and non-standard policies used by the CommCare exchange – including direct price regulation (ceilings and floors), subsidies that varied across plans, and full subsidies for below-poverty enrollees. We believe this gives us the best estimate of the potential effect of price-linking subsidies in this type of market; it will not give us an exact estimate of the counterfactual outcome in this specific market.

Though our demand estimation does not rely on assumptions about firm pricing strategies, the simulations require assumptions about how firms set prices. We assume that they maximize static expected profits; we are ruling out intertemporal pricing considerations by the firms. We also require that firms have the same quality and product positioning under fixed and price-linked subsidies. We do not allow the subsidy structure to affect the set of firms in the market. While these are strong assumptions, they are consistent with other empirical work on insurance markets (e.g. Curto et al. (2014); Tebaldi (2016); Decarolis et al. (2015)) and let us quantify a benchmark estimate of the relevance of price-linked subsidies for pricing.

Risk Adjustment

To mitigate adverse selection, both Massachusetts and the ACA risk-adjust 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. Following Massachusetts’ rules, a plan with price P_{jt} receives $\phi_{it}P_{jt}$ for enrolling consumer i in year t . We assume that ϕ_{it} perfectly captures the individual (non-plan) portion of costs: $\phi_{it} = \exp(\hat{\mu}X_{it} + \psi_{0t})/\bar{c}_t$, where \bar{c}_t is the average of $\exp(\hat{\mu}X_{it} + \psi_{0t})$ across individuals who buy insurance in the data. Thus, risk adjustment fully captures individuals’ expected costs in a proportional way: net revenue on consumer i ($= \phi_{it}P_{jt} - E(c_{ijt})$) simplifies to $\phi_{it}(P_{jt} - \bar{c}_t \exp(\psi_{j,r,t}))$.

Equilibrium

We look for a static, full information, Nash equilibrium, where each insurer sets its price to maximize profits:

$$\pi_{jt} = \sum_i (\phi_{it}P_{jt} - E[c_{ijt}]) \cdot Q_{ijt}(P^{Cons}(P_t)) \quad (9)$$

¹⁸Although ACA consumers can choose across tiers, in practice, the vast majority of subsidized consumers choose a silver plan because “cost-sharing reduction” (CSR) subsidies are linked to silver plans.

where $P^{Cons}(\cdot)$ is the subsidy function mapping prices into consumer premiums. This equilibrium is defined by the first-order conditions (FOCs) $\partial\pi_{jt}/\partial P_{jt} = 0$ 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 counterfactual markets with only two insurers.

Note that the insurers' FOCs depend on $\frac{dQ_{jt}}{dP_{jt}}$, the total derivative that incorporates an effect of changing P_{jt} on the subsidy if plan j is pivotal (cheapest) under price-linked subsidies (as shown in Equation (1)). 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 $\frac{dQ_{jt}}{dP_{jt}} = \frac{\partial Q_{jt}}{\partial P_{jt}} + \frac{\partial Q_{jt}}{\partial M}$) while above it, they are not (so $\frac{dQ_{jt}}{dP_{jt}} = \frac{\partial Q_{jt}}{\partial P_{jt}}$). This means that if the incentive distortion causes the cheapest plan to raise its price up to the level of the second cheapest plan, there will likely be multiple equilibria – a range of prices at which the two plans tie for cheapest 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.

Prices

Table 5 reports the prices, subsidy and the share of eligible individuals that are insured for each equilibrium. 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, leading to a range of possible equilibria where these two plans tie for cheapest. 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 preferred). 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.

In 2011, CeltiCare's cost are enough lower than the other plans that it is the unique cheapest plan. CeltiCare's price is \$24 (6%) higher under price-linked subsidies; this represents a 43% increase in its markup over average costs. The other insurers actually lower their prices very slightly in moving from fixed to price-linked subsidies. Nevertheless, the change in the average price is only slightly smaller than the change in the cheapest price because there is a substantial change in market share: since CeltiCare is no longer as cheap, its in-market share drops from about 61% to about 41%. The ~20% of consumers who switch away from CeltiCare choose more expensive plans, increasing the average price paid. Because prices are higher while (income-specific) subsidies are held constant, fewer people buy insurance under price-linked subsidies.

The estimated \$24 price change is smaller than the \$36 we calculated in Section 3.3 using the first-order approximation. We can try to separate out how much of the difference is due to (1) different elasticities in 2008 (when the mandate penalty introduction occurred) versus for CultiCare in 2011 (the simulation year), (2) adverse selection and risk adjustment, (3) pass-through of cost shocks into prices not equaling 1 (which is implicitly assumed by the first-order approximation), and (4) strategic responses from the other firms. We analyze the role of each of these factors in Appendix C.3 and Appendix Table 13. We find that about half of the gap between estimates is explained by different elasticities in 2008 vs. 2011, a fourth is explained by adverse selection and risk adjustment, and nearly all the rest is explained by cost pass-through being less than 1.

Table 5: Equilibrium under Price-Linked and Fixed Subsidies

	2009					2011		
	Price-Linked		Fixed	Difference		Price-Linked	Fixed	Diff
	Min	Max		Min	Max			
Minimum Price	\$415.3	\$437.3	\$411.7	\$3.5	\$25.6	\$403.8	\$379.9	\$23.8
Profit Margin	14.6%	19.1%	14.3%	0.3%	4.8%	20.0%	15.0%	5.0%
Avg Price	\$418.6	\$441.2	\$416.6	\$2.0	\$24.6	\$427.3	\$405.4	\$21.9
Avg Subsidy	\$354.4	\$375.8	\$354.2	\$0.2	\$21.7	\$345.7	\$346.7	-\$1.0
Share Insured	57.6%	58.5%	58.3%	-0.7%	0.2%	40.0%	45.9%	-5.9%

Note: This is a comparison of the market equilibria under price-linked and fixed subsidies (with the “Diff” columns referring to price-linked minus fixed). In 2009, there are only four plans in the market (CultiCare has not entered); there are a range of equilibria under price-linked subsidies; this reports statistics for the equilibrium with the minimum and maximum cheapest plan price. The profit margin ($=100 \cdot (\text{revenue} - \text{cost}) / \text{revenue}$) is reported for the cheapest plan. For 2009, the margin is the share-weighted average of BMC and Network’s margins. The average price is across all plans, weighted by plan shares. The average subsidy is across all 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.

Our simulations do not include medical loss ratio (MLR) rules, which could mitigate the price distortion. The ACA requires that insurers spend at least 80% of revenue on medical care and related services. If binding, MLR rules could prevent price-linked subsidies from raising markups. However, the MLRs implied by our simulations – calculated conservatively as the ratio of medical costs to revenue – are typically above 80%. (The lone exception is BMC in the max-price equilibrium in 2009, whose MLR is 77%; see table in Appendix C.4).

Another factor that could mitigate the pricing distortion is if market regulators are directly involved in setting prices via negotiations with insurers, something that has occurred in a handful of state-based ACA exchanges (e.g., California; see Tebaldi (2016)). We do not know how much this could limit the distortion, and modeling this interaction is beyond the

scope of our paper. But we note that it seems that much of the potential benefits of price-linking relative to fixed subsidies (which arise from insurers’ informational advantage) could also be diminished if the regulator is involved in price setting.

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 (which cost 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 2009.¹⁹

Welfare

Our main welfare analysis focuses on “public surplus,” introduced in Section 1.2 – consumer surplus, plus the avoided externality of uninsurance, minus government costs. We also consider adding insurer profits to this. For each simulation, we use the simulated equilibrium prices along with the estimated demand parameters to calculate the fixed component of utility for each consumer for each plan, $\hat{u}_{ij} \equiv u_{ij} - \epsilon_{ij}$. With logit demand, 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 expenditure on each individual is $S(1 - \widehat{Pr}_{i0}) - M \cdot \widehat{Pr}_{i0}$.

Uninsured people can generate a negative externality through the uncompensated care that they receive (whose cost is born by the providers or government); if society or the regulator has a paternalistic desire for people to have health insurance, then that is another negative effect of people being uninsured, independent of the medical costs they incur. We assume that the regulator sets the affordable amounts for price-linked subsidies optimally given the total externality. This allows us to back out the level of the paternalistic externality.²⁰ This calibration is conceptually important. Without the paternalistic externality, subsidies are “too high” from a welfare-maximizing perspective so there are “too many” people buying insurance. This makes fixed subsidies seem bad because they lead to more people

¹⁹These 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.

²⁰For these calculations, we assume that uncompensated care costs are equal to 80% of expected cost of care with insurance and that the paternalistic component of the externality does not vary with consumers’ expected costs. See Section 5.1 for further discussion of uncompensated care costs and Appendix D for robustness checks with the externality equal to 60% and 100% of expected cost.

buying insurance, whereas we want to compare fixed and price-linked subsidies assuming they are set at the right level, given the externality.

Table 6 shows the different components of welfare under fixed and price-linked subsidies. 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 CeltiCare’s price applied to about $\sim 20\%$ of the eligible population who purchase CeltiCare.²¹

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

	\$ Per Eligible Consumer per Month							
	2009					2011		
	Price-Linked		Fixed	Difference		Price-Linked	Fixed	Diff
Min	Max	Min		Max				
Consumer Surplus	-0.3	0.9	0.7	-1.0	0.1	-14.4	-9.3	-5.1
+ Saved Externality	246.2	249.7	249.0	-2.9	0.7	165.5	188.2	-22.7
- Gov Costs	-177.2	-193.2	-179.9	2.7	-13.3	-103.2	-125.3	22.1
= Public Surplus (PS)	68.6	57.3	69.8	-1.2	-12.5	47.9	53.6	-5.7
Insurer Profits	33.9	45.3	33.4	0.5	11.9	22.3	22.7	-0.4
PS + Profits	102.6	102.6	103.3	-0.7	-0.7	70.2	76.3	-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. (This number is negative because consumers have low or negative value for insurance.) Government costs are mandate revenue minus subsidy expenditures. Saved externality is the sum across consumers of the probability that they buy insurance times their externality. Public surplus adds consumer surplus and the saved externality, and subtracts government costs. The last rows add firm profits to public surplus.

Consistent with our theoretical result, public surplus is higher under fixed subsidies. When the level of subsidy is equal under the two policies – as it is for 2011 and the minimum price-linked equilibrium in 2009 – the difference in public surplus is approximately equal to the difference in consumer surplus. This is because the differences in government costs and the externality approximately cancel out due to an envelope theorem argument – since the regulator set fixed subsidies (affordable amounts) optimally, the value of covering an

²¹We 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 the program has little (or negative) effect on consumer surplus.

additional person equals the additional public costs. For 2009’s maximum price equilibrium, the subsidy amount is also larger under price-linked subsidies, so there is an additional increase in government spending, lowering public surplus.

The last two rows of Table 6 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 profits. The effect of price-linked subsidies on profits is ambiguous because the higher prices decrease the quantity insured, particularly among healthier, low-cost consumers, so profits may decrease. Even when profits are higher, the sum of public surplus and profits is still lower under price-linked subsidies.

Less Competitive Markets

Many of the ACA exchanges have only one or two insurers. To understand the implications of price-linked subsidies in less competitive markets, we again simulate equilibria using our parameters for 2011, but with the market limited to two available plans. We first note that less competition implies higher prices even under fixed subsidies: for instance, CeliCare’s price averages \$26 higher than in the baseline fixed subsidy simulation with five plans. This is comparable to the price difference between fixed and price-linked subsidies that we observed in Table 5. Thus, switching from fixed to price-linked subsidies has about the same effect on the price of the cheapest plan as removing all but one of its competitors.

We next compare fixed and price-linked subsidies in a market with two insurers. Table 7 shows the the price distortion – the increase in the cheapest price under price-linked relative to fixed subsidies – for each pair of the four main insurers (excluding Fallon, which is a smaller regional plan). The first insurer in the pair is listed in the row and the second in the column. The plans are listed in order of increasing costs: CeliCare 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 price distortion amounts are larger than what we estimate for the full market with five insurers (\$23.80). 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 CeliCare and Network Health, two low-cost plans, are competing. But it is more than twice as big (\$50.1) when CeliCare 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. This also reflected in welfare (not shown in the table); depending on whether the second cheapest firm is low or high cost, public surplus per eligible consumer is between \$8 and \$39 higher under fixed subsidies than under price-linked subsidies.

With only two insurers, medical loss ratio rules become more binding. CeltiCare’s cost-to-revenue ratio under price-linked subsidies is below 80% against any competitor and below 70% against NHP. (CeltiCare also has an MLR below 80% even for fixed subsidies when competing with NHP). Network Health and BMC also have MLRs below 80% when competing against NHP.

Therefore, for these two-insurer markets, the distortion from price-linked subsidies may be mitigated by MLR rules. In addition, other regulatory pressures (e.g., “rate review” by state insurance departments) may limit extremely high plan markups. These regulatory pressures are clearly necessary to explain prices in ACA counties with a single insurer, for which our model (taken literally) predicts infinite markups. For these extreme cases, our model predicts what insurers have an incentive to do – and may find ways of doing if they can elude regulatory barriers.

Table 7: Price Distortion with Two Insurers (2011 Parameters)

		(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 insurers 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 (in parentheses) under fixed subsidies, for markets with two insurers. These estimates should be compared to the increases in prices with five insurers from Table 5. Each set of figures reported corresponds to simulations where the market includes just the two insurers listed in the row and column headers. 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.

We focus 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 rules 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.

5 Discussion

5.1 Why do we see price-linked subsidies?

Given their existence in real-world markets, are there benefits to price-linked subsidies that justify their use despite the pricing distortion we highlight? There are two categories of constraints on regulators that we omitted from the model that may make price-linked subsidies more attractive: informational constraints (uncertainty) and political economy constraints.

Uncertainty

Our comparison of price-linked subsidies and fixed subsidies assumed that regulators had full information about cost and prices in the market so they could ex-ante set an optimal fixed subsidy. There are two reasons that price-linked subsidies may be desirable in the presence of uncertainty. First, if the regulator is uncertain about the costs of health care, price-linked subsidies allow them to more closely match the subsidy for insurance to the externality of uninsurance. If there is a market-level shock to cost growth that insurers can observe (e.g., because of their up-to-date data and regular interactions with providers) and regulators cannot, then that shock will be reflected in prices; prices will contain useful information about the externality of uninsurance and therefore about the optimal subsidy, so price-linking subsidies may get them closer to the optimal level. Second, if the regulator is uncertain about insurance prices – either because of cost uncertainty or because of uncertainty about the equilibrium given costs – then price-linked subsidies allow the regulator to avoid the “affordability risk” that consumers incur from the variance in the post-subsidy prices that occurs with fixed subsidies.

Uncertainty about market level costs can be incorporated into our simulations in a fairly straight forward way; we take the observed cost levels (estimated in our model) as the ‘expected’ cost level and recalculate what the equilibrium would be if costs were between 20% lower and 20% higher than expected, holding fixed the affordable amounts and subsidy levels set for the ‘expected’ costs. To calculate the externality of uninsurance for each equilibrium, we start by assuming that the paternalistic component of the uninsurance externality does not depend on health care cost and that the expected social cost of the care received by the uninsured is proportional to the expected cost of their care with the average plan. We expect uncompensated care costs to be lower than insurance costs, so a \$1 increase in insurance cost would result in a $\lambda \leq \$1$ increase in the cost of uncompensated care. Since price-linking is more attractive when the externality moves more with prices, we consider the $\lambda = 1$ as a best-case scenario for price-linked subsidies. In Appendix D we discuss and analyze other reasonable benchmarks, give more details on the uncertainty framework and simulations,

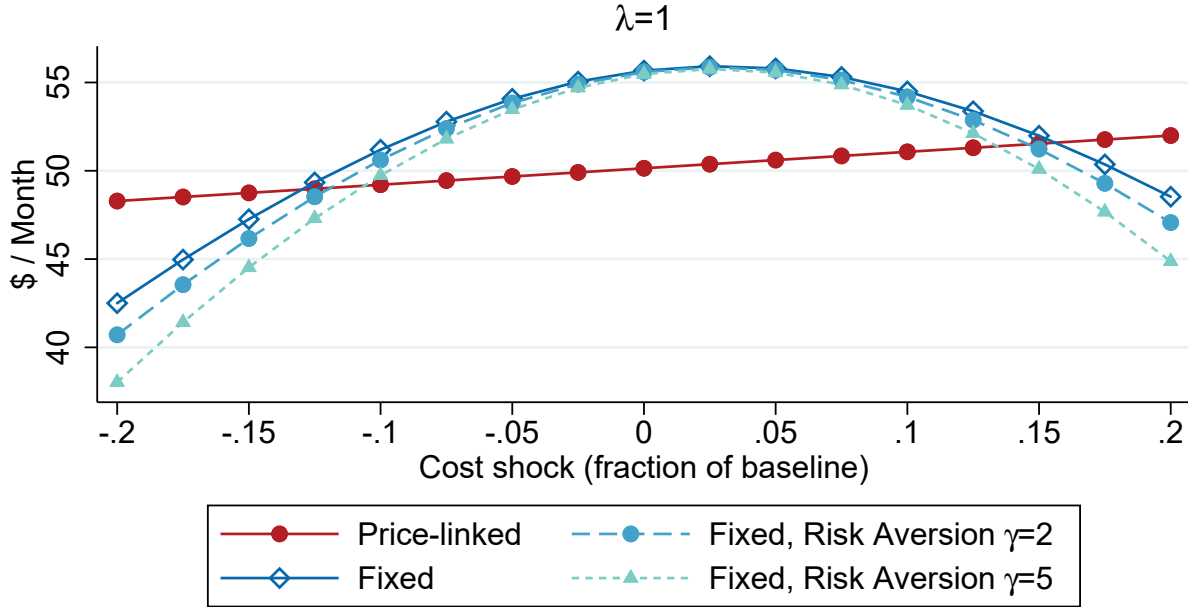


Figure 3: Public surplus under cost shocks in 2011

Note: These graphs show public surplus (the regulator’s objective function, in dollars per month per eligible member) under price-linked and fixed subsidies for cost shocks of -20% to +20% of baseline. The dashed and dotted lines subtract a cost of pricing risk to consumers (using the method in Equation (6)) with coefficients of relative risk aversion of 2 and 5, respectively.

and show how prices, subsidies, and the insurance coverage rate vary with the cost shocks.

Figure 3 shows public surplus (the regulator’s objective function) as costs diverge from the regulator’s expectations, for cost shocks from -20% to +20%. As expected, based on the lower prices, fixed subsidies result in higher surplus at a cost shock of 0. For non-zero cost shocks, the gap between fixed and price-linked subsidies narrows because the optimal subsidy gets farther from the fixed subsidy level as costs – and therefore the externality of uninsurance – diverge from expectations. For the baseline public surplus measure (solid curves), price-linked subsidies only do better for cost shocks greater than 15% or less than -12.5%. Using Equation (6), we then adjust the public surplus measure to include a cost of price risk to consumers, shown for a coefficient of relative risk aversion of $\gamma = 2$ (dashed lines) and a more extreme case of $\gamma = 5$ (dotted lines), 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). Even with fairly high risk aversion, the cost shock must be over 12.5% or below -10% for price-linked subsidies to have higher public surplus than fixed subsidies.

The basic results also hold if we use public surplus plus profits as the welfare metric. These results are also shown in Appendix D. The range of shocks under which fixed subsidies do better is somewhat smaller, -7.5% to +10%, but still substantial.

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 a ballpark magnitude under different assumptions about the subsidy-setting process.

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. If 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), then 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) greater than 5% (in absolute value) occur less than 1% of the time. 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 greater than 12.5% – the minimum shock for which price-linked subsidies do better – still occur less than 1% of the time.

We conclude that fixed subsidies do better across the range of cost shocks that are “reasonable” in a mature market similar to the one we study. However, uncertainty is likely to be much larger in a new market, like the ACA in 2014, where past data is not available or in markets where regulations are changing or enrollment has not reached equilibrium levels. Moreover, in more competitive markets or markets where there is less substitution to the outside option, the distortion from price-linked subsidies would be smaller, so the amount of cost uncertainty necessary to overcome the higher prices would also be smaller.

Political economy constraints

Another motivation for price-linked subsidies is that they may be less susceptible to lobbying or regulatory capture. Affordable amounts require less year-to-year adjustment than fixed subsidies, which would likely need to grow with health care costs. Moreover, they can be set at the same level across all regions, whereas fixed subsidies may need to vary with local health care costs. The joint federal-state nature of the program (subsidies are federally funded but exchanges are state-regulated) also presents complications. With fixed subsidies, state regulators might unilaterally set subsidies above the optimal level to bring additional federal subsidy dollars into their states. With price-linked subsidies the regulator would have to try to coordinate higher pricing (collusion) across insurers. This is similar to the “fiscal shenanigans” concerns that have been documented for federal matching funds in Medicaid (Baicker and Staiger, 2005).

Another potential rationale for insurance subsidies might be redistribution. We choose not to model distributional objectives (instead adopting a surplus standard), since we assume these are addressed elsewhere in the tax/transfer system, but it is possible that political constraints make it easier to do redistribution through health insurance than via tax credits or direct payments. Political constraints could also create extreme affordability concerns that go beyond standard risk aversion. If the political ramifications of insurance not being affordable in one county are very large, than even a tiny bit of uncertainty could make price-linked subsidies necessary.

Political constraints could also limit the government’s ability to set fixed subsidies optimally under uncertainty. If local regulators were too subject to capture, Congress might feel the need to control the level of the fixed subsidies, setting an initial level of fixed subsidies and indexing 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. In the medium-to-long term, just indexing fixed subsidies will frequently lead to worse outcomes than price-linked subsidies.

5.2 Implications for Other Markets

The tradeoffs with price-linked subsidies apply more broadly than the ACA. They apply in any market where (1) firms have market power – i.e., a small price increase does not cause a firm’s demand to fall to zero – and (2) there is the possibility of substitution to an unsubsidized outside option. These conditions apply in a variety of programs, including Medicare Advantage, Medicare Part D, and employer-sponsored insurance.

In Medicare Advantage, the distinction between price-linked and fixed subsidies is relevant for comparing “competitive bidding” and “premium support” reform proposals.²² Both reforms propose explicitly linking subsidies to insurer prices. Under competitive bidding proposals, the price-linked subsidy applies only to Medicare Advantage 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 (see Section 1.3) 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’

²²Medicare Advantage’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).²³

Employers typically pick a small menu of insurance 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 type of pricing distortion. Since tax rules limit employers’ ability to subsidize employees’ outside options, employers who want to keep prices down should consider making their subsidies not depend, even implicitly, on insurers’ prices.

In theory, the logic of price-linked subsidies applies to market-based public programs other than health insurance. The markets tend to be less centralized, but housing subsidies, school vouchers, and Pell Grants have many of the same properties of price-linked 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.²⁴ Again, price-linked subsidies would only be beneficial if the optimal subsidy increased with prices.

6 Conclusion

This paper considers the effects on pricing incentives generated by price-linked subsidies in health insurance exchanges, an important topic for economists analyzing these markets and policymakers designing and regulating them. We highlight the incentive 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. This would translate to \$46 million across all CommCare enrollees in 2011 or \$3 billion for ACA subsidized enrollees in 2016. While we do not view these numbers as a precise estimate of either the historical distortion in Massachusetts or the distortion in the ACA exchanges, we think that they indicate that the pricing incentives we identify in theory should be of practical concern.

²³Decarolis (2015) discusses additional pricing effects arising from the design of low-income subsidies.

²⁴Linking 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 also show that price-linked subsidies may be a response to political economy constraints or uncertainty about health care costs. The right balance between the insurers' pricing incentives on the one hand and affordability and consumer incentive concerns on the other will vary from market to market. In general, price-linked subsidies are more likely to be beneficial on net when regulators face more uncertainty, since the information about costs that prices contain will be more valuable. They are also more likely to be beneficial in more competitive markets, where price-linking does not distort prices as much. We hope our analysis contributes to a better understanding of the tradeoffs involved. In addition to further analysis of the ACA, future research could 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.

References

- 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? *American Economic Review: Papers & Proceedings*, 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. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488.
- 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.

- Cutler, D. M. and Reber, S. J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466.
- 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.
- Finkelstein, A., Hendren, N., and Shepard, M. (2017). Subsidizing health insurance for low-income adults: Evidence from Massachusetts.
- Garthwaite, C., Gross, T., and Notowidigdo, M. J. (2017). Hospitals as insurers of last resort. *American Economic Journal: Applied Economics*. Forthcoming.
- 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 Scott Morton, F. (2017). The impact of consumer inattention on insurer pricing in the Medicare Part D program. *The RAND Journal of Economics*, 48(4):877–905.
- Ho, K. and Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, 85(2):379–417.
- Kaiser Family Foundation (2018). Insurer participation in the 2018 individual marketplace.
- 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.
- Ryan, P. (2017). A better way to fix health care.
- 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

A.1 Multi-Plan Insurers (ACA Case)

In this section, we show how the basic theory of price-linked subsidies' effect on markups (see Section 1) carries over to the multi-plan setting of the ACA. In the ACA insurers may offer plans in each of multiple tiers – bronze, silver, gold, and platinum.²⁵ Subsidies are set equal to the price of the second-cheapest silver plan minus a pre-specified “affordable amount.”²⁶ In general, 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.²⁷

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:

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

The second term reflects a standard effect for multi-product firms: when the insurer raises the price of one plan, it captures revenue from consumers who switch to its other plans.

²⁵Platinum plans cover 90% of medical costs (comparable to a generous employer plan); 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.

²⁶The 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.

²⁷If an insurer does not offer a bronze plan *and* consumers on the margin of uninsurance mostly pick bronze plans, that would mitigate the distortion. Participating insurers are required to offer gold and silver plans but may choose not to offer bronze; empirically, however, nearly all participating insurers do offer at least one bronze plan.

However, with price-linked subsidies, the markup for the subsidy-pivotal plan is:

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

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.

A.2 Adverse Selection

Our basic theory in Section 1 abstracted from adverse selection (beyond what the exchange's risk adjustment accounted for) for expositional simplicity. In this section, we show that the basic logic of the distortion of price-linked subsidies generalizes to a model with risk selection. With risk selection, a plan's average costs, $\bar{c}_j(P^{cons}, M)$, depend on the set of consumers who select a plan, which is affected by premiums and the mandate penalty. 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:

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

It's first-order condition is

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

Just like the total derivative of demand, the total derivative of cost, $\frac{dc_j^{Net}}{dP_j}$, differs with the subsidy rule. With fixed subsidies

$$\frac{dc_j^{Net}}{dP_j} = \frac{\partial c_j^{Net}}{\partial P_j^{cons}}.$$

With price-linked subsidies

$$\frac{dc_j^{Net}}{dP_j} = \frac{\partial c_j^{Net}}{\partial P_j^{cons}} + \frac{\partial c_j^{Net}}{\partial M}.$$

The respective markups are

$$Mkup_j^F = \frac{1}{\eta_j} \left(1 - \frac{\partial c^{Net}}{\partial P_j^{cons}} \right)$$

$$Mkup_{\underline{j}}^{PLink} = \frac{1}{\eta_{\underline{j}} - \eta_{\underline{j},M}} \left(1 - \frac{\partial c^{Net}}{\partial P_j^{cons}} - \frac{\partial c^{Net}}{\partial M} \right).$$

where $\eta_j \equiv -\frac{1}{Q_j} \frac{\partial Q_j}{\partial P_j^{cons}}$ is the own-price semi-elasticity of demand and $\eta_{j,M} \equiv \frac{1}{Q_j} \frac{\partial Q_j}{\partial M}$ is the semi-elasticity of demand for \underline{j} with respect to the mandate penalty. The difference is

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

Adverse selection implies lower markups under either fixed or price-linked subsidies, since $\left(1 - \frac{\partial c^{Net}}{\partial P_j} \right) < 1$. This decrease in markups has the effect of decreasing the difference in markups between the subsidy policies. However, with adverse selection, a higher subsidy brings healthier people into the market, decreasing costs ($\frac{\partial c^{Net}}{\partial M} < 0$), so the second term is positive, raising the difference in markups. The change as a fraction of the markup under fixed subsidies

$$\frac{Mkup_{\underline{j}}^{PLink} - Mkup_{\underline{j}}^F}{Mkup_{\underline{j}}^F} = \frac{\eta_{\underline{j},M}}{(\eta_{\underline{j}} - \eta_{\underline{j},M})} - \frac{\frac{\partial c^{Net}}{\partial M}}{1 - \frac{\partial c^{Net}}{\partial P_j^{cons}}} \frac{\eta_{\underline{j}}}{\eta_{\underline{j}} - \eta_{\underline{j},M}}$$

is larger with adverse selection because of the second term (compared to Equation (4) in the main text).

A.3 Optimality of Fixed Subsidies

To show that fixed subsidies are optimal, we model there being a total subsidy, S that is a mixture between a fixed subsidy \hat{S} and a subsidy that depends on the cheapest plan's price with weight α :

$$S = \hat{S} + \alpha P_{\underline{j}},$$

where firm \underline{j} is the subsidy-pivotal plan. We show that under the conditions Section 1.2, as long as the level of \hat{S} is set optimally, the derivative of welfare with respect to α is negative, implying that it is optimal to not link subsidies to price. (This actually suggests that the optimal subsidy would be negatively linked to prices, but that is not a policy we have seen proposed.)

Consumers only care about the total subsidy, which is also what determines government expenditures. Therefore, α and \hat{S} affect welfare only via their effects on the total subsidy

and their effects on prices.²⁸ So,

$$\begin{aligned}\frac{dW}{d\hat{S}} &= \frac{\partial W}{\partial S} + \sum_j \frac{\partial W}{\partial P_j} \frac{\partial P_j}{\partial \hat{S}} \\ \frac{dW}{d\alpha} &= \frac{\partial W}{\partial S} \cdot P_{\underline{j}} + \sum_j \frac{\partial W}{\partial P_j} \frac{\partial P_j}{\partial \alpha}\end{aligned}$$

If \hat{S} is set optimally, then the first line equals zero. Substituting the resulting expression for $\frac{\partial W}{\partial S}$ into the second line, it becomes

$$\frac{dW}{d\alpha} = \sum_j \frac{\partial W}{\partial P_j} \left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial \hat{S}} \cdot P_{\underline{j}} \right).$$

Condition (iii) of the proposition is that $\frac{\partial W}{\partial P_j}$ is negative for all j , so we just need the price changes $\left(\frac{\partial P_j}{\partial \alpha} - \frac{\partial P_j}{\partial \hat{S}} \cdot P_{\underline{j}} \right)$ to be positive.

For price changes, we turn to the firms' first-order conditions, which must continue to hold as α or \hat{S} changes. If π_i is each firm's profit function, the effect of \hat{S} and α on prices follow

$$\begin{aligned}\left[\frac{\partial^2 \pi_i}{\partial P_i \partial P_j} \right] \cdot \frac{\partial P}{\partial \hat{S}} + \frac{\partial^2 \pi_i}{\partial P_i \partial \hat{S}} &= 0 \\ \left[\frac{\partial^2 \pi_i}{\partial P_i \partial P_j} \right] \cdot \frac{\partial P}{\partial \alpha} + \frac{\partial^2 \pi_i}{\partial P_i \partial \alpha} &= 0\end{aligned}$$

so

$$\frac{\partial P}{\partial \alpha} - P_{\underline{j}} \cdot \frac{\partial P}{\partial \hat{S}} = - \left[\frac{\partial^2 \pi_i}{\partial P_i \partial P_j} \right]^{-1} \left(\frac{\partial^2 \pi_i}{\partial P_i \partial \alpha} - P_{\underline{j}} \frac{\partial^2 \pi_i}{\partial P_i \partial \hat{S}} \right). \quad (\text{A.2})$$

For all firms except \underline{j} , both \hat{S} and α only enter the profit function via S so

$$\frac{\partial^2 \pi_i}{\partial P_i \partial \alpha} = P_{\underline{j}} \frac{\partial^2 \pi_i}{\partial P_j \partial S} = P_{\underline{j}} \frac{\partial^2 \pi_i}{\partial P_j \partial \hat{S}},$$

making the term in parentheses in Equation (A.2) zero. For firm \underline{j} ,

$$\begin{aligned}\frac{\partial \pi_{\underline{j}}}{\partial P_{\underline{j}}} &= \left(1 - \frac{\partial c_{\underline{j}}^{Net}}{\partial P_{\underline{j}}} \right) + \frac{\partial Q_{\underline{j}}}{\partial P_{\underline{j}}} (P_{\underline{j}} - c_{\underline{j}}^{Net}) + \alpha \left(\frac{\partial Q_{\underline{j}}}{\partial \hat{S}} (P_{\underline{j}} - c_{\underline{j}}^{Net}) - Q_{\underline{j}} \frac{\partial c_{\underline{j}}^{Net}}{\partial \hat{S}} \right) \\ \frac{\partial^2 \pi_{\underline{j}}}{\partial P_{\underline{j}} \partial \alpha} - P_{\underline{j}} \frac{\partial^2 \pi_{\underline{j}}}{\partial P_{\underline{j}} \partial \hat{S}} &= \left(\frac{\partial Q_{\underline{j}}}{\partial \hat{S}} (P_{\underline{j}} - c_{\underline{j}}^{Net}) - Q_{\underline{j}} \frac{\partial c_{\underline{j}}^{Net}}{\partial \hat{S}} \right).\end{aligned}$$

²⁸Even if profits enter welfare, ex-post these only depend on the total subsidy, though ex-ante, α and \hat{S} obviously create different incentives.

This is equal to $= P_{\underline{j}} \frac{\partial \pi_{\underline{j}}}{\partial \hat{S}}$, which Condition (i) requires to be positive.

Condition (ii) ensures that the effect on other firm's prices when the pivotal plan has a marginal cost increase,

$$\frac{\partial P_i}{\partial c_j^{Net}} = - \left[\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} \right]_{j,\underline{j}}^{-1},$$

is positive. Since all the terms in the \underline{j}^{th} column of $\left[\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} \right]^{-1}$ are negative, and $\frac{\partial P_i}{\partial \alpha} - P_{\underline{j}} \cdot \frac{\partial P_i}{\partial \hat{S}}$ is positive for all j , $\frac{dW}{d\alpha}$ is negative.

Uncertainty

The above logic breaks down under uncertainty. With uncertainty, the first-order condition for the fixed component of the subsidy only holds in expectation:

$$\frac{dW}{d\hat{S}} = E \left[\frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial \hat{S}} \frac{\partial W}{\partial P_j} \right] = 0.$$

This means that

$$\begin{aligned} E[P_{\underline{j}}] \cdot E \left[\frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial \hat{S}} \frac{\partial W}{\partial P_j} \right] &= 0 \\ cov \left(P_{\underline{j}}, \frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial \hat{S}} \frac{\partial W}{\partial P_j} \right) - E \left[P_{\underline{j}} \left(\frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial \hat{S}} \frac{\partial W}{\partial P_j} \right) \right] &= 0. \end{aligned}$$

So the effect of α is

$$\begin{aligned} \frac{dW}{d\alpha} &= E \left[\frac{\partial W}{\partial S} \cdot P_{\underline{j}} + \sum_j \frac{\partial P_j}{\partial \alpha} \frac{\partial W}{\partial P_j} \right] \\ &= cov \left(P_{\underline{j}}, \frac{\partial W}{\partial S} + \sum_j \frac{\partial P_j}{\partial \hat{S}} \frac{\partial W}{\partial P_j} \right) + E \left[\sum_j \frac{\partial W}{\partial P_j} \left(\frac{\partial P_j}{\partial \alpha} - P_{\underline{j}} \frac{\partial P_j}{\partial \hat{S}} \right) \right]. \end{aligned}$$

In order for price-linking to have a positive effect on welfare, the covariance between the subsidy-pivotal price and the marginal effect of the subsidy on welfare must be sufficiently positive to outweigh the negative effect on welfare of the higher prices.

A.4 Risk Aversion

Let consumers have concave utility over non-health insurance consumption, $u(\cdot)$, with a (local) coefficient of relative risk aversion $\gamma = -\frac{u''(y)y}{u'(y)}$. If Y is income and P is the premium

for insurance, then the change in utility from a premium difference of ΔP is

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

The last 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 first term is the effect on risk neutral consumers which is included in our baseline calculation. The second term is the additional cost from risk aversion.

B Data and Reduced Form Estimation Details

B.1 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. For each CommCare year, we pool observations for the two ACS calendar years in which it occurred. For instance, for CommCare year 2009 (which ran from July 2008 to June 2009), we pool observations from ACS 2008 and 2009 (and divide weights by two) so that the uninsured sample reflects an average of the two relevant calendar years. We restrict to Massachusetts residents 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 all ACS

uninsured observations but only new enrollees for the CommCare data – so that the uninsured rate calculated in this sample matches the population uninsured rate.

B.2 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 (a proxy for the steady-state size of each plan). 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 a slightly expanded version of the difference-in-difference specification shown in the text in Section 3.1:

$$NewEnroll_{y,t} = \sum_{m \in DM} ((\alpha_0^m + \alpha_1^m Treat_y) \cdot 1_m \cdot MandIntro_t + \zeta_{ym}) + \xi_y + \delta_y X_t + \varepsilon_{y,t},$$

where DM is the calendar months December through March, $Treat_y$ is a dummy for the treatment group, 1_m is an indicator for month m , $MandIntro_t$ is a dummy for the mandate penalty introduction period (Dec. 2007 to March 2008), and X_t is a vector of time polynomials and CommCare-year dummies.²⁹ The difference-in-difference coefficients of interest are the α_1^m 's. In our main specifications, we use two income groups (y): 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. People 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.

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

²⁹The CommCare-year starts in July, so these dummies will not conflict with the treatment months of December to March. All results are robust to excluding these year dummies and time polynomials.

Table 8: Introduction of the Mandate Penalty.

	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)	0.102** (0.007)	0.101** (0.007)	0.112** (0.007)	0.091** (0.008)
x Jan2008	0.073** (0.004)	0.067** (0.006)	0.067** (0.007)	0.059** (0.007)	0.084** (0.007)	0.086** (0.007)
x Feb2008	0.043** (0.005)	0.033** (0.006)	0.032** (0.006)	0.026** (0.007)	0.042** (0.007)	0.049** (0.006)
x Mar2008	0.025** (0.018)	0.027** (0.023)	0.021** (0.024)	0.020* (0.026)	0.021** (0.023)	0.022** (0.023)
Total	0.253*** (0.018)	0.237*** (0.023)	0.222*** (0.024)	0.205*** (0.026)	0.259*** (0.023)	0.248*** (0.023)
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	0.967	0.921	0.923	0.925	0.920	0.919

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 < 100% 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.

all years, forming the triple difference specification (shown in the equation above) that nets 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 – consistent with the dramatic spike shown in Figure 1. In our preferred triple difference estimates in column (3), the mandate penalty increases enrollment in the cheapest plan by 22.2% 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 1.17% for the 150-200% of poverty group, 0.74% for the 200-250% poverty group, and 0.47% for the 250-300% of poverty group, with a (group-size) weighted average of 0.95%.

We interpret 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. However 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 (though the total *annual* penalty for 2008 was actually larger than in 2007 for all but the 150-200% poverty group). 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 A1a and A1b 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.)³⁰ 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 (but without financial

³⁰The 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.

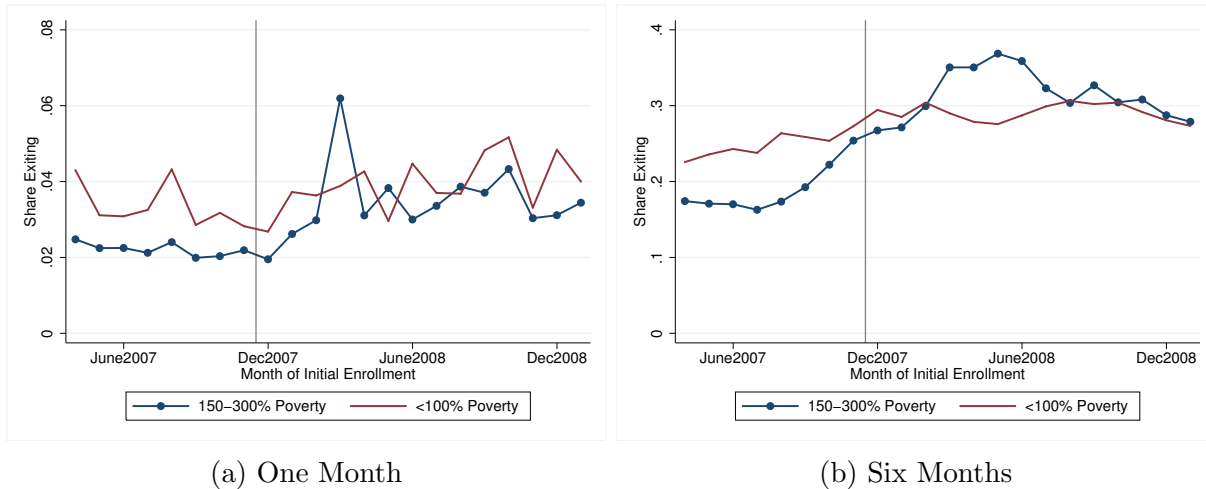


Figure A1: Share of New Enrollees Exiting Within the Specified Number of Months.

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 in 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.

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 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 government 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 2007.³¹ 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

³¹The 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.




Figure A2: Decrease in the Affordable Amount Experiment

Note: Analogously to Figure 1, this graph 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 decrease in the affordable amount, which affected the “100-150% Poverty” income group. The “100-150% Poverty (Other Years)” combines all years in our data *except* July 2007–June 2008. The spike in the control group (“200-300% Pov”) in December 2007 is due to the introduction of the mandate penalty (our other natural experiment), and we dummy out Dec 2007 to March 2008 for the control group so that this does not affect the regression estimates.

auto-enrollment. Figure A2 shows enrollment in the cheapest plan as a share of June 2008 enrollment for the 100-150% poverty treatment group and the 200-300% poverty control group as well as other years for the treatment group. There is a jump in treatment group enrollment in July 2007. (The jump in the control group in January 2008 is due to the introduction of the mandate penalty.)

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 (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 17.4% increase in the cheapest plan’s demand. Dividing by the \$18 reduction in its premium implies a semi-elasticity of 0.97%. 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 very close to the 1.17% semi-elasticity for the 150-200% poverty group found in the mandate penalty introduction experiment.

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	0.065** (0.005)	0.061** (0.008)	0.065** (0.009)	0.014** (0.005)	0.017* (0.007)	0.019* (0.008)
x Aug2007	0.053** (0.005)	0.032** (0.008)	0.036** (0.009)	0.029** (0.005)	0.016* (0.007)	0.022** (0.008)
x Sep2007	0.022** (0.005)	0.017* (0.008)	0.013* (0.008)	0.011 (0.005)	0.016* (0.006)	0.010 (0.007)
x Oct2007	0.063** (0.005)	0.054** (0.008)	0.057** (0.008)	0.008 (0.005)	0.009 (0.006)	0.010 (0.006)
Total	0.202*** (0.021)	0.157*** (0.031)	0.174*** (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	52	104	104
R-Squared	0.986	0.980	0.980	0.958	0.955	0.956

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.

Own-price Semi Elasticity

As described in Section 3.2, we use within plan variation in consumer premiums generated by the subsidy rules to estimate the semi-elasticity of demand with respect to own-price. Figure A3 shows the same changes as Figure 2 in the text, but for plans that increase their prices. Table 10 reports the coefficient on premium (α) from the regression in Equation (7):

$$\ln(NewEnroll_{j,y,r,t}) = \alpha \cdot P_{j,y,r,t}^{cons} + \xi_{j,r,t} + \xi_{j,r,y} + \epsilon_{j,y,r,t}.$$

The first column has no controls. Since price and quality tend to be positively correlated, it is not surprising that adding plan-region-income and plan-region-year controls (second column) increases the coefficient by about 50%. Since the regions in Massachusetts are different sizes, the last column weights by the region's average monthly new enrollment, generating our semi-elasticity estimate of 2.16% used in the calculations.




Figure A3: Market share around price increases

Note: This graph illustrates the source of identification for the estimation of the own-price semi-elasticity of demand. The graph shows average monthly plan market shares among new enrollees for plans that increased their prices at time 0. The identification comes from comparing demand changes for above-poverty price-paying (new) enrollees (for whom premium changes at time 0) versus below-poverty zero-price enrollees (for whom premiums are always \$0). The sample is limited to fiscal years 2008-2011, the years we use for demand estimation.

Table 10: Differential Premium Changes.

Effect on Price on Log New Enrollment			
	(1)	(2)	(3)
Premium	-0.0129*** (.00077)	-0.0193*** (.00168)	-0.0216*** (.00175)
Plan X region X income dummies		X	X
Plan X region X year dummies		X	X
Weight: Region avg enrollment			X
Observations	4713	4713	4713
R-Squared	0.238	0.891	0.882

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 a plan’s premium on the number of new enrollees for 2008 through 2011. Column (1) has no controls. Column (2) adds plan-region-year dummies and plan-region-income dummies. Enrollees below 100% of the federal poverty line pay zero premium for all plan, so they control for changes in plan quality over time. Other enrollees pay higher premiums for more expensive plans. Column (3) weights by the region’s average new enrollees each month. Standard errors are clustered at the region-income-year level.

C Structural Model Details and Additional Results

C.1 Moments

Let θ refer to all the parameters to be estimated. Given logit errors, and the systematic part of utility $\hat{u}_{ij} = u_{ij} - \epsilon_{ij}$, the plan choice probabilities are

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

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 expected share of consumers in that group who chose that plan given θ , to the observed share, s_j^{Obs} . 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. Analogous to the moments in F^1 , we match the observed share of uninsured to that predicted for the model. We do it separately by age-gender groups, income groups, regions, and year, but because the uninsured data are based on relatively small samples in the ACS, we do not interact these categories. The corresponding

moments are

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

To identify the different price-sensitivity parameters (α 's), we match the covariance of plan premium and individual attributes, following Berry et al. (2004). For each income or demographic group, h , we use

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

with $P_{i,0}^{Cons} = M_i$.

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 demand for the cheapest plan to the observed 22.2% change in demand for the mandate penalty introduction experiment:

$$G^3(\theta) = \sum_i ((1 + 22.2\%) Pr(j_i^* = j_{min} | Z_i, \nu_i, \theta, M_i^{Pre}) - Pr(j_i^* = j_{min} | Z_i, \nu_i, \theta, M_i^{Post})).$$

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.

The cost model we use is called a GLM with a Poisson error and a log link function; it is equivalent to a Poisson regression, although allowing for a continuous outcome variable. Mathematically, the model is estimated by MLE as if the data were generated by a Poisson process with mean equal to the expression in (8). However, in practice, this is simply a technique to estimate the parameters in (8), which describe the conditional mean of costs. These conditional means are what we use for costs in our simulations; we do not use the Poisson distribution. This method has the advantage over a log-linear regression of allowing for \$0 observations, which occur regularly in our data.

C.2 Parameter Estimates

To supplement the averages reported in the text, we report demand and cost coefficients for different groups. Table 11 shows the premium coefficients by age-gender group. Table 12 shows the average utility of uninsurance by age-gender group. Table 14 shows the average cost differences across income groups. Table 15 shows the cost differences across age and

gender groups.

Table 11: Interaction terms for Premium Coefficients

Baseline by income group (For omitted group: males 40-44)			Interacted with demographic group (To be added to baseline value)				
	Coef	S.E.		Male		Female	
				Coeff	S.E.	Coeff	S.E.
100-150% Pov	-0.046***	0.0054					
150-200% Pov	-0.029***	0.0061	Age 19-24	-0.0086*	0.0049	0.0011	0.0026
200-250% Pov	-0.027***	0.0072	Age 25-29	-0.0097**	0.0047	-0.0003	0.0033
250-300% Pov	-0.023**	0.0010	Age 30-34	-0.0012	0.0028	-0.0048	0.0042
			Age 35-39	-0.0016	0.0033	-0.0016	0.0039
			Age 40-44			0.0013	0.0035
			Age 45-49	0.0077**	0.0037	0.0006	0.0037
			Age 50-54	0.0050	0.0030	0.0013	0.0034
			Age 55-59	0.0122***	0.0037	0.0084**	0.0034
			Age 60-64	0.0111***	0.0038	0.0164***	0.0037

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Note: This table shows the the premium coefficient by income (left) and for each gender x age bin (right). The full premium coefficient for any sub-group is the sum of the coefficients for the corresponding income group and demographic group. Every income group has a premium coefficient < -0.02 , and no demographic group’s coefficient is $> .012$, so every group has a negative total premium coefficient.

C.3 Comparison of reduced form and structural distortion estimates

Our paper includes two main estimates of the distortion from price-linked subsidies: a “reduced form” estimate of \$36.35 based on the first-order approximation formula (reported in Section 3.3) and a “structural” estimate of \$23.38 from our simulation model for 2011 (reported in Section 4.3). In this appendix, we seek to separate out how much of the difference is due to (1) different estimated elasticities, (2) adverse selection and risk adjustment, (3) pass-through of costs into prices not equaling 1 (which is implicitly assumed by the first-order approximation), and (4) strategic responses from the other firms.

Table 13 breaks down the role of each of these factors. Line (a) shows our reduced form estimate of \$36.35, while the final line (e) is the headline structural estimate of \$23.38 for 2011. Bridging these two estimates are the following factors:

- **Different time period (line (b)):** The reduced form estimate uses elasticities from 2008 (when the natural experiment occurred), while the headline structural estimate is from 2011. Line (b) takes demand elasticities for 2011 simulated from our structural model and plugs these into the reduced form distortion formula. The result is a distortion of \$30.02 – bridging about half of the original gap. From this we conclude that the differing time periods is a major source of the gap.
- **Incorporating adverse selection (line (c)):** Our main distortion formula ignores adverse selection, but the structural model allows for it. To incorporate adverse selection, we use an augmented distortion formula derived in Appendix A.2 (Equation

Table 12: Relative Utility of Uninsurance

Baseline by income group (For omitted group: males 40-44)			Interacted with demographic group (To be added to baseline value)				
	Coef	S.E.		Male		Female	
				Coeff	S.E.	Coeff	S.E.
<100% Pov	0.952***	0.142					
100-150 % Pov	0.667***	0.174	Ages 19-24	0.236**	0.109	-0.596***	0.201
150-200 % Pov	0.150	0.319	Ages 25-29	0.095	0.117	-0.792***	0.240
200-250 % Pov	0.042	0.526	Ages 30-34	0.066	0.128	-0.889***	0.259
250-300 % Pov	0.517	0.825	Ages 35-39	-0.067	0.162	-0.844***	0.251
			Ages 40-44			-1.073***	0.324
			Ages 45-49	-0.043	0.135	-1.097***	0.329
			Ages 50-54	-0.229	0.156	-0.874***	0.251
			Ages 55-59	-0.299*	0.161	-0.778***	0.256
			Ages 60-64	-0.590**	0.253	-0.785***	0.298

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Note: This table shows the the relative utility of uninsurance by income (left) and for each gender x age bin (right). The total average utility of uninsurance (relative to the average CommCare plan) for a group is the sum of the coefficients for the corresponding income group and demographic group. Some groups have a positive utility of uninsurance (negative value of insurance), which is necessary to explain the fact that there are many people for whom insurance is free who nonetheless choose not to purchase insurance.

(A.1)). Simulating the necessary additional statistics using our structural model and plugging these into the formula, we get a distortion of \$26.71 (bridging another fourth of the original gap).

- **Moving to a structural simulation (no strategic interactions) (line (d)):** This line implements a structural simulation of the distortion for CeltiCare in 2011 that shuts down strategic pricing responses by other firms. The difference is accounted for by the fact that the sufficient stats formula (which, e.g., implicitly assumes constant semi-elasticities of demand and a pass-through rate of 1). The estimate in line (d) is \$24.34, closing another 17% of the original gap.
- **Adding strategic interactions:** Line (e) shows our final structural simulation that allows for strategic pricing responses by competitors. The difference between (d) and (e), however, is tiny (just \$0.51), indicating that strategic effects explain very little of the original gap.

C.4 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 16 shows estimated costs as a fraction of revenue for each firm for each policy simulation. The only MLR that is below the 80% threshold is BMC's under the maximum-price equilibrium for price-linked subsidies in 2009.

Table 13: Comparison of reduced form and structural estimates

Calculation method		Estimate (\$)
Reduced-form formula		
(a)	Baseline	36.35
(b)	With structural model elasticities	30.02
(c)	+ Adverse selection	26.71
Structural model		
(d)	Without strategic interactions	24.34
(e)	Full model	23.83

Table 14: Cost by income group

	Income Group			
	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: This table reports the costs for each income group relative to the below poverty group. The reported percentage for group g is $\exp(\mu_g) - 1$, when the μ_g for the ‘Less than Poverty’ group is set to zero.

Table 15: Costs by Age and Gender

	Male			Female		
	Poisson Coefficient	S.E.	Percent	Poisson Coefficient	S.E.	Percent
Ages 19-24				0.312	0.00031	36.6%
Ages 25-29	0.336	0.00032	39.9%	0.488	0.00028	62.9%
Ages 30-34	0.424	0.00035	52.8%	0.558	0.00031	74.7%
Ages 35-39	0.526	0.00035	69.2%	0.609	0.00036	83.9%
Ages 40-44	0.632	0.00033	88.1%	0.752	0.00036	112.1%
Ages 45-49	0.807	0.00031	124.1%	0.850	0.00032	134.0%
Ages 50-54	0.943	0.00030	156.8%	0.886	0.00030	142.5%
Ages 55-59	1.065	0.000302	190.1%	0.907	0.00029	147.7%
Ages 60-64	1.193	0.00031	229.7%	1.073	0.00030	192.4%

All coefficients are significant at the $p < .01$ level.

Note: This table shows cost coefficients by age and gender. Males 19-24 are the omitted group. For each gender, the first two columns are the poisson coefficient and its standard error. The third coefficient converts this to a percent difference, $\exp(\mu) - 1$.

This suggests MLR rules may have had some bite for avoiding the more expensive equilibrium in 2009, but would not have limited the distortion 2011.

Table 16: 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 ($100 \cdot \text{costs}/\text{revenue}$) 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 MLR rules become more binding. Table 17 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 and for fixed subsidies when it competes against NHP (the highest cost insurer). Network Health and BMC also have MLRs below 80% when competing against NHP. MLR rules are more relevant for extremely uncompetitive markets.

Table 17: 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 ($100 \cdot \text{costs}/\text{revenue}$) for the insurer of a given row, when competing only against the insurer of a given column. The first number in each cell is the MLR under price-linked subsidies; the number in parentheses is the MLR under fixed subsidies. The plans are listed in order of increasing costs: CeltiCare, Network Health, BMC, NHP. 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.

D Uncertainty simulations

To model the potential benefits of price-linked subsidies, we consider the problem of a regulator setting subsidies for an upcoming year. We assume that regulators observe lagged demand and cost data from insurers (e.g., from state insurance department rate filings),

which give them all the relevant information for the optimal subsidy except for a *market-level* shock to cost growth.³² This proportional shock applies to all plans' costs and the cost of charity care, multiplying them by a factor of $(1 + \Delta)$. Fixed subsidies have to be set in advance, prior to the realization of health care costs, but the regulator would like to set a higher subsidy when costs are higher because of the externalities of charity care and the desire to insure consumers against price variations (see discussion in Section 1.2). If insurers observe cost shocks before setting prices, then prices contain information about optimal subsidies; price-linked subsidies can make use of this information.³³

We assume the externality that is avoided when an individual chooses to purchase insurance has two components: charity care costs and the paternalistic component of the externality. Charity care costs are, unfortunately, quite difficult to measure. Instead of attempting to estimate them in our Massachusetts setting, we assume they are a proportion λ of an individual's expected costs in the average plan. For a cost-shock Δ ,

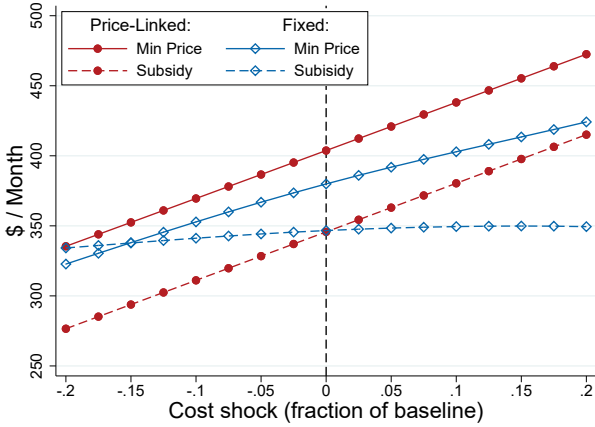
$$C_i^{Charity} = \lambda \cdot \underbrace{\exp(\hat{\mu}X_i + \psi_0)}_{\text{Baseline cost in avg. plan}} \cdot (1 + \Delta) \quad \lambda = \{.6, .8, 1\}.$$

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, implying $\lambda = .6$. The best case for price-linked subsidies, which we report in the text is $\lambda = 1$. We also show results for $\lambda = .8$ as an intermediate point.


We assume the regulator sets (income-specific) fixed subsidies to be optimal if costs are at the expected level (i.e., zero cost shock) and sets affordable amounts for price-linked subsidies to replicate the fixed subsidy amounts at $\Delta = 0$. In practice, we implement this by calibrating the paternalistic externality, E_0 as a residual to rationalize the ACA's affordable amounts (and therefore subsidy levels) as optimal at $\Delta = 0$. 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 largely 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. This result is consistent with the finding of Finkelstein et al. (2017) that demand is less than marginal cost for most CommCare enrollees. With $E_0 = 0$ 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.

³²It would also be possible to model insurer-specific cost shocks. However, market-level cost shocks are what matters for the welfare-relevant correlation between insurer costs and the externality of uncompensated care.

³³We assume that consumers (like the regulator) do not observe the cost shock, so their utility of insurance (as reflected in ξ_j and β) is held constant. Even 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.



(a) Cheapest prices and subsidies



(b) Share insured

Figure A4: Equilibrium under cost shocks

Note: These graphs show the equilibrium under price-linked and fixed subsidies for cost shocks of -20% to +20% of baseline. The left graph shows the price of the cheapest plan and the subsidy level. The right graph shows the share of consumers who purchase insurance.

Prices and Coverage Rates

Figure A4a shows the price of the cheapest plan and the average subsidy for 2011 across values of the cost shock (on the x-axis) for each subsidy policy. Unsurprisingly, prices (shown in solid lines) increase with costs under both policies. They increase slightly faster under price-linked subsidies.³⁴ 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).³⁵

Figure A4b 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 since they 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 70%) with a negative cost shock and lower (down to 19%) with a positive shock. Of course, it is not clear whether price-linked subsidies' stabilization of coverage in the face of cost shocks is desirable: it may be optimal for fewer people to buy insurance when it is more expensive.




Figure A5: Public surplus under cost shocks in 2011

Note: These graphs show public surplus (in dollars per month per eligible member) under price-linked and fixed subsidies for cost shocks of -20% to +20% of baseline. Each graph corresponds to a different assumption about how much the externality changes with costs. The dashed and dotted lines subtract a cost of pricing risk to consumers (using the method in Equation (6)) with coefficients of relative risk aversion of 2 and 5, respectively.

Welfare

The three panels of Figure A5 show how public surplus changes with cost shocks under different assumptions on λ , which captures how the externality scales with costs. (The bottom panel is the same as Figure 3 in the text.) Fixed subsidies (in blue) are preferred at zero cost shock in all cases.³⁶ For non-zero cost shocks, the gap between fixed and price-linked subsidies narrows, 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.

We then adjust the public surplus measure to include a cost of price variation to consumers, based on Equation (6). This adjustment depends on the level of consumption, which we approximate using monthly income at the median of a consumer’s income bin, minus the average premium (or mandate penalty) they pay. For (ΔP) , we use the premium variation for the average plan relative to its premium at $\Delta = 0$. The cost of price variation also depends on the coefficient of relative risk aversion, γ . Chetty (2006) argues that $\gamma \leq 2$; to consider the case most favorable to price-linked subsidies, we use $\gamma = 2$ (dashed lines). We also consider a more extreme case of $\gamma = 5$ (dotted lines), 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). For both values of γ , the welfare losses from greater price risk to consumers are fairly small and do not substantially change our results – i.e., the dashed and dotted lines in A5 are not very different from the solid lines. This is explained by the well-known fact that for reasonable risk aversion, utility is locally quite close to linear (Rabin, 2000). The maximum premium variation we consider (about +/- \$75 per month) represents only about +/- 5% of average income, even for this low-income population. Based on this analysis, it seems that 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.

Results for 2009

Figure A6 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

³⁴On average a 20% cost shock is about \$75, so prices increase a little less than one-for-one with costs.

³⁵We note that for the largest negative cost shocks that we simulate, prices under fixed subsidies fall slightly below the subsidy amount. In these cases, our calculations assume that insurers can charge “negative premiums” – i.e., rebate the difference to consumers. The ACA does not allow negative premiums, but implementing these does not seem infeasible. Rebates to consumers could be accomplished either via direct payments or via tax rebates (similar to the way the ACA’s mandate penalty works). If rebates were not allowed, then fixed subsidies could also distort pricing incentives for highly negative cost shocks, since insurers would have no incentive to lower prices below the subsidy amount. This is not a major issue for the range of shocks we consider (where average subsidies are always less than the minimum price), but it could become important for even larger negative shocks.

³⁶With the differing values of λ , the calibrated fixed externality E_0 differs, so the externality of inframarginal consumers is not the same under the 3 scenarios, which is why 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.

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 – typically a shock of between 5-10%.

Including insurer profits in welfare

Regulators may also include profits in their objective function. We repeat the comparison of fixed and price-linked subsidies under uncertainty for this more inclusive welfare definition. Doing so requires recalibrating the fixed component of the externality of uninsurance, E_0 , since it is calculated based on the assumption that equilibrium subsidies are optimal and what is optimal depends on the objective function. In practice, we find 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, the uncertainty analysis only makes sense if baseline subsidies are optimal at zero cost shock, so we proceed with the estimated values.

Figure A7 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 $\geq -10\%$.

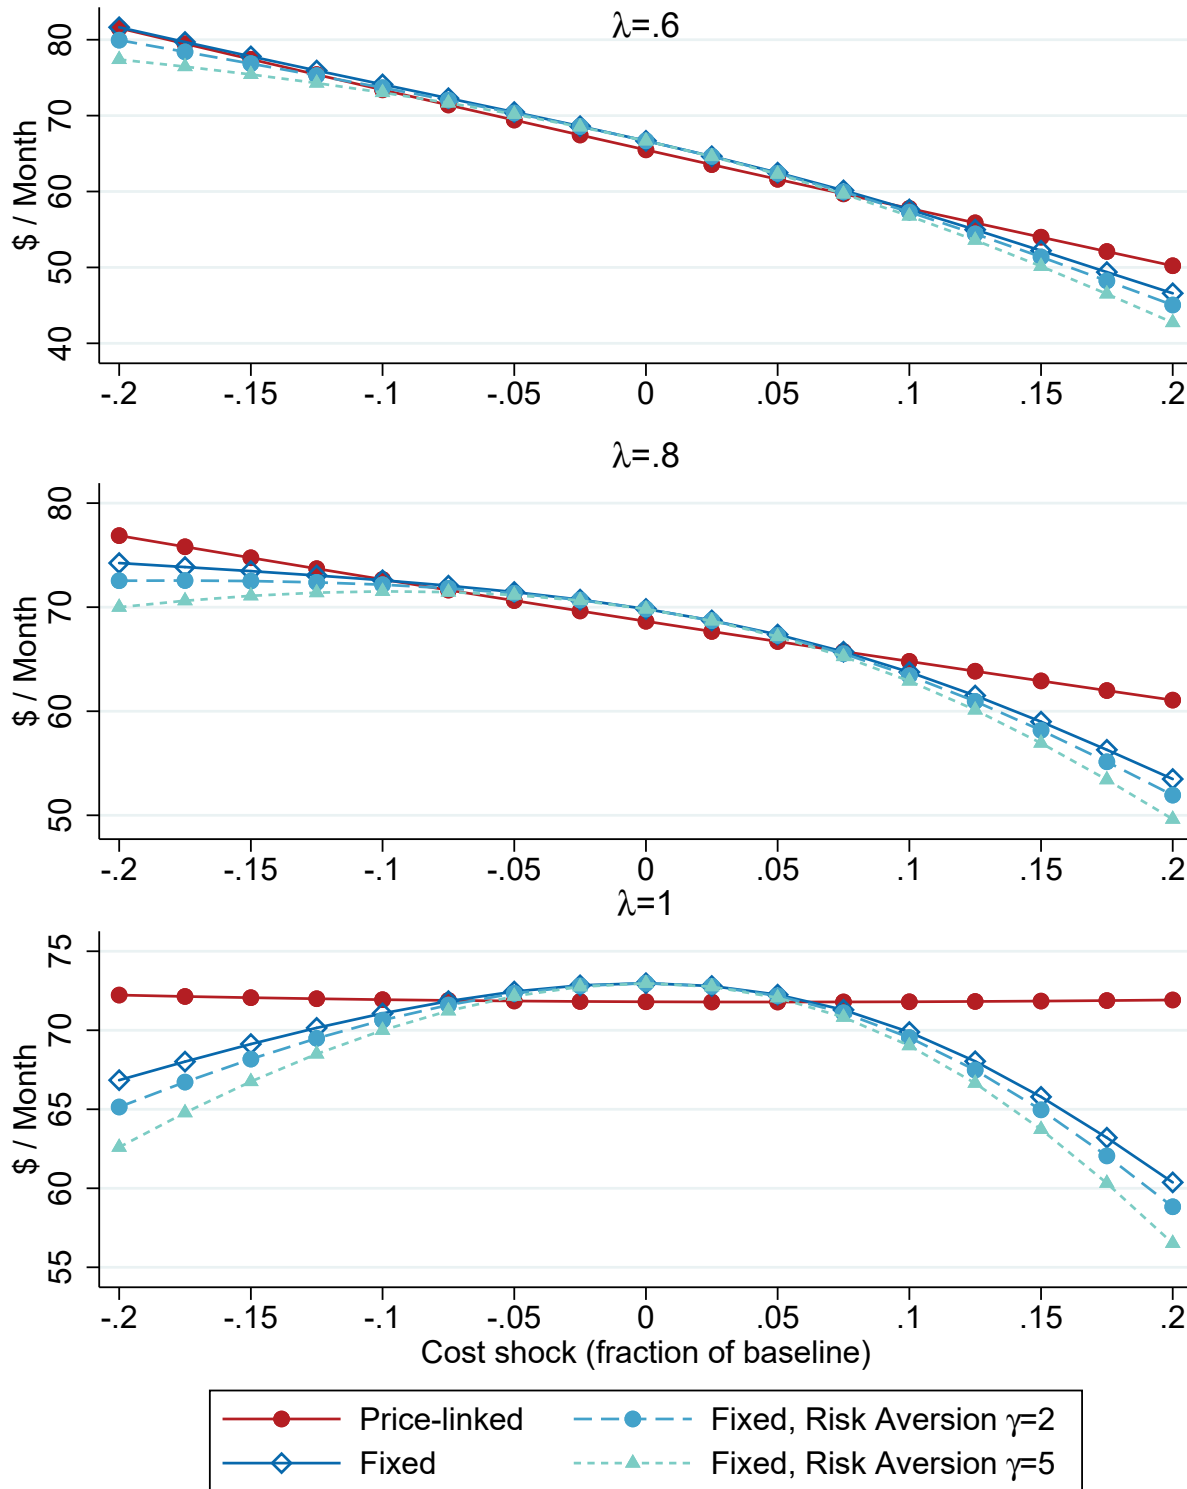


Figure A6: Public surplus under cost shocks in 2009

Note: These graphs show public surplus under price-linked and fixed subsidies for cost shocks of -20% to +20% 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.

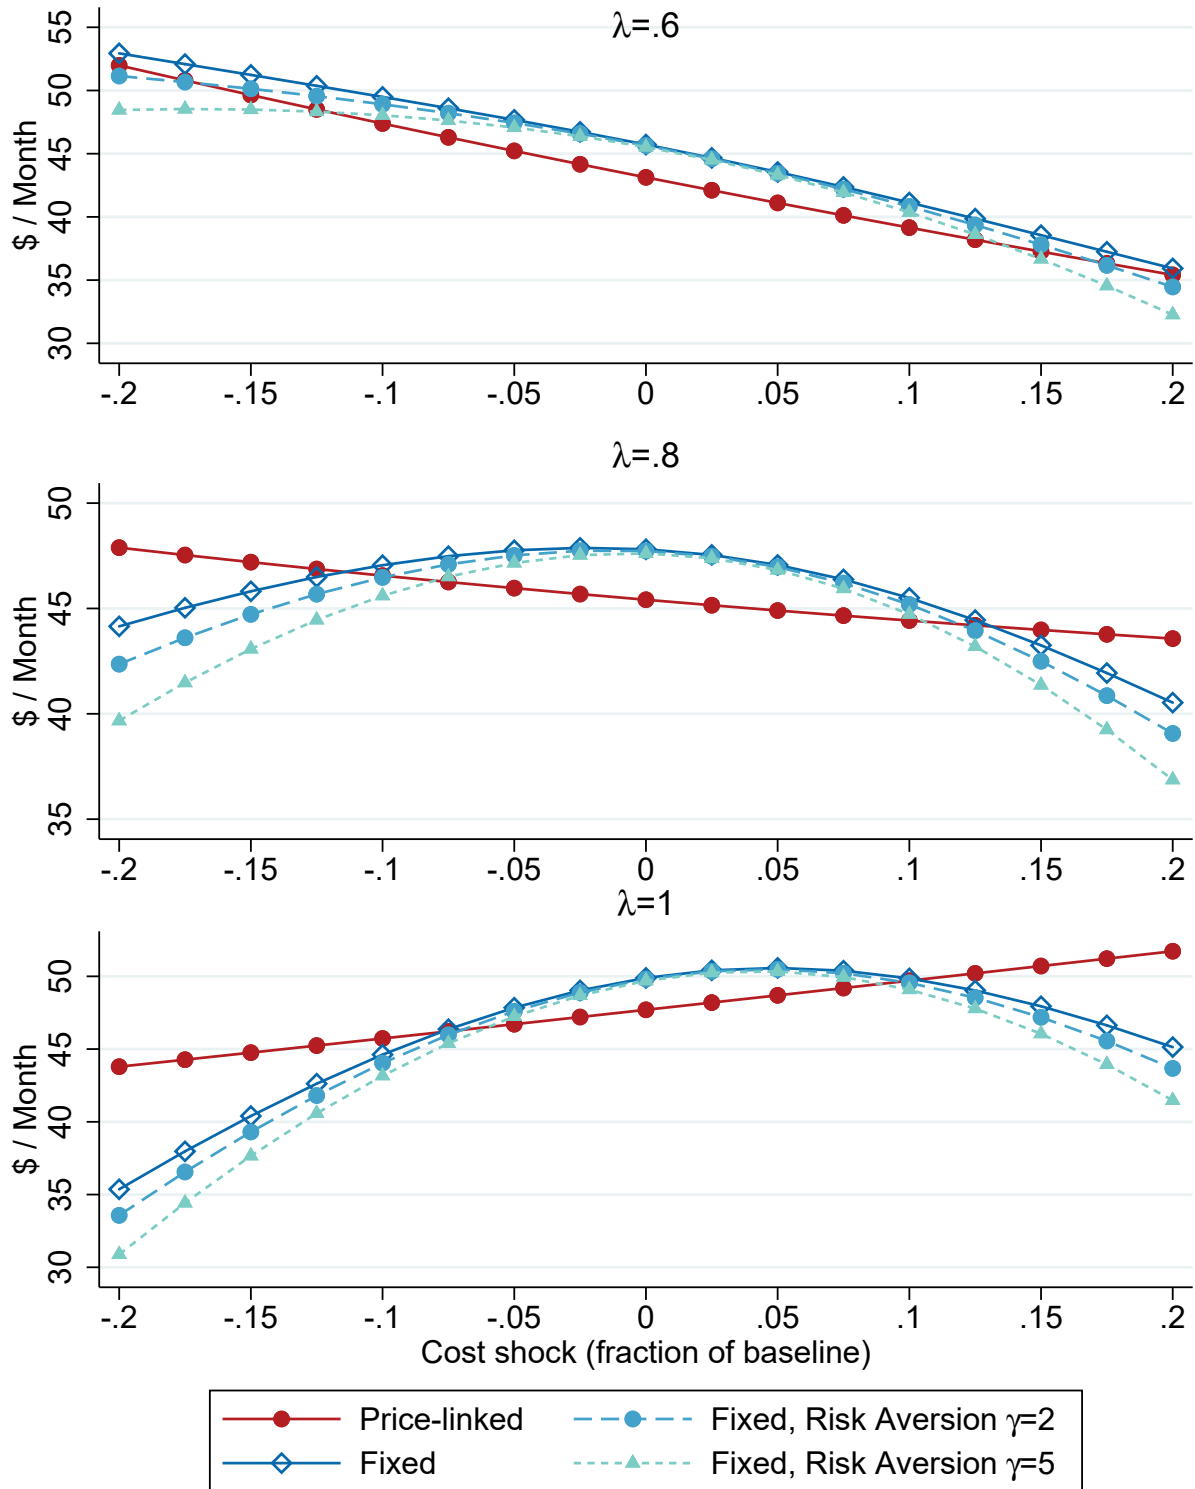


Figure A7: Public surplus plus profits under cost shocks in 2011

Note: These graphs show the sum of public surplus and profits under price-linked and fixed subsidies for cost shocks of -20% to +20% 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.